Packing Perfect Matchings in Random Graphs at the Optimal Threshold
UROP+ Final Paper, Summer 2019

Maxwell Fishelson
maxfish@mit.edu
Mentor: Vishesh Jain
Project suggested by Vishesh Jain

Abstract

The celebrated resolution of Shamir's conjecture by Johannson, Kahn, and Vu [2] shows that, with high probability, perfect matchings exist in the binomial k-uniform hypergraph $H^k_{n,p}$ for $p = \omega_k(\log n/n^{k-1})$. Here, we consider the problem of approximately decomposing $H^k_{n,p}$ into edge-disjoint perfect matchings. Using a novel method of generating random hypergraphs, known as 'online sprinkling', it was shown by Ferber and Vu that this can be done with high probability, provided that $p = \omega(\log^{4k}(n)/n^{k-1})$. By refining the Ferber-Vu analysis, we show that, given a conjectural local resilience version of the Johannson-Kahn-Vu result, this can be done at the asymptotically optimal threshold $p = \omega_k(\log n/n^{k-1})$.

1 Introduction

In 2008, Johannson, Kahn, and Vu [2] resolved the long-standing so-called Shamir's problem in probabilistic combinatorics by determining the asymptotic threshold $p(n)$ for which the binomial random k-uniform hypergraph $H^k_{n,p(n)}$ has a perfect matching with high probability (where we take n to be a multiple of k). A standard 'coupon-collector' argument shows that unless $p = \Omega(\log n/n^{k-1})$, $H^k_{n,p}$ has isolated vertices, and hence, cannot contain a perfect matching. Using a highly ingenious argument, they showed that this lower bound is essentially tight i.e. for $p \geq C(k)\log n/n^{k-1}$, $H^k_{n,p}$ contains a perfect matching with high probability.

Given this result, and motivated by the long line of work in graph decompositions (see, e.g., the survey [7] and the references therein), it is natural to ask when $H^k_{n,p}$ can approximately be decomposed into edge-disjoint perfect matchings. This question was first considered by Frieze and Krivelevich [8], who showed that such a decomposition exists with high probability for $p = \omega_k(\log^2 n/n)$. Subsequently, using a novel method for generating binomial hypergraphs known as 'online sprinkling', Ferber and Vu [3] considerably improved this bound to $p = \omega_k(\log^{5k} n/n^{k-1})$, which is optimal up to the exponent of the logarithmic factor. More specifically in this paper, we focus our study on the random k-partite hypergraph model $H^k_{n\times k,p}$. It is the random hypergraph on the vertex set consisting of k parts with n vertices each, where each possible edge (formed by selecting 1 vertex from each part) is included in the hypergraph independently with probability p. On this more specific model, Ferber and Vu's arguments give a bound of $p = \omega_k(\log^{4k} n/n^{k-1})$.

The main open problem in the paper of Ferber and Vu is whether the factor $\log^{4k} n$ can be replaced by $\log^C n$ for some constant C independent of k (ideally $C = 1$). Here, by carrying out a variation of the Ferber-Vu analysis, we show that this can indeed be done, provided that one has a 'local resilience' version of the Johansson-Kahn-Vu result, which we now discuss.

The study of the local resilience of random (hyper)graph properties was initiated by Sudakov and Vu in 2007 [9], and is centered around the following question: given a monotone increasing (hyper)graph property $\mathcal{P}$ and a probability $p(n)$, determine the largest value r such that, almost surely, $H^k_{n,p(n)}$ possesses property $\mathcal{P}$ even after removing an arbitrary hypergraph of maximum vertex-degree r. Note that this is equivalent to determining the largest value of r such that for almost all samples of $H^k_{n,p}$, an adversary cannot 'destroy' the property $\mathcal{P}$ by removing at most r edges incident to each vertex. Whereas a lot of local resilience results have been proved for random graphs for the property of containing a specified large substructure (see, e.g., the survey [10] and the references therein), much less is known for the case of large substructures in random hypergraphs. In particular, at the time of writing, we are not aware of *any* local resilience results in random hypergraphs at the 'correct threshold' i.e. when $p(n)$ is the threshold for the appearance of a single copy of the large substructure.

Here, we conjecture that a local resilience version of the Johansson-Kahn-Vu result should hold for $p = \omega_k(\log n/n^{k-1})$ and $r = \alpha n^{k-1}p$, for some sufficiently small constant $\alpha > 0$.

Conjecture 1.1. *Fix $k \geq 2$. Then there exists a sufficiently large constant $K > 0$ and a sufficiently small constant $\alpha > 0$ such that, if $p = \omega_k(\log n/n^{k-1})$, then*

$$\lim_{n\to\infty} \Pr(H^k_{n\times k,p} \setminus E(H^*) \textit{ has a 1-factor}) = 1$$

for all H^*, where H^* is a hypergraph with maximum degree $r = \alpha n^{k-1}p$ that is in terms of the sampled $H^k_{n\times k,p}$.

The evidence in support of this conjecture is the widespread belief that the 'only' obstruction to the presence of perfect matchings in random hypergraphs is the presence of isolated vertices; indeed, this is formalized in very recent (and as of now unpublished) work of Kahn, who shows that when we generate a random hypergraph sequentially by adding edges uniformly at random from among all the non-edges, then almost surely, the first time a perfect matching appears coincides with the first time when there are no isolated vertices. Given this intuition, along with the fact that deleting only $\alpha n^{k-1}p$ edges incident to each vertex in $H^k_{n,p}$ cannot create isolated vertices (as each vertex has expected degree $n^{k-1}p$), we are led to believe that the above conjecture should hold.

We do not go through Ferber and Vu's arguments fully, as the most wasteful part of their argument occurs at the end of their proof. In order to emphasize our contribution to the problem, we begin, in Section 2, by outlining the first half of their proof and then delving more deeply into the second half of the proof, where we believe there is room for improvement that will enable us to close the gap on the 1-factor packing threshold. In Section 3, we prove a generalized version of the Johannson-Kahn-Vu solution that is necessary for our arguments in Section 2. Lastly, in Section 4, we prove a collection of lemmas necessary for our arguments in Section 3.

2 Recycled Sprinkles

Fundamental to Ferber and Vu's arguments is the technique of "online sprinkling" (sometimes referred to as "multiple exposures"). It is a means of exposing some information about a random event but not all the information, leaving some randomness to be exposed later. More specifically, if we have $1 - p = \prod_{i=1}^{\ell}(1 - p_i)$ for some p_i, then we see $H^k_{n,p}$ is distributed identically to $H^k_{n,p_1} \cup H^k_{n,p_2} \cup \cdots \cup H^k_{n,p_\ell}$. In this way, we can sample $H^k_{n,p}$ incrementally, first sampling H^k_{n,p_1} revealing some information about the edges of the random hypergraph while leaving other parts of the hypergraph unobserved. In our argumentation here, the p_i's we will use will be sufficiently small such that, with high probability, no edge will be sampled in more than one H^k_{n,p_i}, and so the contribution each H^k_{n,p_i} has on the $H^k_{n,p}$ we sample will be totally independent and we will not have to condition on the parts of $H^k_{n,p}$ we've already exposed.

As stated previously, we focus our study on the random k-partite hypergraph model $H^k_{n \times k,p}$. Of the n^k possible edges in $H^k_{n \times k,p}$, we expect to have pn^k, and since each perfect matching, or 1-factor, contains n edges, we wish to find $N = (1 - o(1))pn^{k-1}$ edge disjoint perfect matchings to achieve the desired packing. We outline Ferber and Vu's arguments here so we can focus on the piece of the argument that we believe to be the weakest and the area where improvement could lead to the closing of the probability gap asymptotically.

This 1-factor packing is generated via an incremental construction that takes place over two phases. In the first phase, we generate N almost perfect matchings: each matching only containing $(1 - \alpha)n$ edges for a small constant α. Then, in the second phase we complete these almost perfect matchings into perfect matchings. This idea of first constructing an almost spanning structure before completing it comes from the "nibbling" idea introduced by Ajtai-Komlos-Szemerédi [4] and Rödl [5].

