MATCHINGS AND LOOSE CYCLES IN THE SEMIRANDOM HYPERGRAPH MODEL

MICHAEL MOLLOY, PAWEŁ PRAŁAT, AND GREGORY B. SORKIN

Abstract. We study the 2-offer semirandom 3-uniform hypergraph model on n vertices. At each step, we are presented with 2 uniformly random vertices. We choose any other vertex, thus creating a hyperedge of size 3. We show a strategy that constructs a perfect matching, and another that constructs a loose Hamilton cycle, both succeeding asymptotically almost surely within $\Theta(n)$ steps. Both results extend to s-uniform hypergraphs. Our methods are qualitatively different from those that have been used for semirandom graphs. Much of the analysis is done on an auxiliary graph that is a uniform k-out subgraph of a random bipartite graph, and this tool may be useful in other contexts.

1. Introduction and Main Results

1.1. **The model and our contribution.** We consider the **semirandom graph process** suggested by Peleg Michaeli and introduced formally in [3]. This model, that can be viewed as a "one player game", generalizes many well-known and important processes such as the Erdős-Rényi random graph process and the k-out process, and has received a lot of attention recently. We concentrate on its natural generalization to hypergraphs, as proposed in [1].

In the r**-offer** s**-uniform semirandom model**, the player constructs a sequence of s-uniform hypergraphs H_t on vertex set $[n] := \{1, \ldots, n\}$, with the goal that H_t satisfies a given property $\mathcal{P}$ as quickly as possible. In each **round** $t \in \mathbb{N}$, a uniformly random set of r vertices is **offered**, and the player **chooses** $s - r$ additional vertices according to some **strategy**, to make a hyperedge that is added to H_{t-1} to form H_t. When choosing the $s-r$ vertices, the player has full knowledge of the r vertices offered and of the past, including the hypergraph H_{t-1} at the end of the previous round. The hypergraph results in [1] are direct extensions of the graph results, and are only for the 1-offer model.

Our contribution is to resolve some natural first questions for the hypergraph semirandom process, notably in the case where each offer specifies two vertices rather than just one. We believe that this is the first work on semirandom hypergraphs that is qualitatively different from anything that has been done for semirandom graphs. In this work, the property $\mathcal{P}$ is that the semirandom hypergraph should contain as a subgraph a desired structure, namely a perfect matching or a loose Hamilton cycle, and constructing such spanning structures has also been one focus of the study of semirandom graphs. A main point of difference is that for graphs, it has always been possible to quickly and simply construct most of the structure desired, then cleverly fill in the final gaps. That approach does not seem to work here. Instead, we first generate a large absorbing structure of some sort, build a small part of the desired structure using the few vertices absent from the absorber, then use the absorber to complete the desired structure.

To illustrate the profoundly greater challenges of semirandom hypergraphs compared with graphs, let us mention another possible approach to matching, which works for graphs but seemingly not for hypergraphs. In both cases, it is easy to semirandomly construct a random k-out graph (or hypergraph). In the graph case, by [9], the k-out almost surely contains a perfect

Date: 31 March 2024.

matching. In the hypergraph case, by contrast, Devlin and Kahn [5] observe that the corresponding statement is "almost certainly correct [but] likely to be difficult [to show]", as it is stronger than the 2008 resolution of Shamir's Problem by [16] and would imply a (still open) "natural guess regarding a beautiful problem of Frieze and Sorkin" (from [7]). Devlin and Kahn make progress towards this by proving the presence of a perfect *fractional* matching, but that is not sufficient for our purpose. Even if this approach did work, it would be existential; we prefer (and obtain) an efficient construction.

Results presented in this paper are asymptotic by nature. We say that a property $\mathcal{P}$ holds **asymptotically almost surely** (**a.a.s.**) if the probability that H has property $\mathcal{P}$ tends to 1 as n goes to infinity.

In our first main result, we concentrate on the 2-offer semirandom construction of a 3-uniform hypergraph H on n vertices, n divisible by 3, so that H will contain a perfect matching (a partition of the set of vertices into $n/3$ hyperedges).

Theorem 1. *For some constant C, the 2-offer 3-uniform semirandom model has a strategy that, for n divisible by 3, in time $t = Cn$, gives a hypergraph H_t a.a.s. containing a perfect matching.*

We present the strategy explicitly, and show how to construct a matching in H (not merely proving its existence). This result is clearly best possible up to the constant C, which we have made no effort to optimise.

Our linear-time construction extends to perfect matchings in an s-uniform hypergraph when the player is offered $r = 2$ vertices and is able to select $s-2$ vertices. Moreover, it immediately implies the same result for a "1-offer" semirandom model, as the player can simply simulate a 2nd random offer vertex with the first of the $s-1$ chosen vertices.

Theorem 2. *For any $s \geq 3$ and $r \in \{1, 2\}$, for some constant C, the r-offer s-uniform semirandom model has a strategy that, for n divisible by s, in time $t = Cn$, gives a hypergraph H_t a.a.s. containing a perfect matching.*

Our second main result is a somewhat similar strategy to construct a 3-uniform hypergraph H on n vertices, n divisible by 2, so that H will contain contain a loose Hamilton cycle.

Theorem 3. *For some constant C, the 2-offer 3-uniform semirandom model has a strategy that, for n divisible by 2, in time $t = Cn$, gives a hypergraph H_t a.a.s. containing a loose Hamilton cycle.*

This result also extends.

Theorem 4. *For any $s \geq 3$ and $r \in \{1, 2\}$, for some constant C, the r-offer s-uniform semirandom model has a strategy that, for n divisible by $s-1$, in time $t = Cn$, gives a hypergraph H_t a.a.s. containing a loose Hamilton cycle.*

1.2. **Outline.** Section 2 gives some background on the semirandom process. Section 3 sketches the approach we will take. The approach relies greatly on a certain auxiliary graph, and the section motivates, states, and proves the lemmas governing the auxiliary graph. (Roughly: a random bipartite graph contains a large uniformly random k-out subgraph. A.a.s., every large subgraph of such a graph, if it has minimum degree at least $k-1$ and parts of equal size, contains a perfect matching.) Section 4 proves Theorem 1 by presenting a 2-offer strategy for constructing a 3-uniform hypergraph matching, and analysing it. Subsection 4.4 proves the extension to s-uniform hypergraphs. This is a simple modification of the argument for the 3-uniform case but requires reference to the strategy: Theorem 2 is not a black-box corollary of Theorem 1. Section 5 proves Theorems 3 and 4. The strategy for constructing a loose Hamilton cycle is different in its details from that for matching, with some new complications, but it has the same general structure and shares many key elements.

2. Background

We begin with a clarification. Recall that our r-offer, s-uniform semirandom hypergraph is formed as follows: At each round t, the player is offered r uniform vertices and then chooses $s-r$ additional vertices to form a hyperedge. In our proofs and previous works, there are situations where, upon seeing the offered vertices, the player decides not to build any hyperedge on them and so does not bother choosing $s-r$ vertices. We say that the player *ignores* the offer. Formally, we can think of the player adding an arbitrary set of $s-r$ vertices and then never using the resulting hyperedge. This counts as a round and so, e.g., is one of the Cn steps in Theorem 1.

Since the semirandom process is still relatively new, let us highlight a few results on this model, focusing first on perfect matchings.

In the very first paper [3], it was shown that the semirandom process on graphs is general enough to simulate several well-studied random graph models (using suitable strategies), including the extensively studied k-out process. In the k-out process on n vertices, each vertex independently connects to k randomly selected vertices, resulting in a random graph with kn edges (see, for example, Chapter 18 in [17]). This is easily simulated with the semirandom process by "choosing" vertex 1 to accompany each of the first k offered vertices, choosing vertex 2 for the next k, etc.

Since the 2-out process has a perfect matching a.a.s. [9], we immediately get that a.a.s. one can create a perfect matching in $(2+o(1))n$ rounds. Using the semirandom process to construct a different sort of random graph known to have a perfect matching a.a.s. [18], the upper bound can be improved to $(1+2/e+o(1))n < 1.73576n$ [3]. This bound was subsequently improved by investigating a fully adaptive algorithm. The current best upper bound is $1.20524n$ [13]. On the other hand, [3] observed that $(\ln(2)+o(1))n > 0.69314n$ is a lower bound on the number of rounds needed to create a perfect matching, and this has since been improved to $0.93261n$ [13].