Phase 1 is broken into N rounds, and each round is further divided into steps. During round each i, we construct an almost perfect matching. We start with an empty matching M_{i0}, and then at each step j, we generate the matching M_{ij} from the current matching $M_{i(j-1)}$ by adding a collection of vertex-disjoint edges via a very low probability sprinkling. After $\ell = O(\log \log n)$ steps, we will have generated our almost perfect matching $M_{i\ell} = M_i$ containing $(1 - \alpha)n$ edges. The key here is to show that the sprinkling over the course of this phase does not leave any edge with too much weight. That is, if $p_1, \cdots, p_s$ denote all the edge-probabilities used during this phase in order to "expose" a particular k-tuple e, we wish to show that $1 - \prod_{i=1}^{s}(1 - p_i) \leq (1 - \epsilon)p$, leaving at least a constant factor of the total exposure probability p for later. Then, we show that if p is not too large (i.e. just above the

$\Theta\left(\frac{\log n}{n^{k-1}}\right)$ 1-factor existence threshold), this procedure gives us edge-disjoint almost perfect matchings whp, that is it is very unlikely for any edge to be sprinkled twice in two separate rounds. Later, we deal with the case of p being much larger than the threshold, which becomes negligible due to a monotonicity result.

For each almost perfect matching M_i, we let U_i be the "residual" of round i, or the set of αkn vertices not yet matched. Phase 2 constitutes matching these residual vertices in an edge-disjoint way. The technique from Phase 1 can't simply be used to complete the matchings since, as the set of remaining vertices for a matching get smaller and smaller, we must use exposures with increasingly large probability, which reaches an exposure total surpassing the p threshold. Another technique must be used to complete these matchings, but the approach originally used by Ferber and Vu was somewhat wasteful. For each residual, they simply expose each k-tuple of vertices with probability $\frac{\log(\alpha n)}{(\alpha n)^{k-1}} = (1+o(1))\frac{\log n}{(\alpha n)^{k-1}}$. By Johannson-Kahn-Vu, this guarantees the existence of a 1-factor in this residual and therefore a way to complete this perfect matching.

Importantly, in how we execute the first phase, the residual vertices will be distributed uniformly at random. That is, the αn vertices belonging to the residual in each of the k parts will be uniformly distributed. In this way, the probability that a single k-tuple belongs to a residual is α^k and this phase 2 sprinkling contributes a probability of

$$\alpha^k \cdot \frac{\log n}{(\alpha n)^{k-1}} = \frac{\alpha \log n}{n^{k-1}}$$

to each k-tuple. Over the course of $N = (1-o(1))pn^{k-1}$ rounds, the total porbability sprinkled during this phase is $\alpha \log(n) p$, necessitating $\alpha = o(\log^{-1} n)$ in order to have the sprinkled probability not surpass p. Using an α this small necessitates larger probability exposures in the first phase, forcing us to use a larger p. Namely, the smallest p achievable with this strategy is $\Theta\left(\frac{\log^{4k} n}{n^{k-1}}\right)$, giving a polylog factor (in terms of k) worse than the 1-factor existence threshold lower bound.

However, theoretically there is room to optimize this second phase; an optimization that will enable us to remove the $\log^{4k}$ factor, lower the upper bound on the 1-factor packing existence threshold to $p = \frac{K \log n}{n^{k-1}}$ with $K = O(1)$, and close the gap asymptotically on 1-factor packings. Specifically, each k-tuple of vertices belongs to $\alpha^k N = O(\log n)$ residuals in expectation, and so it gets sprinkled in phase 2 many times. However, intuitively, if a k-tuple is sprinkled in a previous residual but not used in the matching on that residual, we shouldn't have to sprinkle it again if it belongs to another residual and we should be able to simple reuse (or recycle) that previous sprinkling. The benefit of this is that, before phase 1 even begins, we can sprinkle every k-tuple in the graph with some probability q where $q < p$ and $q > \frac{\log n}{(\alpha n)^{k-1}}$. As long as we choose an $\alpha < \epsilon$ with $\alpha = O(1)$, this will leave enough probability to execute phase 1 without issue. Now, by Johannson-Kahn-Vu, any possible residual (any collection of vertices containing αn vertices from each part) will contain a perfect matching whp. So, when we get to phase 2, we can form matchings from these q sprinkled edges for each residual, as long as we do so in an edge disjoint way.

This edge disjointness condition does lead to some difficulties, though. Now that a q sprinkled edge cannot be used if it was used in a previous residual's matching, we cannot look at the q sprinkled edges within a residual in isolation and simply couple it to a random hypergraph with probability q. Our new model is a random hypergraph with edge probability q where, after the edges are sampled, certain edges are deemed unusuable in the 1-factor (or are equivalently deleted). These deleted edges are in terms of the sampled edges, and so if we are to prove any claims about this random hypergraph with edge deletions, we must assume that these deletions are performed adversarially. Any claims about the deleted edges would place a conditional probability on the random sampling, and so we must argue that any way to delete edges (that corresponds to edges belonging to 1-factors in previous residuals) will not destroy the existence of a 1-factor.

In this sense, we cannot directly apply the Johannson-Kahn-Vu (JKV) result, as they consider a random hypergraph model and now we seek to prove a result on a random hypergraph with adversarial edge deletions. However, there is hope that such a generalization of JKV is possible. The threshold edge probability for the existence of a 1-factor: $p = \frac{O(\log n)}{n^{k-1}}$ is also the threshold probability for our random hypergraph to not contain any isolated vertices. Having no isolated vertices is a necessary, but not sufficient condition for the existence of a 1-factor. One could imagine enumerating a collection of conditions, each individually necessary for the existence of a 1-factor, and together sufficient. As we increase p from 0 to 1, we cross the threshold probability for each of these conditions at some point, and the final condition satisfied is the isolated vertex condition (as evidenced by the fact that it's threshold probability is that of the 1-factor existence). Thus, as long as our adversarial deletions do not produce isolated vertices, it seems plausible that the arguments of JKV still hold.

In fact, we will see that the deletions will not produce any isolated vertices. Since the vertices of a residual are uniformly random and independent of previous residuals, each vertex belongs to any given residual with probability α, belongs to αN total residuals in expectation, and by Chernoff, belongs to at most $2\alpha N$ residuals total whp. Since the deleted edges correspond to perfect matchings in previous residuals, no vertex can lose more than one edge per previous residual. So, for the qn^{k-1} edges containing a vertex v created by this sprinkling, at most $2\alpha N = 2\alpha pn^{k-1}$ are destroyed whp.

For a new residual with vertex v, the edges containing v in that residual are uniformly random, and so at most a $\frac{2\alpha p}{q}$ portion of them are deleted whp. So, choosing a $q > 2\alpha p$ will make it so a vertex's degree can only decrease by a constant factor due to the deletions and therefore produce no new isolated vertices. For asufficiently small $\alpha = O(1)$, we can still achieve the necessary $q < p$. Thus, given a resilience version of JKV, we can show that this alternative strategy for Phase 2 succeeds whp, and only using $\Theta(\log n/n^{k-1})$ total probability sprinkling as desired. In the next section, we will attempt to show JKV's arguments are generalizable to this adversarial deletions setting, as well as rehash the arguments in a more intuitive way than they were originally presented.

3 Johansson-Kahn-Vu and the Adversarial Deletions Conjecture

In this section, we present a proof of the JKV result and we look at the short comings of the argument that prevent us from extending the result to the adversarial context.

First, we formalize some notation on random k-partite hypergraphs. $H^k_{n\times k,p}$ and $H^k_{n\times k,m}$ are both k-uniform hypergraphs with vertex sets consisting of k parts with n vertices each. $H^k_{n\times k,p}$ is the standard Erdős-Rényi random hypergraph where each possible edge (formed by selecting 1 vertex from each part) is included in the hypergraph independently with probability p. $H^k_{n\times k,m}$, on the other hand, is the random graph in which a random set of m of the n^k possible edges is included uniformly at random. The expected number of edges in $H^k_{n\times k,p}$ is pn^k. We will see that when $m = pn^k$, these two random hypergraphs behave similarly, somewhat interchangeably. This enables us to prove results on $H^k_{n\times k,m}$ using the $H^k_{n\times k,p}$ model. We are concerned with the inflection point $p = \frac{K\log n}{n^{k-1}}$, and so we investigate $m = Kn\log n = pn^k$. The original statement of JKV is

Theorem 3.1. *Johannson,Kahn, Vu [2]*
Fix $k \geq 2$. Then there exists a sufficiently large constant $K > 0$ such that, if $m \geq$

$Kn\log n$, then

$$\lim_{n\to\infty} \Pr(H^k_{n\times k,m} \textit{ has a 1-factor}) = 1$$

We applied JKV on the $H^k_{n\times k,p}$ random graph model but we prove a theorem on the $H^k_{n\times k,m}$ model. This will be sufficient as

$$\lim_{n\to\infty} \Pr\left(|E(H^k_{n\times k,p})| \geq \frac{1}{2}np^k\right) = 1$$

by Chernoff. And so, for $m \leq \frac{1}{2}pn^k$,

$$\begin{aligned}\Pr(H^k_{n\times k,m} \text{ has a 1-factor}) &\leq \Pr\left(H^k_{n\times k,p} \text{ has a 1-factor}\middle| |E(H^k_{n\times k,p})| \geq \frac{1}{2}np^k\right) \\ &\leq \frac{\Pr(H^k_{n\times k,p} \text{ has a 1-factor})}{\Pr(|E(H^k_{n\times k,p})| \geq \frac{1}{2}np^k)}\end{aligned}$$

which gives

$$\lim_{n\to\infty} \Pr(H^k_{n\times k,m} \text{ has a 1-factor}) \leq \lim_{n\to\infty} \Pr(H^k_{n\times k,p} \text{ has a 1-factor})$$

Our conjecture on adversarial deletions can be stated as follows

Conjecture 3.2. *Fix $k \geq 2$. Then there exists a sufficiently large constant $K > 0$ and a sufficiently small constant $\alpha > 0$ such that, if $m \geq Kn\log n$, then*

$$\lim_{n\to\infty} \Pr(H^k_{n\times k,m} \setminus E(H^*) \textit{ has a 1-factor}) = 1$$

for all H^*, where H^* is a k-partite hypergraph with maximum degree $2\alpha\frac{m}{n}$ that is in terms of the sampled $H^k_{n\times k,m}$.