In the 2-offer 3-uniform semirandom model it is equally easy to generate a random k-out hypergraph, i.e., a hypergraph with k random edges on (and associated with) each vertex: to each of the first k offered pairs add vertex 1 (or a random vertex, if vertex 1 was offered), and so on. However, as discussed in the Introduction, it is not known whether a random k-out hypergraph a.a.s. contains a perfect matching, and we take a different approach to constructing a perfect matching.

Let us briefly comment on other results on the semirandom graph model. Unsurprisingly, Hamilton cycles have received a lot of attention [3, 11, 12, 8, 6], and currently the best upper and lower bounds are $1.81701n$ and $1.26575n$ [6]. More general spanning subgraphs were considered in [2]. It showed that any graph with maximum degree Δ can be constructed a.a.s. in $(3\Delta/2+o(\Delta))n$ rounds; it also showed that if $\Delta = \omega(\log(n))$, this upper bound improves to $(\Delta/2+o(\Delta))n$ rounds, close to the trivial lower bound of $\Delta n/2$. Note, however, that these results are meaningful only for $\Delta \to \infty$; for any fixed Δ they say only that $O(n)$ rounds suffice. Cliques, chromatic number, and independent sets were considered in [10].

Finally, while our interest is in hypergraphs, let us mention some other variants of the semirandom process. In [19], sharp thresholds were studied for a more general class of processes that includes the semirandom process. In [4], a random spanning tree of K_n is presented, and the player keeps one of the edges. In [14], vertices are presented by the process in a random permutation. In [21], the process presents k random vertices, and to create an edge the player selects one of them, and freely chooses a second vertex.

3. Outline, and lemmas on an auxiliary bipartite graph

A *uniformly random k-out bipartite multigraph* (or just "uniform k-out" graph) on disjoint node sets $\mathcal{A}, \mathcal{B}$ is a directed graph $\mathcal{D}$ in which every node in $\mathcal{A}$ has outedges to k nodes in $\mathcal{B}$ chosen uniformly at random with replacement, and symmetrically for $\mathcal{B}$.

A uniform k-out auxiliary graph plays a central role in our semirandom construction. To motivate the lemmas in this section, we briefly explain the role each will play. Proper explanations will come in later sections.

Throughout, we will refer to vertices of the auxiliary graph as "nodes", calling vertices of the semirandom process itself "points", "elementary vertices", or occasionally "vertices". We will use calligraphic letters in this section, for general node sets $\mathcal{A}$ and $\mathcal{B}$ and a corresponding k-out graph $\mathcal{D}$. In context, later, we will define particular node sets A and B and a corresponding k-out graph D.

The number of nodes in the auxiliary graph is different from the number of points in the hypergraph. We will reserve n for the latter, writing N in the lemmas below to make it easy later to define N in terms of n. The k in the lemmas is the same as that in the main argument; indeed, we will take $k = 10$ throughout.

Our semirandom constructions of both hypergraph matchings and loose Hamilton cycles will begin by defining node sets $\mathcal{A}$ and $\mathcal{B}$ (differently for the two cases).

Phase 1 will form a uniform k-out graph from a subset of a set of uniformly random edges on $\mathcal{A} \times \mathcal{B}$. Because the edges supplied are uniformly random, a few nodes are bound to have too few edges, but Lemma 6 shows that we can get a uniform k-out graph $\mathcal{D}$ on *most* of the nodes in $\mathcal{A}$ and $\mathcal{B}$. The few nodes missing from $\mathcal{D}$ will correspond to points of the hypergraph we will have to deal with specially.

Phase 2 will deal with these points. The "actions" addressing them will have the side effect of deleting more nodes from $\mathcal{D}$, correspondingly decreasing the degrees of the remaining nodes. We will arrange that the outdegree of a node in $\mathcal{D}$ is always either k or $k-1$, by ensuring that each outdegree is reduced at most once (i.e., by one action), and by at most 1. The "by at most 1" comes from ensuring that we never delete a node with parallel inedges (Claim 7 shows that such nodes are rare), and that we never simultaneously delete a set of nodes sharing an inneighbour (Claim 8 shows that such sets are rare). (Alternatively, "by at most constant" is immediate from the nature of the actions used, and we could have used a larger k to end with the same outdegree bound.) The "at most once" is achieved by "blocking" a deleted node's inneighbour's outneighbours from future deletion, adding these nodes to a set Q of nodes not to be deleted; Lemma 9 is the basis for showing that the set Q remains small. After all the actions, $\mathcal{D}$ still contains most of $\mathcal{A}$ and $\mathcal{B}$.

Phase 3 will complete the hypergraph construction using $\mathcal{D}$ as an absorber. Walkup [22] showed that a random 2-out bipartite graph has a perfect matching a.a.s., and Frieze [9] showed the same without the bipartiteness condition, for 2-out graphs and multigraphs. Here, despite the fact that $\mathcal{D}$ is neither exactly k-out nor uniformly random, Lemma 10 shows that almost certainly Hall's condition is satisfied, so it has a perfect matching. The hypergraph construction will be completed by hyperedges corresponding to the matching edges.

Throughout, we rely on McDiarmid's form of the Azuma-Hoeffding inequality [20].

Theorem 5 (Azuma)**.** *Let $X_1, \dots X_n$ be independent random variables with X_k taking values in a set A_k for each k. Suppose that the (measurable) function $f\colon \prod A_k \to \mathbb{R}$ satisfies $|f(x) - f(x')| \le c_k$ whenever the vectors x and x' differ only in the kth coordinate. Let Y be the random variable $f(X_1, \dots, X_n)$. Then for any $t > 0$, $\Pr(|Y - \mathbb{E}Y| \ge t) \le 2\exp(-2t^2/\sum c_k^2)$.*

Lemma 6 (k-out)**.** *For all integers $k > 0$ and reals $\epsilon > 0$, there exists a constant $C > 0$ such that the following holds. Let $\mathcal{A}$ and $\mathcal{B}$ be disjoint nonempty sets with $\mathcal{A} \cup \mathcal{B} = [N]$.*

Let Ψ be a multiset of CN (or more) uniformly random undirected edges on $\mathcal{A} \times \mathcal{B}$. Then, a.a.s., we can delete some edges from Ψ, and orient the remaining ones, to give a directed bipartite *multigraph $\mathcal{D}$ on parts $\mathcal{A}' \subseteq \mathcal{A}$ and $\mathcal{B}' \subseteq \mathcal{B}$, where $|\mathcal{A}'| \geq (1-\epsilon)|\mathcal{A}|$, $|\mathcal{B}'| \geq (1-\epsilon)|\mathcal{B}|$, and $\mathcal{D}$ is a uniformly random k-out bipartite multigraph on $\mathcal{A}' \times \mathcal{B}'$. The algorithm for constructing $\mathcal{D}$ is polynomial time.*

Proof. Fix $\delta = \min\left\{\epsilon/2, 0.99\left(16e^2\binom{k+3}{4}\right)^{-1/2}\right\}$. Let λ be the value such that the probability that a Poisson random variable $Z \sim \mathrm{Po}(\lambda)$ is less than $k+3$ is $\delta/4$; let us write this as $\Pr(\mathrm{Po}(\lambda) < k+3) = \delta/4$. Choose $C = 2.1\lambda$.

Assume that $\mathrm{Po}(2\lambda N) \leq CN$; the probability that this fails to hold is exponentially small in N. Randomly orient the first $\mathrm{Po}(2\lambda N)$ edges of Ψ, so there are $\mathrm{Po}(\lambda N)$ edges out of each part.

Suppose part $\mathcal{A}$ has m nodes. Choose t to be the power of 2 for which $\frac{1}{2}N/t < m \leq N/t$. The number of outedges $Z(v)$ from a node $v \in \mathcal{A}$ is Poisson with intensity $\lambda N/m \geq \lambda N/(N/t) = t\lambda$. We may thus write $Z(v)$ as $Z(v) = Z_1(v) + \cdots + Z_t(v) + Z'(v)$ where every $Z_i(v) \sim \mathrm{Po}(\lambda)$, $Z'(v)$ is also Poisson, and all are independent. The event $Z(v) < k+3$ occurs only if every $Z_i(v) < k+3$.

Let X be the number of nodes of $\mathcal{A}$ with degree less than $k+3$. From the above, $X = \sum_{v\in\mathcal{A}} \mathbf{1}(Z(v) < k+3) \leq \sum_{v\in\mathcal{A}} \mathbf{1}[(Z_1(v) < k+3) \wedge \cdots \wedge (Z_t(v) < k+3)]$, so the event $X \geq \delta m$ requires that, of all tm independent probability-$\delta/4$ events, at least $t\delta m$ hold. Then, $\Pr(X \geq \delta m) \leq \Pr(\mathrm{Bin}(tm, \delta/4) \geq t\delta m) \leq \Pr(\mathrm{Bin}(N, \delta/4) \geq \delta N/2)$. With δ constant, this has probability exponentially small in N.