The edges of this H^* hypergraph that we are removing represent the adversarial deletions. Our random k-uniform hypergraph $H^k_{n\times k,m}$ on nk vertices with m edges has average degree $\frac{k\cdot m}{nk} = \frac{m}{n}$, and so subtracting the edges of H^* from $H^k_{n\times k,m}$ is equivalent to deleting at most a 2α fraction of the expected number of edges incident to each vertex. Since this theorem holds for every possible H^*, it means that no matter how an adversary chooses to delete edges, given that he can delete at most $2\alpha\frac{m}{n}$ edges belonging to any one vertex, he will be unable to destroy all the 1-factors whp.

Our arguments to prove this theorem and investigate this conjecture closely parallel those made by Bal and Frieze in their paper "The Johansson-Kahn-Vu solution of the Shamir problem" [1] and those presented in the textbook "Introduction to

Random Graphs" by Frieze and Karoński [6]. While the Shamir problem is solved in the original paper by Johannson Kahn and Vu, they prove a much more general result, and this paper by Bal and Frieze uses tools specifically tailored to the Shamir problem that simplify the arguments, though the revolutionary techniques from the original paper are still the bedrock of the argument.

The main technique is a backwards analysis. Fundamental to the JKV solution is to view the sampling of the random hypergraph $H^k_{n\times k,m}$ as an incremental deletion (or should I say decremental; should downward escalators be called de-escalators?). We start with the complete k-uniform hypergraph $H_0 := K^k_{n\times k}$ and then produce H_1 by deleting an edge from H_0 uniformly at random. Then, we produce H_2 by deleting a random edge from H_1, and repeat this process until we produce H_T where $T = n^k - m$. We see that H_i is identically distributed to $H^k_{n\times k,n^k-i}$ as the order of edge deletions is irrelevant. This is where the backwards analysis stops in the original argument.

In our pursuit to prove the conjecture, we continue to delete edges to form $H_{T+1}, H_{T+2}, \cdots, H_{T+s}$. These additional deletions are not random as the first T deletions were, but rather, are performed in an adversarial fashion. They are the edges in H^* that belong to H_T and the only claim we can make about the edges deleted here is that no $2\alpha\frac{m}{n}$ of them share the same vertex. Nonetheless, we want to show H_{T+s} has a 1-factor whp. The idea is that, no single edge deletion can affect the number of 1-factors appreciably, and even over the course of $T + s$ deletions, many 1-factors must remain at the end.

At each step in the backwards analysis i, we keep track of 3 events, $\mathcal{R}_i, \mathcal{B}_i, \mathcal{A}_i$. Roughly, $\mathcal{R}_i$ is the event that no vertex in H_i has degree deviating too far from the average/expected degree. $\mathcal{B}_i$ is the event that no edge belongs to many more 1-factors than the average edge. And $\mathcal{A}_i$ is the event that H_i has many 1-factors. Essentially, we show that $\mathcal{R}_i$ holds for all i whp. Then, we proceed somewhat inductively. We show that, if $\mathcal{R}_i$ and $\mathcal{A}_i$ hold, then $\mathcal{B}_i$ holds whp. And, we show that, if $\mathcal{B}_i$ and $\mathcal{R}_i$ hold for all $i < t$, then $\mathcal{A}_t$ holds whp.

Let $p_i = \frac{n^k-i}{n^k}$, the fraction of the original edges that remain in H_i. Formally, we define

$$\mathcal{R}_i : \quad \text{for each } x \in V, \left|\deg(x, H_i) - p_i n^{k-1}\right| \leq \frac{1}{K^{1/2}} p_i n^{k-1}$$

where $p_i n^{k-1}$ is the average degree in H_i. For a hypergraph H, we let $\Phi(H)$ be the number of 1-factors of H. We let V_k be the set of collections of k vertices in H with one vertex from each part, that is the possible edges in H_i.

We define the function $w_i : V_k \to \mathbb{R}$ as $w_i(Z) = \Phi(H_i - Z)$. That is, we remove the vertices in Z from H_i and count how many 1-factors there are on the remaining vertices. If Z is an edge, then this is exactly a count of the 1-factors that contain Z: if Z is used in the 1-factor, we must find a way to match the remaining vertices. For a finite set A, we define

$$\bar{w}(A) = \frac{1}{|A|}\sum_{a\in A} w(a), \quad \max w(A) = \max_{a\in A} w(a), \quad \text{maxr } w(A) = \frac{\max w(A)}{\bar{w}(A)}$$

and we define

$$\mathcal{B}_i : \quad \text{maxr } w_i(E_i) \leq K^{1/2}$$

Lastly, we define,

$$\mathcal{A}_i : \quad \log \Phi(H_i) \geq n\left(\log(n^k - i) - \log(n) - O(1)\right)$$

and so $\mathcal{A}_t$ implies that H_T has many 1-factors, $\Omega(n \log\log n)$, and is sufficient to prove the existence of a 1-factor in the final random hypergraph.

We consider the first time $t \leq T$, if any, where $\mathcal{A}_t$ fails. We see that this "first failure at time t" event can be written as a subset of the following event

$$\bar{\mathcal{A}}_t \cap \bigcap_{i<t} \mathcal{A}_i \subseteq \left[\bigcup_{i<t} \bar{\mathcal{R}}_i\right] \cup \left[\bigcup_{i<t} \mathcal{A}_i\mathcal{R}_i\bar{\mathcal{B}}_i\right] \cup \left[\bar{\mathcal{A}}_t \cap \bigcap_{i<t}(\mathcal{B}_i\mathcal{R}_i)\right]$$

And so, by a union bound,

$$\begin{aligned}\Pr\left(\bar{\mathcal{A}}_t \cap \bigcap_{i<t}\mathcal{A}_i\right) &\leq \sum_{i<t}\Pr\left(\bar{\mathcal{R}}_i\right) + \sum_{i<t}\Pr\left(\mathcal{A}_i\mathcal{R}_i\bar{\mathcal{B}}_i\right) + \Pr\left(\bar{\mathcal{A}}_t \cap \bigcap_{i<t}(\mathcal{B}_i\mathcal{R}_i)\right) \\ &\leq \sum_{i<t}\Pr\left(\bar{\mathcal{R}}_i\right) + \sum_{i<t}\Pr\left(\mathcal{A}_i\mathcal{R}_i\bar{\mathcal{B}}_i\right) + \Pr\left(\bar{\mathcal{A}}_t \middle| \bigcap_{i<t}(\mathcal{B}_i\mathcal{R}_i)\right)\end{aligned}$$

We show that each of these three probabilities is $n^{-\omega_K(1)}$ and so, for sufficiently large K, we can take a union bound over the possible t and we see whp there is no time $t \leq T$ where $\mathcal{A}_t$ fails. We also see that in the adversarial deletion context, that is $t \leq T + s$, we are able to bound $\sum_{i<t}\Pr\left(\bar{\mathcal{R}}_i\right)$ and $\Pr\left(\bar{\mathcal{A}}_t \middle| \bigcap_{i<t}(\mathcal{B}_i\mathcal{R}_i)\right)$ but that additional tools are required to bound $\sum_{i<t}\Pr\left(\mathcal{A}_i\mathcal{R}_i\bar{\mathcal{B}}_i\right)$. Bounding this final quantity would be the last step to showing that our hypergraph still has many 1-factors even after the adversarial deletions.