Then, with exponentially small failure probability, all but a $1-\delta$ fraction of nodes in both $\mathcal{A}$ and (by the same argument, and a union bound) in $\mathcal{B}$ have at least $k+3$ outedges. For any node with more than $k+3$ outedges, randomly delete some to leave exactly $k+3$.

Let $\mathcal{D}_0$ be the resulting directed bipartite graph. Let X_0 be the nodes failing to have outdegree $k+3$. Let $\mathcal{D}$ be the k-core of $\mathcal{D}_0 \setminus X_0$. $\mathcal{D}$ can be constructed from $\mathcal{D}_0$ by initialising $X = X_0$, successively adding to X any node with less than k outneighbours in $\mathcal{D}_0 \setminus X$, or equivalently with at least 4 outneighbours in X. When there are no more such nodes, $\mathcal{D}$ is the subgraph of $\mathcal{D}_0$ induced by $\mathcal{A}' = \mathcal{A} \setminus X$ and $\mathcal{B}' = \mathcal{B} \setminus X$.

We claim that, at the end, $|X \cap \mathcal{A}| < 2\delta|\mathcal{A}|$ and $|X \cap \mathcal{B}| < 2\delta|\mathcal{B}|$ a.a.s. By hypothesis, $\delta \leq \epsilon/2$, so this condition immediately implies the lemma's conclusion. Suppose the condition does not hold. At any time, let S and T be the nodes added in parts $\mathcal{A}$ and, respectively, $\mathcal{B}$ that bring X_0 to X. (I.e., $S = (X\cap\mathcal{A})\setminus(X_0\cap\mathcal{A})$.) For the condition to be violated, at the end we must have $|S| \geq \delta|\mathcal{A}|$ or $|T| \geq \delta|\mathcal{B}|$. Consider the first moment that either event first occurs. There are two symmetric cases, and we consider just the one where, at that moment, $|S| = \delta|\mathcal{A}|$ (up to integrality) while $|T| < \delta|\mathcal{B}|$. Note that every node in S has at least 4 outneighbours in $T \cup (X_0 \cap \mathcal{B})$. This remains true if we expand T arbitrarily (but disjoint from X_0) so that $|T| = \delta|\mathcal{B}|$. The probability that any such pair S, T exists is at most the expected number of such pairs, over all fixed subsets and subject to the randomness of $\mathcal{D}$. For a fixed $v \in S$, the probability v has at least 4 outneighbours in $T \cup (X_0 \cap \mathcal{B})$ is at most $\binom{k+3}{4}\left(|T \cup (X_0 \cap \mathcal{B})| \,/\, |\mathcal{B}|\right)^4 = \binom{k+3}{4}(2\delta)^4$. Then, the expected number of pairs S, T is at most

$$\binom{N}{\delta N}\binom{N}{\delta N}\left(\binom{k+3}{4}(2\delta)^4\right)^{\delta N} \leq \left(\frac{eN}{\delta N}\right)^{2\cdot\delta N}\left(\binom{k+3}{4}(2\delta)^4\right)^{\delta N}$$
$$= \left(e^2\binom{k+3}{4}2^4\delta^2\right)^{\delta N} \leq 0.99^{\delta N} = o(1),$$

using the hypothesis on δ.

So far, we have a graph $\mathcal{D}$ on parts $\mathcal{A}'$ and $\mathcal{B}'$ of the claimed cardinalities. Our implementation of the core process revealed only edges directed into X, so the edges in $\mathcal{D}$ remain completely random. Each node in $\mathcal{D}$ has outdegree between k and $k+3$. Reveal the outdegree of each node;

if a node has outdegree greater than k, randomly delete outedges so that it has degree exactly k. The remaining edges are still uniformly random, and $\mathcal{D}$ is precisely k-out. ☐

Claim 7. *In a uniformly random k-out digraph with parts $\mathcal{A}$ and $\mathcal{B}$ both of size $\Theta(N)$, the number of nodes with any double inedge is at most $N^{2/3}$, with failure probability of order* $\exp(-\Omega(N^{1/3}))$.

Proof. We consider double inedges to part $\mathcal{B}$; the symmetric argument applies to part $\mathcal{A}$. Call a node in $\mathcal{B}$ "bad" if it has a multiple inedge. Consider the outedges of a node $a \in \mathcal{A}$ in sequence. If an edge duplicates a previous one (and is the first to do so), it results in a bad $\mathcal{B}$ node. The expected number of such duplications on $a \in \mathcal{A}$ is $O(1/N)$, for a total expected $O(1)$ duplications over all $a \in \mathcal{A}$. Changing an outedge of any $a \in \mathcal{A}$ changes the number of bad $\mathcal{B}$ nodes by at most 1. Azuma's inequality yields the claim. ☐

Claim 8. *For any constants k and d, in a uniformly random k-out digraph with parts $\mathcal{A}$ and $\mathcal{B}$ both of size $\Theta(N)$, the probability that d random nodes in $\mathcal{B}$ have a common inneighbour in $\mathcal{A}$ is $O(1/N)$, and the symmetric condition also holds.*

Proof. If d nodes in $\mathcal{B}$ have $a \in \mathcal{A}$ as a common inneighbour, then at least two of them must be in $N^+(a)$. Since $|N^+(a)| = k$ while $|\mathcal{B}| = \Theta(N)$, this has probability $O(1/N^2)$. Taking the union bound over all $a \in \mathcal{A}$, the probability the d nodes have any common inneighbour is $O(1/N)$. ☐

Lemma 9 (expansion)**.** *For any $c, k, \delta > 0$, let $\eta_{c,k}(\delta) = \max\left\{\frac{17}{3}\delta\ln(e/\delta), 2\frac{1-c}{c}k\delta\right\}$. Let $\mathcal{D}$ be a uniformly random k-out bipartite multigraph on $\mathcal{A} \times \mathcal{B}$, with $\mathcal{A} \cup \mathcal{B} = [N]$ and $|\mathcal{A}|, |\mathcal{B}| \geq cN$. Then, a.a.s., for every $S \subseteq [N]$ with $|S| \leq \delta N$,*

$$\left|N_{\mathcal{D}}^-(S)\right| \leq \sum_{v\in S}\left|N_{\mathcal{D}}^-(v)\right| \leq \eta_{c,k}(\delta)\, N.$$

That is, any set of up to a δ fraction of the nodes expands to at most an $\eta_{c,k}(\delta)$ fraction. Note:

- The lemma refers to all subsets S of size at most δN simultaneously, not merely any given subset.
- We can make $\eta_{c,k}(\delta)$ arbitrarily small by choosing δ sufficiently small.
- The lemma implies that sets of size δN expand by a factor at most $\eta_{c,k}(\delta)/\delta$. This is false for smaller sets: there will typically be a handful of nodes of arbitrarily high indegree.
- The proof below uses the first-moment method. A slightly smaller $\eta_{c,k}(\delta)$ can be obtained by choosing (in the Poisson limit) a degree threshold above which there are at least δN nodes (thus encompassing the worst-case set S), and considering the total of those nodes' degrees (concentrated, and at least $|N^-(S)|$).

Proof. Assume, without loss of generality, that $|\mathcal{A}| \leq |\mathcal{B}|$. Let $\deg^-(v)$ denote the indegree of v. For any set $S \in [N]$, define the volume $\operatorname{vol}(S) := \sum_{v\in S}\deg^-(v)$ to be the number of edges directed into S.

For any set S, the neighbourhood size $\left|N_{\mathcal{D}}^-(S)\right|$ is at most $\operatorname{vol}(S)$; this establishes the lemma's first inequality. For the second, it is enough for prove the statement for sets S of size δN.