First, we analyze $\Pr\left(\bar{\mathcal{A}}_t \middle| \bigcap_{i<t}(\mathcal{B}_i\mathcal{R}_i)\right)$. We let $\mathcal{F}_i$ denote the set of 1-factors of H_i. We express

$$|\mathcal{F}_t| = |\mathcal{F}_0|\frac{|\mathcal{F}_1|}{|\mathcal{F}_0|}\cdots\frac{|\mathcal{F}_t|}{|\mathcal{F}_{t-1}|} = |\mathcal{F}_0|(1-\xi_1)\cdots(1-\xi_t)$$

or

$$\log|\mathcal{F}_t| = \log|\mathcal{F}_0| + \sum_{i=1}^{t}\log(1-\xi_i) \geq \log|\mathcal{F}_0| + \sum_{i=1}^{t}(\xi_i+\xi_i^2) \tag{1}$$

where $\xi_i = \frac{|\mathcal{F}_i|}{|\mathcal{F}_{i-1}|}$ represents the fractional decrease in the number of perfect matchings from H_{i-1} to H_i.

We can explicitly compute $|\mathcal{F}_0| = (n!)^{k-1}$. In the complete k-partite graph, every possible way to partition the vertices into n groups with one vertex from each of the k parts in each group is a perfect matching. If we label the n groups by the vertex belonging to it from the first part, there are $(n!)^{k-1}$ ways to permute the vertices in each of the $k-1$ other parts to assign them groups. And so,

$$\log|\mathcal{F}_0| = (k-1)n\log n - O(n)$$

We can also explicitly compute the average value of ξ_i. This can also be seen as the expected value of ξ_i for the random deletion steps as the deleted edge is selected uniformly. We see that the expected number of perfect matchings that are destroyed when a random edge is removed from H_{i-1}

$$\begin{aligned}
&= \frac{1}{|E(H_{i-1})|}\sum_{e\in E(H_{i-1})}\text{number of 1-factors containing } e \\
&= \frac{1}{n^k-i+1}\sum_{e\in E(H_{i-1})}\sum_{F\in\mathcal{F}_{i-1}} 1[e\in F] \\
&= \frac{1}{n^k-i+1}\sum_{F\in\mathcal{F}_{i-1}}\sum_{e\in E(H_{i-1})} 1[e\in F] \\
&= \frac{1}{n^k-i+1}\sum_{F\in\mathcal{F}_{i-1}}\text{number of edges in } F \\
&= \frac{n|\mathcal{F}_{i-1}|}{n^k-i+1}
\end{aligned}$$

and so

$$\mathbb{E}[\xi_i] = \gamma_i = \frac{n}{n^k-i+1}$$

We want to show that, in general, the ξ_i do not deviate too greatly from their average/expected values γ_i.

We see that $\mathcal{B}_{i-1}$ implies $\xi_i \leq K^{1/2}\gamma_i$. So, in a random deletion step, ξ_i is a random variable that takes on values $\in [0, K^{1/2}\gamma_i]$ with expected value γ_i. We will see that, over these random deletion steps, $\sum_{i=1}^{T}(\xi_i - \gamma_i) \leq n$ whp. If this is the case, we have

$$\begin{aligned}
\sum_{i=1}^{t} \xi_i &\leq \sum_{i=1}^{t} \gamma_i + \sum_{i=1}^{T}(\xi_i - \gamma_i) \\
&\leq \sum_{i=1}^{t} \gamma_i + n \\
&= n\left(\frac{1}{n^k} + \frac{1}{n^k - 1} + \cdots + \frac{1}{n^k - i}\right) + O(n)
\end{aligned}$$

Using the fact that $\sum_{i=1}^{x} \frac{1}{i} = \log x + e + O(1/x)$, we have

$$\begin{aligned}
\sum_{i=1}^{t} \xi_i &\leq n(\log(n^k) - \log(n^k - i) + O(1/n)) + O(n) \\
&= kn\log(n) - n\log(n^k - i) + O(n)
\end{aligned}$$

Therefore, again using $\xi_i \leq K^{1/2}\gamma_i \leq K^{1/2}/(K\log n)$, we have

$$\sum_{i=1}^{t} \xi_i^2 \leq \frac{1}{K^{1/2}\log n}\sum_{i=1}^{t} \xi_i = O(n)$$

and from (1)

$$\begin{aligned}
\log\Phi(H_i) &= \log|\mathcal{F}_t| \\
&\geq \log|\mathcal{F}_0| - \sum_{i=1}^{t} \xi_i - \sum_{i=1}^{t} \xi_i^2 \\
&= (k-1)n\log n - (kn\log(n) - n\log(n^k - i)) - O(n) \\
&= n\left(\log(n^k - i) - \log(n) - O(1)\right)
\end{aligned}$$

as desired.

Lastly, we show $\sum_{i=1}^{T}(\xi_i - \gamma_i) \leq n$ whp. We see that

$$\Pr\left(\sum_{i=1}^{T}(\xi_i - \gamma_i) \geq n\right) = \Pr\left(e^{h\sum_{i=1}^{T}(\xi_i-\gamma_i)} \geq e^{hn}\right)$$
$$\leq \mathbb{E}\left[e^{h\sum_{i=1}^{T}(\xi_i-\gamma_i)}\right] e^{-hn}$$

Even though the ξ_i are not independent, the analysis we did to compute the $\mathbb{E}[\xi_i]$ only relied on the number of edges in the hypergraph H_i. Therefore, our argument holds regardless of the conditional outcome of the previous ξ variables, and we have

$$\mathbb{E}[\xi_i|\xi_1, \cdots, \xi_{i-1}] = \gamma_i$$

Thus,

$$\mathbb{E}[e^{h\xi_i}|\xi_1, \cdots, \xi_{i-1}] \leq \left(\frac{1}{K^{1/2}}\right) e^{hK^{1/2}\gamma_i} + \left(1 - \frac{1}{K^{1/2}}\right) e^{0}$$

The random variable $e^{h\xi_i}$ will have largest expected value when ξ_i takes on extreme values. Due to the bounds on ξ_i and its expectancy, this extreme case occurs when

$$\xi_i = \begin{cases} K^{1/2}\gamma_i & \text{with probability } \frac{1}{K^{1/2}} \\ 0 & \text{with probability } 1 - \frac{1}{K^{1/2}} \end{cases}$$

This gives us

$$\mathbb{E}[e^{h(\xi_i-\gamma_i)}|\xi_1, \cdots, \xi_{i-1}] \leq e^{-h\gamma_i}\left(1 - \frac{1}{K^{1/2}} + \frac{1}{K^{1/2}} e^{hK^{1/2}\gamma_i}\right) \leq e^{K^{1/2}h^2\gamma_i^2}$$

and so

$$\Pr\left(\sum_{i=1}^{T}(\xi_i - \gamma_i) \geq n\right) \leq \mathbb{E}\left[e^{h\sum_{i=1}^{T}(\xi_i-\gamma_i)}\right] e^{-hn}$$
$$\leq e^{-hn} e^{K^{1/2}h^2\sum_{i=1}^{t}\gamma_i^2}$$
$$\leq e^{K^{-1/2}h^2\log^{-1}(n)\sum_{i=1}^{t}\gamma_i - hn}$$

and since $\sum_{i=1}^{t}\gamma_i = O(n\log n)$, we see that this event holds with low probability for sufficiently large constant h.

This part of the proof does hold for the adversarial deletion setting. For an adversarial deletion step i, we cannot bound ξ_i as a deviation from its expectancy as it is no longer a random variable, but we still have $\xi_i \leq K^{1/2}\gamma_i$ from $\mathcal{B}_i$. And so, for $t \in [T+1, T+s]$,

$$\begin{aligned}
\sum_{i=1}^{t} \xi_i \leq& \sum_{i=1}^{T} \xi_i + \sum_{i=T+1}^{t} \xi_i \\
\leq& \sum_{i=1}^{T} \gamma_i + n + K^{1/2} \sum_{i=T+1}^{t} \gamma_i \\
=& n\left(\frac{1}{n^k} + \frac{1}{n^k-1} + \cdots + \frac{1}{m}\right) + K^{1/2} n \left(\frac{1}{m-1} + \cdots + \frac{1}{m-s}\right) + O(n) \\
\leq& n\left(\frac{1}{n^k} + \frac{1}{n^k-1} + \cdots + \frac{1}{Kn\log n}\right) + K^{1/2} n \left(\frac{1}{m-1} + \cdots + \frac{1}{(1-2\alpha)m}\right) + O(n) \\
\leq& n(\log(n^k) - \log(Kn\log n) + O(1/n)) \\
&+ K^{1/2} n(\log(m-1) - \log((1-2\alpha)m) + O(1/n)) + O(n) \\
=& (k-1)n\log(n) - n\log\log(n) + O(n)
\end{aligned}$$

This is from the fact that $m \geq Kn\log n$ and since $s \leq |E(H^*)|$ and the maximum degree of H^* is $2\alpha\frac{m}{n}$ it can have at most $2\alpha m$ edges. The rest of the proof follows identically.