Let $\rho = \frac{1-c}{c}$. The outneighbours of each node are chosen uniformly with replacement from the other part. So, the indegree of a node $v \in \mathcal{A}$ is distributed as $\deg^-(v) \sim \operatorname{Bin}(k\,|\mathcal{B}|, 1/|\mathcal{A}|)$, with expectation $k\,|\mathcal{B}|/|\mathcal{A}| \leq \rho k$. For $v \in \mathcal{B}$, $\deg^-(v) \sim \operatorname{Bin}(k\,|\mathcal{A}|, 1/|\mathcal{B}|)$, again with expectation at most ρk since $|\mathcal{A}| \leq |\mathcal{B}|$. For a given set S of size δN, $Z = \operatorname{vol}(S)$ is a sum over the elements of S of binomials of these two types, all independent, so Z is given by a sum of independent random Bernoulli variables and $\mathbb{E}Z \leq \rho k\delta N$. For convenience we may artificially add a few more independent Bernoulli variables to give a similar sum Z', with Z' stochastically dominating Z and $\mathbb{E}Z' = \rho k\delta N$.

Let $\rho = \frac{1-c}{c}$ and let $\eta = \eta_{c,k}(\delta)$. Apply the Chernoff-type inequality $\Pr(Z \ge \mathbb{E}Z + t) \le \exp\left(-\frac{t^2}{2(\mathbb{E}Z+t/3)}\right)$ [15, eq. (2.5) and Theorem 2.8] to $\Pr(Z' \ge \eta N)$. This means taking $t = \eta N - \mathbb{E}Z' = (r-1)\rho k\delta N$, where $r = \eta/(\rho k\delta) \ge 2$ by definition of $\eta_{c,k}(\delta)$. Then,

$$\Pr(Z' \ge \eta N) \le \exp\left(-\frac{(r-1)^2(\mathbb{E}Z')^2}{2(1+\frac{r-1}{3})\mathbb{E}Z'}\right) \le \exp\left(-\frac{(r/2)^2}{2(\frac{2r}{3})}\mathbb{E}Z'\right) = \exp\Big(-\tfrac{3}{16}r\cdot\rho\delta kN\Big).$$

The number of choices for the set S is

$$\binom{N}{\delta N} \le \left(\frac{eN}{\delta N}\right)^{\delta N} = \exp(\delta N \ln(e/\delta)).$$

By the first-moment method, the probability that any set S violates the claimed condition is at most

$$\begin{aligned}\binom{N}{\delta N}\Pr(Z' \ge \eta N) &\le \exp\left(\delta N\cdot\left(\ln(e/\delta) - \frac{3}{16}r\rho k\right)\right)\\ &= \exp\left(\delta N\cdot\left(\ln(e/\delta) - \frac{3}{16}\frac{\eta}{\delta}\right)\right) = \exp(-\Omega(N)) = o(1),\end{aligned}$$

the penultimate equality by definition of $\eta_{c,k}(\delta)$. ☐

Lemma 10 (Hall)**.** *Given $k \ge 10$ and $\rho < 1 - 2/(1+3/e)$. (Having $\rho \le 0.049$ suffices.) Let $\mathcal{D}$ be a uniformly random k-out bipartite digraph on parts $\mathcal{A}$ and $\mathcal{B}$. Consider the subgraph $\mathcal{D}'$ induced by a pair of subsets $\mathcal{A}' \subseteq \mathcal{A}$, $\mathcal{B}' \subseteq \mathcal{B}$, where*

(i) $|\mathcal{A}'| + |\mathcal{B}'| \ge (1-\rho)(|\mathcal{A}| + |\mathcal{B}|)$,
(ii) $|\mathcal{A}'| = |\mathcal{B}'|$; and
(iii) every node in $\mathcal{D}'$ has outdegree at least $k-1$.

Then, a.a.s. in $|\mathcal{A}|$, every such induced subgraph $\mathcal{D}'$ contains a perfect matching.

Note that the asymptotic statement in terms of $|\mathcal{A}|$ is equivalent to one in terms of $|\mathcal{B}|$, $|\mathcal{A}'|$, or $|\mathcal{B}'|$. The hypothesis $k \ge 10$ can be weakened to $k \ge 7$, at the expense of a slightly more complicated proof, and likely further.

Proof. Let $n := |\mathcal{A}'| = |\mathcal{B}'|$. If $\mathcal{D}'$ does not have a perfect matching (ignoring the edge directions) then there are Hall sets $S \subseteq \mathcal{A}'$ and $T \subseteq \mathcal{B}'$ such that $N(S) \subseteq T$ and $|T| = |S| - 1$. Write $\bar{S}$ for $\mathcal{A}' \setminus S$ and $\bar{T}$ for $\mathcal{B}' \setminus T$. If S, T is a Hall pair, so is $\bar{T}, \bar{S}$. Since $|S| + |\bar{T}| = n+1$, at least one of them must be at most $(n+1)/2$. By symmetry, it suffices to consider a Hall pair S, T where $S \subseteq \mathcal{A}$ has cardinality $s \le (n+1)/2$.