To bound $\sum_{i<t} \Pr\left(\bar{\mathcal{R}_i}\right)$ and $\sum_{i<t} \Pr\left(\mathcal{A}_i\mathcal{R}_i\bar{\mathcal{B}_i}\right)$, instead of using the random graph model $H^k_{n\times k;m_i}$, where a collection of $m_i = n^k - i$ edges are selected uniformly at random, we use the $H^k_{n\times k;p_i}$ model where each edge is included independently at random with probability $p_i = \frac{n^k-i}{n^k} = \frac{m}{\text{total possible edges}}$. The edge independence of this model will make our argumentation easier. We are able to switch to this model for two reasons. Firstly, for any condition X, the event "X is satisfied on model $H^k_{n\times k;m_i}$" is equivalent to the event "X is satisfied on model $H^k_{n\times k;p_i}$ conditioned on the fact that exactly m_i edges get sampled", as this conditioning makes the two random models equivalent. Secondly, $\Pr(H^k_{n\times k;p_i}$ has exactly m_i edges$) = \Omega(m_i^{-1/2}) = \Omega(n^{-k/2})$. Therefore,

$$\begin{aligned}
\Pr(X \text{ on model } H^k_{n\times k;m_i}) &= \Pr(X \text{ on model } H^k_{n\times k;p_i} \mid \text{exactly } m_i \text{ edges sampled}) \\
&\leq \Pr(X \text{ on model } H^k_{n\times k;p_i})/\Pr(\text{ exactly } m_i \text{ edges sampled}) \\
&\leq O(n^{k/2})\Pr(X \text{ on model } H^k_{n\times k;p_i})
\end{aligned}$$

And so, if we can show that X holds on model $H^k_{n\times k;p_i}$ with sufficiently low probability (an upper bound of the form $n^{-\omega_K(1)}$), setting K to be sufficiently large will also ensure that X holds with low probability on $H^k_{n\times k;m_i}$.

This new model enables us to bound $\sum_{i<t}\Pr\left(\bar{\mathcal{R}}_i\right)$ quite easily. For a specific vertex $x\in V$ and a specific $i\le T$, we have that $\deg(x,H_i)$ is a sum of independent indicator events representing the sampling of the n^{k-1} possible edges containing x. We have $\mathbb{E}[\deg(x,H_i)]=p_in^{k-1}$ and so, by Chernoff,

$$\Pr\left(|\deg(x,H_i)-p_in^{k-1}|\le\frac{1}{K^{1/2}}p_in^{k-1}\right)\le e^{-\Omega(K\log n)}=n^{-\omega_K(1)}$$

and $\sum_{i<t}\Pr\left(\bar{\mathcal{R}}_i\right)=n^{-\omega_K(1)}$ by union bound. We can also bound this term in the adversarial context. Again by Chernoff,

$$\Pr\left(|\deg(x,H_i)-p_in^{k-1}|\le\left(\frac{1}{K^{1/2}}-2\alpha\right)p_in^{k-1}\right)\le e^{-O(K\log n)}=n^{-\omega_K(1)}$$

Even though we have certain requirements to make K sufficiently large, we can make α even smaller so that $\frac{1}{K^{1/2}}>2\alpha$. Taking a union bound over the vertices tells us that $\mathcal{R}_i$ holds whp for all $i\le T$.

We know that the degree of any single vertex can decrease by at most $2\alpha\frac{m}{n}$. And there can be at most $2\alpha m$ total edge deletions, so $p_in^{k-1}\ge(1-2\alpha)p_Tn^{k-1}$. Therefore, since we know $|\deg(x,H_i)-p_in^{k-1}|\le\left(\frac{1}{K^{1/2}}-2\alpha\right)p_in^{k-1}$ holds whp for $i=T$, we must have

$$|\deg(x,H_i)-p_in^{k-1}|\le\frac{1}{K^{1/2}}p_in^{k-1}$$

for all $i\le T+s$, as desired.

Bounding $\sum_{i<t}\Pr\left(\mathcal{A}_i\mathcal{R}_i\bar{\mathcal{B}}_i\right)$ is where we run into problems with the adversarial deletion model, at least in the original proof of JKV. To bound $\sum_{i<t}\Pr\left(\mathcal{A}_i\mathcal{R}_i\bar{\mathcal{B}}_i\right)$, we introduce a new event $\mathcal{C}_i$ that, similar to $\mathcal{B}_i$, bounds the number of 1-factors a single edge can belong to. For $\ell\le k$, we define V_ℓ to be the set of ℓ-tuples of vertices in $V(H)$ such that no two vertices belong the same part. For Y, a set of vertices with at most one vertex in each part ($|Y|\le k$), we define

$$V_{k,Y}=\{Z\in V_k:Z\supseteq Y\}$$

to be the set of collections of k vertices in our hypergraph (one from each part) that contain Y. We let

$$\mathcal{C}_i=\left\{\max w_i(V_{k,Y})\le\max\left(2\text{med }w_i(V_{k,Y}),n^{-(k+1)}|\mathcal{F}_i|\right)\text{ for all }Y\in V_{k-1}\right\}$$

where med $w(A)=$ med $_{a\in A}w(a)$. There are four important distinctions between $\mathcal{C}_i$ and $\mathcal{B}_i$. Firstly, we have replaced deviation from the average number of 1-factors to deviation form the median number. This will be necessary as our arguments will require a guarantee that the "most 1-factor popular" edge is a constant factor more popular than at least half of the other edges. Secondly, instead of comparing the number of 1-factors an edge belongs to to all the other edges, we compare specifically to other edges that share $k-1$ of the same vertices. We will see that this way of quantifying 1-factor popularity deviations is essentially the same as the original way, and it will simply make our argumentation easier. Thirdly, we add a $n^{-(k+1)}|\mathcal{F}_i|$ to the event which makes it so $\mathcal{C}_i$ still holds when med $w_i(V_{k,Y})=0$ or is very small but has very little effect on the event otherwise. Lastly, $\mathcal{B}_i$ is an event that concerns how far the number of 1-factors containing a single edge can deviate from that of the average edge. However, $\mathcal{C}_i$ is concerned with bounding this deviation for k-tuples of vertices that are not guaranteed to be edges (we delete a k-tuple of vertices and find a 1-factor on the residual graph, acting like the K-tuple is an edge when it is not necessarily). We will see that the random edge distribution makes this distinction negligible.

We will show

$$\Pr\left(\bar{\mathcal{C}}_i|\mathcal{A}_i\mathcal{R}_i\right)=n^{-\omega_K(1)} \tag{2}$$

$$\Pr\left(\bar{\mathcal{B}}_i\mathcal{C}_i\mathcal{R}_i\right)=n^{-\omega_K(1)} \tag{3}$$

Then using

$$\mathcal{A}_i\mathcal{R}_i\bar{\mathcal{B}}_i\subseteq\bar{\mathcal{C}}_i\mathcal{A}_i\mathcal{R}_i\cup\bar{\mathcal{B}}_i\mathcal{C}_i\mathcal{R}_i$$

gives

$$\Pr(\mathcal{A}_i\mathcal{R}_i\bar{\mathcal{B}}_i)\le\Pr(\bar{\mathcal{C}}_i\mathcal{A}_i\mathcal{R}_i)+\Pr(\bar{\mathcal{B}}_i\mathcal{C}_i\mathcal{R}_i)\le\Pr(\bar{\mathcal{C}}_i|\mathcal{A}_i\mathcal{R}_i)+\Pr(\bar{\mathcal{B}}_i\mathcal{C}_i\mathcal{R}_i)=n^{-\omega_K(1)}$$

as desired. Thus, by another union bound, $\mathcal{A}_i\mathcal{R}_i\bar{\mathcal{B}}_i$ holds for all i whp.

Proof of (2). To argue $\Pr\left(\bar{\mathcal{C}}_i|\mathcal{A}_i\mathcal{R}_i\right)=n^{-\omega_K(1)}$, we require two lemmas that we will prove in the following section. First, for a vertex $y\in V=V(H)$, we define the random variable $X(y,H)$ to be the edge e containing y in a uniformly random 1-factor of H. We let

$$h(y,H)=-\sum_{e\ni h}\Pr(X(y,H)=e)\log\Pr(X(y,H)=e)$$

denote the entropy of $X(y,H)$.

Lemma 3.3. *For P one of the k vertex parts in V:*

$$\log|\mathcal{F}_i| \leq \sum_{y\in P} h(y, H_i)$$

This lemma encapsulates the randomness of a uniform 1-factor that must be preserved in these vertex-specific random variables.

Lemma 3.4. *For a function $w : S \to \mathbb{R}^+$ from some finite set S to the positive reals, we let X be the random variable with*

$$\Pr(X = x) = \frac{w(x)}{w(S)}$$

where $w(S) = \sum_{s\in S} w(s)$. If $h(X) = -\sum_{x\in S} \Pr(x)\log\Pr(x)$ represents the entropy of X, and

$$h(X) \geq \log|S| - O(1)$$

then there exists $a, b \in range(w)$ with $a \leq b \leq O(a)$ such that for $J = w^{-1}[a, b]$, we have

$$|J| = \Omega(|S|) \text{ and } w(J) > .7w(S)$$

The maximum possible entropy of a random variable with range S is $\log|S|$: the uniform distribution. This lemma demonstrates how a random variable almost maximal entropy must be almost uniformly distributed: the probabilities of almost any two outcomes being at most a constant factor away form each other.

Armed with these two lemmas, we can show $\mathcal{C}_i$ holds whp given $\mathcal{A}_i$ and $\mathcal{R}_i$. Say that $\mathcal{C}_i$ fails at $Y \in V_{k-1}$ where Y contains one vertex from each of the k parts except for the part P. Let $x \in P$ satisfy $w_i(Y \cup \{x\}) = \max w_i(V_{k,Y})$. We note that there is no guarantee that $e = Y \cup \{x\}$ is actually an edge in H_i. Choose $y \in P$ with $w_i(Y \cup \{y\}) \leq \text{med } w_i(V_{k,Y})$ and with $h(y, H_i - e)$ maximum subject to this restriction. Here $H_i - e$ refers to deleting the vertices in e from H_i. We set $f = Y \cup \{y\}$ and also note that $y \neq x$ by its definition.

We see

$$\log(w_i(e)) > \log(n^{-(k+1)}|\mathcal{F}_i|) = n\left(\log(n^k - i) - \log(n) - O(1)\right) \tag{4}$$

by $\mathcal{A}_i$ holding and $\mathcal{C}_i$ failing. Lemma 3.3 implies that

$$\log(w_i(e)) = \log\Phi(H_i - e) \leq \sum_{z\in P\setminus\{x\}} h(z, H_i - e) \tag{5}$$

We can upper bound $h(z, H_i - e) \le h(y, H_i - e)$ for at least half the z's in P, as y maximizes $h(y, H_i - e)$ over all vertices $z \in P$ with $w_i(Y \cup \{z\})$ below the median value. For the other z, we can bound

$$h(z, H_i - e) \le \log(\deg(z, H_i - e)) \le \log\left((1 + o(1))p_i n^{k-1}\right) = \log(n^k - i) - \log(n) + O(1)$$

by $\mathcal{R}_i$. And so,

$$\sum_{z \in P \setminus \{x\}} h(z, H_i - e) \le \frac{n}{2}\left(h(y, H_i - e) + \log(n^k - i) - \log(n) + O(1)\right) \qquad (6)$$

Combining (4), (5), and (6), we get

$$h(y, H_i - e) > \log(n^k - i) - \log(n) - O(1) = \log\deg(y, H_i - e) - O(1)$$

This means that the edge containing y from a uniformly random 1-factor of $H_i - e$ will be well-distributed over all the possible edges containing y. We apply Lemma 3.4 to get J, a, b.

We consider the set W_{k-1} of $(k-1)$-tuples of vertices in $V \setminus Y$ where each $Z \in W_{k-1}$ has exactly one vertex in each part except for P. For each $Z \in W_{k-1}$, $Z \cup \{y\}$ is a possible edge in $H_i - e$ and $Z \cup \{x\}$ is a possible edge in $H_i - f$. That is,

$$w_i(e) = \sum_{Z \in W_{k-1}} 1_{Z \cup \{y\} \in E_i} \cdot \Phi(H_i - (Y + Z + x + y))$$

and

$$w_i(f) = \sum_{Z \in W_{k-1}} 1_{Z \cup \{x\} \in E_i} \cdot \Phi(H_i - (Y + Z + x + y))$$

Since $\mathcal{C}_i$ fails, we must have $w_i(e) \ge 2w_i(f)$. However, returning to the random graph model $H^k_{n \times k; p_i}$ where each edge is sampled independently, we will see that it is very unlikely for these two sums to deviate so greatly, even with adversarial deletions.

We define

$$w'_i(Z) = \Phi(H_i - (Y + Z + x + y))$$

Let $W^*_{k-1} \subseteq W_{k-1}$ be the set of vertex $(k-1)$-tuples Z such that $Z \cup \{y\} \in J$. That is,

$$\Phi(H_i - (e + Z + y)) = w'(Z) \in [a, b] \qquad \forall Z \in W^*_{k-1}$$

We also have

$$\sum_{Z \in W^*_{k-1}} w'_i(Z) > .7 \sum_{Z \in W_{k-1}} w'_i(Z)$$

and so we must have

$$\sum_{Z\in W^*_{k-1}} 1_{Z\cup\{y\}\in E_i}\cdot w'_i(Z) > .7\sum_{Z\in W_{k-1}} 1_{Z\cup\{y\}\in E_i}\cdot w'_i(Z) = .7w_i(e) \tag{7}$$

as any term removed from the left sum (due to the fact that $Z\cup\{y\}\notin E$) must also be removed from the right sum. Additionally,

$$\sum_{Z\in W^*_{k-1}} 1_{Z\cup\{x\}\in E_i}\cdot w'_i(Z) \le \sum_{Z\in W_{k-1}} 1_{Z\cup\{x\}\in E_i}\cdot w'_i(Z) = w_i(f) \le .5w_i(e) \tag{8}$$

In the random graph model $H^k_{n\times k;p_i}$, each event $Z\cup\{y\}\in E_i$ and $Z\cup\{x\}\in E_i$ is an independent Bernuolli distribution with probability p_i. So, the two sums that appear in (7) and (8) have the same expectation, but since one is $> .7w_i(e)$ and the other is $\le .5w_i(e)$, one must deviate from expectation by a factor of at least $\sqrt{.7/.5}$. Since all the $w'_i(Z)$'s are within $O(1)$ of each other and these sums have $|W^*_{k-1}| = |J| = \Omega(|W_{k-1}|)$ terms in them, the probability that these sums deviate by a constant factor is $e^{-\Omega(|W_{k-1}|)} = n^{-\omega(1)}$ by a Chernoff-type bound.

Unfortunately, in the adversarial deletions context, the adversary can almost always force these two sums to deviate greatly. Since the deletions are in terms of the random sampling, the adversary can observe the outcomes of the random events $1_{Z\cup\{x\}\in E_T}$ and alter W^*_{k-1} accordingly. Lemma 3.4 shows that $|W^*_{k-1}| = \Omega(|W_{k-1}|)$, but which $Z\in W_{k-1}$ belong to W^*_{k-1} can be altered. After observing the outcomes of the $Z\cup\{x\}\in E_T$ events, the adversary can make it so $1_{Z\cup\{x\}\in E_T} = 0$ for all $Z\in W^*_{k-1}$, making this sum deviation always hold.

□

Proof of (3). We assume $\Pr(\mathcal{C}_i\mathcal{R}_i)\ge\epsilon = n^{-\omega_K(1)}$. Otherwise, $\Pr\left(\bar{\mathcal{B}}_i\mathcal{C}_i\mathcal{R}_i\right) = n^{-\omega_K(1)}$ is trivialized. We show

$$\Pr\left(\bar{\mathcal{B}}_i|\mathcal{C}_i\mathcal{R}_i\right) = n^{-\omega_K(1)}$$

which is sufficient as $\Pr\left(\bar{\mathcal{B}}_i\mathcal{C}_i\mathcal{R}_i\right) = \Pr\left(\bar{\mathcal{B}}_i|\mathcal{C}_i\mathcal{R}_i\right)\Pr(\mathcal{C}_i\mathcal{R}_i) \le \Pr\left(\bar{\mathcal{B}}_i|\mathcal{C}_i\mathcal{R}_i\right)$. This argument requires an additional lemma that will also be proved in the following section.

Lemma 3.5. *Let w be a function $w: V_k\to\mathbb{R}^+$. Suppose that for each $Y\in V_{k-1}$ with $\max w(V_{k,Y})\ge B$ we have*

$$\left|\left\{Z\in V_{k,Y} : w(Z)\ge\frac{1}{2}\max w(V_{k,Y})\right\}\right| \ge \frac{n}{2}$$

Then, for any $X \in V_{k-j}$ *(a set of* $k-j$ *vertices with no two vertices in the same part) with* $\max w(V_{k,X}) \geq 2^{j-1}B$ *we have*

$$\left|\left\{Z \in V_{k,X} : w(Z) \geq \frac{1}{2^j} \max w(V_{k,X})\right\}\right| \geq \left(\frac{n}{2}\right)^j$$

Conditioning on $\mathcal{C}_i$ holding, we have $\max w_i(V_{k,Y}) \leq \max\left(2\text{med } w_i(V_{k,Y}), n^{-(k+1)}|\mathcal{F}_i|\right)$ for all $Y \in V_{k-1}$. So, for a $Y \in V_{k-1}$ with $\max w_i(V_{k,Y}) \geq n^{-(k+1)}|\mathcal{F}_i|$, we must have $\max w_i(V_{k,Y}) \leq 2\text{med } w_i(V_{k,Y})$ or equivalently

$$\left|\left\{Z \in V_{k,Y} : w_i(Z) \geq \frac{1}{2} \max w_i(V_{k,Y})\right\}\right| \geq \frac{n}{2}$$

Thus, we can apply Lemma 3.5 on w_i with $B = n^{-(k+1)}|\mathcal{F}_i|$. We consider $X = \emptyset$, that is $j = k$, where $\max w_i(V_{k,X}) = \max w_i(V_k)$. We want to show that there are many $Z \in V_k$ that achieve $w_i(Z)$ at least a constant factor of this $\max w_i(V_k)$. And so, in order to apply the lemma, we want to show $\max w_i(V_k) \geq 2^{k-1}B$.

We can trivially lower bound $\max w_i(V_k) \geq \bar{w}_i(V_k)$ by the average value. We lower bound the average by a double counting argument: counting pairs of edges and 1-factors (e, F) in H_i where $e \in F$. Choosing the 1-factor first, there are $|\mathcal{F}_i|$ choices for F and n choices for $e \in F$. Alternatively, we can choose an edge $e \in E(H_i)$ first and then complete the 1-factor in $w_i(e)$ ways. Thus,

$$n|\mathcal{F}_i| = \sum_{e \in E(H_i)} w_i(e) \leq \sum_{e \in V_k} w_i(e)$$

and so

$$\bar{w}_i(V_k) = \frac{1}{n^k} \sum_{e \in V_k} w_i(e) \geq \frac{1}{n^{k-1}}|\mathcal{F}_i| \geq 2^{k-1}B = \frac{2^{k-1}}{n^{k+1}}|\mathcal{F}_i|$$

which holds as long as $n^2 \geq 2^{k-1}$ which is true for n sufficiently large. Thus, we can apply the lemma and get

$$|S| = \left|\left\{Z \in V_k : w_i(Z) \geq \frac{\max w_i(V_k)}{2^k}\right\}\right| \geq \left(\frac{n}{2}\right)^k \tag{9}$$

This gives us that many $Z \in V_k$ achieve at least a constant factor of the maximal w_i value. The last step is to show that, whp, many actual edges ($Z \in E_i$) also achieve this. Specifically, we define the set

$$E_i^* = \{e \in E_i : w_i(e) \geq \delta \max w_i(V_k)/2\}$$

where $\delta = \frac{1}{2^k}$. We will show $|E_i^*| \geq c\,|E_i|$ for some constant c, which gives

$$\frac{\sum_{e \in E_i} w_i(e)}{\max w_i(V_k)} \geq \frac{\sum_{e \in E_i^*} w_i(e)}{\max w_i(V_k)} \geq \frac{\delta |E_i^*|}{2} \geq \frac{\delta c |E_i|}{2}$$

which implies $\max w_i(E_i) \leq \max w_i(V_k) \leq \frac{2\sum_{e \in E_i} w_i(e)}{\delta c |E_i|} = \frac{2}{\delta c}\bar{w}_i(E_i)$ and property $\mathcal{B}_i$ holds for sufficiently large $K^{1/2} \geq \frac{2}{\delta c}$. This is also where the argument breaks down in the adversarial context.

From (9), we have that $\mathcal{C}_i\mathcal{R}_i$ holding implies there are at least $\delta n/2$ vertices x_1 from the first vertex part that belong to at least $\delta n^{k-1}/2$ k-tuples $Z \in S$. This comes from pigeonhole. For each of the δn^k vertex k-tuples in S, we consider the vertex in the first vertex part it contains $x_1 \in P_1$. Any x_1 can belong to at most n^{k-1} k-tuples in S as there are n^{k-1} total k-tuples in V_k containing x_1. So, in the worst case, we could have $\delta n/2 - 1$ vertices in P_1 that each correspond to n^{k-1} tuples in S, and the remaining $n - \delta n/2 + 1 \leq n$ vertices in P_1 each correspond to $\delta n^{k-1}/2 - 1$ vertices in S. However, this only gives

$$(\delta n/2 - 1)n^{k-1} + n(\delta n^{k-1}/2 - 1) = \delta n^k - n^{k-1} - n$$

k-tuples in S. And any additional tuples added to S will satisfy the condition.

We sample a random graph $H^k_{n\times k, p_i}$ conditioned on $\mathcal{C}_i\mathcal{R}_i$. Ideally, we consider an x_1 that belongs to $\delta n^{k-1}/2$ k-tuples $Z \in S$, and we expect many of these Z to be edges as each edge is included independently at random. However, there is the caveat that we are conditioning on $\mathcal{C}_i\mathcal{R}_i$ holding and on the fact that x_1 is one of the $\delta n/2$ vertices in P_1 belonging to many k-tuples in S. In the following arguments, we will show that E_i^* is still sufficiently large despite this conditioning.

We fix $0 \leq \ell \leq (k-1)n\log n$ and let $L = 2^\ell$. Fix a vertex $x_1 \in P_1$ and let

$$S_L = \{Z \in V_k : x_1 \in Z \text{ and } w_i(Z) \geq L\}$$

Here L will be an approximation to the random variable $\delta \max w_i(V_k)$. Using L in place of $\max w_i(V_k)$ reduces the conditioning. There are not too many choices for L and so we will be able to use the union bound over L.

Note that without the conditioning $\mathcal{C}_i\mathcal{R}_i$, the two events $\{I \subseteq S_L, J \cap S_L = \emptyset\}$ and $\{I \subseteq E_i, J \cap E_i = \emptyset\}$ will be independent for all $I, J \subseteq V_k$. To understand this, we can assume that $Z \in J$ implies $x_1 \in Z$, since otherwise $Z \notin S_L$ regardless of it being in E_i. Then, knowing $\{I \subseteq S_L, J \cap S_L = \emptyset\}$ informs us only that about the values of $w_i(Z), Z \in I \cup J$ and these values depend on edges that do not contain x_1, as $w_i(Z)$ is counting perfect matchings in a graph with the vertex x_1 removed.

Thus, in a random graph $H^k_{n\times k,p_i}$ without conditioning on $\mathcal{C}_i\mathcal{R}_i$, $|S_L \cap E_i|$ will be distributed as $Bin(|S_L|, p_i)$. Hence, if $|S_L| \geq \Delta$, then

$$\Pr(|S_L \cap E_i| \leq p_i\Delta/2) \leq e^{-p_i\Delta/8}$$

by Chernoff. Lastly, we can deal with the conditioning $\mathcal{C}_i\mathcal{R}_i$ using the simple bound $\Pr(A|B) \leq \Pr(A)/\Pr(B)$ and the fact that, at the start of our arguments, we assumed $\Pr(\mathcal{C}_i\mathcal{R}_i) \geq \epsilon$:

$$\Pr(|S_L \cap E_i| \leq p_i\Delta/2|\mathcal{C}_i\mathcal{R}_i) \leq \frac{1}{\epsilon}e^{-p_i\Delta/8}$$

The number of choices for ℓ is $(k-1)n\log n$ and for one of these (ℓ^*) we will have $2^{\ell^*} \leq \delta \max w_i(V_k) \leq 2^{\ell^*+1}$; the $2n\log n$ upper bound coming from the fact that even in the complete k-partite graph, there are only $(n!)^{k-1}$ perfect matchings. For this ℓ^*, taking $L^* = 2^{\ell^*}$, we must have the elements of S that contain x_1 all belong to S_{L^*} as well, as $\delta \max w_i(V_k) \geq L^*$. So, for at least $\delta n/2$ choices of $x_1 \in P_1$, we have $|S_{L^*}| \geq \delta n^{k-1}/2$, and for each of these

$$\Pr(|S_{L^*} \cap E_i| \leq \delta n^{k-1}p_i/4) \leq \frac{1}{\epsilon}e^{-\delta p_i n^{k-1}/16} \leq \frac{1}{\epsilon}n^{-\delta K/16}$$

Since there are at most n choices for x_1 and $(k-1)n\log n$ choices for ℓ, by a union bound, we have that, with probability at least $1 - \frac{1}{\epsilon}n^{2+o(1)-\delta K/16}$, for $(\delta n/2) \cdot (\delta n^{k-1}p_i/4) = \frac{\delta^2}{8}|E_i|$ choices for $Z \in V_k$, Z is an edge and $w_i(Z) \geq 2^{\ell^*} \geq \delta \max w_i(V_k)/2$. Thus,

$$\Pr(|E_i^*| \leq \frac{\delta^2}{8}|E_i|) \leq \frac{1}{\epsilon}n^{2+o(1)-\delta K/16} = \epsilon$$

for $\epsilon = \sqrt{n^{2+o(1)-\delta K/16}} = n^{-\omega_K(1)}$ for sufficiently large K, as desired.

We cannot extend these arguments to the adversarial context, unfortunately. While $\mathcal{C}_i\mathcal{R}_i$ holding guarantees that S_L is large for many x_1, there is no guarantee as to which Z will belong to S_L. The adversary can witness the edges randomly sampled in E_i that contain x_1 and perform deletions to shift the elements of S_L off of these edges.

□

4 Lemmas

4.1 Proof of Lemma 3.3

This follows from Shearer's lemma, which states

Lemma 4.1 (Shearer)**.** *Let $X = (X_1, X_2, \cdots, X_N)$ be a (vector) random variable and $\mathcal{A} = \{A_i : i \in I\}$ be a collection of subsets of a set B, where $|B| = N$, such that each element of B appears in at least k members of $\mathcal{A}$. For $A \subseteq B$, let $X_A = (X_j : j \in A)$. Then,*

$$h(X) \leq \frac{1}{k} \sum_{i \in I} h(X_{A_i})$$

To get Lemma 3.3 from this, we take $B = V_k$ and let X be the indicator of the edges present in a uniformly random 1-factor. We take $\mathcal{A} = (A_v : v \in P)$, where $A_v \subseteq V_k$ whose vertex from part P is v. Thus, each $e \in B$ belongs to at least (exactly) 1 element of $\mathcal{A}$. Finally, note that the entropy $h(X_{A_v}) = h(v, H_i)$ and $h(X) = \log |\mathcal{F}_i|$ since X is (essentially) a random 1-factor. The proof of Shearer's lemma is as follows:

Proof. By the chain rule of entropy,

$$h(X) = \sum_{j \in B} h(X_j | X_1, X_2, \cdots, X_{j-1}) \tag{10}$$

and

$$h(X_{A_i}) = \sum_{j \in A_i} h(X_j | X_\ell, \ell \in A_i, \ell < j) \tag{11}$$

We sum (11) for all $i \in I$ and obtain

$$\sum_{i \in I} h(X_{A_i}) = \sum_{i \in I} \sum_{j \in A_i} h(X_j | X_\ell, \ell \in A_i, \ell < j)$$
$$= \sum_{j \in B} \sum_{A_i \ni j} h(X_j | X_\ell, \ell \in A_i, \ell < j)$$

using entropy inequality $h(X|Y, Z) \leq h(X|Y)$:

$$\geq \sum_{j \in B} \sum_{A_i \ni j} h(X_j | X_1, X_2, \cdots, X_{j-1})$$

since each $j \in B$ appears in at least k A_i's:

$$\geq k \sum_{j \in B} h(X_j | X_1, X_2, \cdots, X_{j-1})$$

using (10):

$$\geq kh(X)$$

as desired.

$\square$

4.2 Proof of Lemma 3.4

This lemma is proved in the textbook "Introduction to Random Graphs" by Frieze and Karoński [6]. We reference the reader to Lemma 14.14, page 284.

Therein, they prove that if $h(X) \geq \log|S| - M$, then there exist $a, b \in range(w)$ with

$$a \leq b \leq 2^{4(M+\log 3)}a$$

such that for $J = w^{-1}[a, b]$, we have

$$|J| \geq e^{-2M-2|S|} \text{ and } w(J) > .7w(S)$$

4.3 Proof of Lemma 3.5

Proof. Our goal is to show that for any $X \in V_{k-j}$ with $\max w(V_{k,X}) \geq 2^{j-1}B$ we have

$$\left|\left\{Z \in V_{k,X} : w(Z) \geq \frac{1}{2^j}\max w(V_{k,X})\right\}\right| \geq \left(\frac{n}{2}\right)^j$$

The assumption of the argument, that for each $Y \in V_{k-1}$ with $\max w(V_{k,Y}) \geq B$ we have

$$\left|\left\{Z \in V_{k,Y} : w(Z) \geq \frac{1}{2}\max w(V_{k,Y})\right\}\right| \geq \frac{n}{2}$$

will form the base case of an induction on j ($j = 1$). Our inductive hypothesis will therefore be that the lemma holds for all $j' < j$.

Consider an $X \in V_{k-j}$ with $\max w(V_{k,X}) \geq 2^{j-1}B$. Let $Z \in V_{k,X}$ be this k-tuple of vertices that maximizes w, that is $w(Z) = \max w(V_{k,X})$. Let P be one of the j vertex parts that contains none of the vertices in X and let $y \in P$ be the vertex from that part that belongs to Z. Taking $Y = X \cup \{y\}$, we see that $Y \in V_{k-(j-1)}$ satisfies

$$\max w(V_{k,Y}) = w(Z) \geq 2^{j-1}B \geq 2^{j-2}B$$

as $Z \in V_{k,Y} \subseteq V_{k,X}$ maximizes w. And so, by our inductive hypothesis, there are at least $\left(\frac{n}{2}\right)^{j-1}$ sets $Z' \in V_{k,Y}$ with

$$w(Z') \geq \frac{1}{2^{j-1}} \max w(V_{k,Y}) = \frac{1}{2^{j-1}} \max w(V_{k,X})$$

For each such $Z', Z' \setminus \{y\}$ belongs to V_{k-1} with $\max w(V_{k,Z'\setminus\{y\}}) \geq w(Z') \geq B$. So, we can again apply our inductive hypothesis, this time with $j = 1$, to obtain that there are at least $\frac{n}{2}$ sets $Z'' \in V_{k,Z'\setminus\{y\}}$ with

$$w(Z'') \geq \frac{1}{2} \max w(V_{k,Z'\setminus\{y\}}) \geq \frac{1}{2} w(Z') \geq \frac{1}{2^j} \max w(V_{k,X})$$

Each Z'' necessarily belongs to $V_{k,X}$. Additionally, the $\frac{n}{2}$ Z'' associated with each of the $\left(\frac{n}{2}\right)^{j-1}$ Z' must all be distinct as $Z' \setminus X$ is distinct for each Z' and the vertex in P from each Z'' associated with a specific Z' is distinct. Thus,

$$\left|\left\{Z \in V_{k,X} : w(Z) \geq \frac{1}{2^j} \max w(V_{k,X})\right\}\right| \geq \left(\frac{n}{2}\right)^{j-1} \cdot \frac{n}{2} = \left(\frac{n}{2}\right)^j$$

as desired.

□

References

[1] Bal, Deepak; Frieze, Alan (2013), "The Johansson-Kahn-Vu solution of the Shamir problem", www.math.cmu.edu/ af1p/Teaching/ATIRS/Papers/FRH/Shamir.pdf

[2] Johansson, Anders; Kahn, Jeff; Vu, Van (2008), "Factors in Random Graphs", arXiv:0803.3406 [math.CO]

[3] Ferber, Asaf; Vu, Van (2016), "Packing Perfect Matchings in Random Hypergraphs", arXiv:1606.09492 [math.CO]

[4] M. Ajtai, J. Komlos, and E. Szemerédi, "A dense infinite Sidon sequence", Eur J Combinat 2 (1981), 111. MR0611925 (83f:10056)

[5] V. Rödl, "On a packing and covering problem", European Journal of Combinatorics 6, no. 1 (1985), 69-78.

[6] Frieze, Alan; Karonski, Michal, "INTRODUCTION TO RANDOM GRAPHS", Sec 14.3 pg 275

[7] D. Kühn, D. Osthus, “Hamilton cycles in graphs and hypergraphs: an extremal perspective”, arXiv:1402.4268 [math.CO]

[8] A. Frieze and M. Krivelevich, Packing hamilton cycles in random and pseudo-random hypergraphs, Random Structures and Algorithms, 41(1) (2012), 1-22.

[9] B. Sudakov, V. Vu, “Local resilience of graphs”, arXiv:0706.4104 [math.CO]

[10] J. Böttcher, “Large scale structures in random graphs”, arXiv:1702.02648 [math.CO]