Henceforth we ignore $\mathcal{A}'$ and $\mathcal{B}'$. We are simply interested in the existence of a Hall pair S, T with $S \subseteq \mathcal{A}$, $T \subseteq \mathcal{B}$, $|S| = s \le (n+1)/2$, and $|T| = s-1$. The probability that $\mathcal{D}$ allows such a pair is at most the expected number of such pairs, which is at most the number of pairs times the probability that a fixed pair has the Hall property.

Since $2n = |\mathcal{A}'| + |\mathcal{B}'| \ge (1-\rho)(|\mathcal{A}| + |\mathcal{B}|)$, we have that $|\mathcal{A}| + |\mathcal{B}| \le \frac{2}{1-\rho}n$ and thus

$$|\mathcal{A}| \le \tfrac{2}{1-\rho}n - |\mathcal{B}| \le \tfrac{2}{1-\rho}n - |\mathcal{B}'| = \left(\tfrac{2}{1-\rho} - 1\right)n \le \tfrac{3}{e}n;$$

the last inequality follows from the hypothesis on ρ and is used in the next calculation. The same holds for $\mathcal{B}'$. Also, $|\mathcal{A}|$ and $|\mathcal{B}|$ are both $\Theta(n)$, and henceforth we work asymptotically in n.

For each $v \in S$, all but one of v's outedges must be directed into T. The probability that there is such a Hall pair S, T, then, is at most the number of choices of $S \subseteq \mathcal{A}$ with $|S| = s$, of $T \subseteq \mathcal{B}$ with $|T| = s-1$, and (by hypothesis (iii)) of one edge per element of S, times the probability that such a combination has all other edges of S directed into T. The probability there is any such pair is at most

$$\begin{aligned}\binom{k}{1}^s \binom{(3/e)n}{s}\binom{(3/e)n}{s-1}\left(\frac{s-1}{n}\right)^{(k-1)s} &\le k^s\left(\frac{3n}{s}\right)^s\left(\frac{3n}{s-1}\right)^{s-1}\left(\frac{s-1}{n}\right)^{(k-1)s}\\ &\le (9k)^s\left(\frac{s-1}{n}\right)^{(k-3)s+1}\\ &\le \frac{s}{n}\left(9k\left(\frac{s-1}{n}\right)^{k-3}\right)^s. \qquad (3.1)\end{aligned}$$

The sum of (3.1) over s from 2 to $10\ln n$, even ignoring the power s, is $O(\ln n)\cdot O\left(\left(\frac{\ln n}{n}\right)^{k-2}\right) = o(1)$. For larger s, up to $s=(n+1)/2$, we have $(s-1)/n<1/2$, so, since $k\ge 10$, $9k(\frac{s-1}{n})^{k-3} \le 9k(\frac{1}{2})^{k-3} < 0.8 < \exp(-1/5)$. Each term in (3.1) is thus of order $O(e^{-s/5}) = O(e^{-2\ln n}) = O(n^{-2})$, for a total of $O(1/n)=o(1)$. □

4. Matching Strategy and Analysis

We construct a hypergraph H, initially empty and eventually containing a perfect matching $\mathcal{M}$, on vertex set $V=[n]$; we assume that n is divisible by 3. Partition V into a set A_0 of vertices called "apexes", and a set B_0 of pairs of vertices, called "base pairs" or just "bases", with

$$|A_0| = \frac{n}{3} + 4\epsilon_D n \qquad\qquad |B_0| = \frac{n}{3} - 2\epsilon_D n, \qquad (4.1)$$

where ϵ_D is a constant chosen small enough to satisfy various claims below. The two members of a base pair are sometimes called "partners". Integrality is not an issue; all that needs to be exact is that $|A_0| + 2|B_0| = n$.

4.1. Phase 1: A robust matching structure.

In this phase we construct the digraph D_1 introduced above, and in tandem the corresponding hypergraph H.

For a sequence of semirandom offers, do the following. For each offer of the form $\{a,b\}$, where $a\in A_0$ and $b\in \mathbf{b}\in B_0$ is a member of a base pair, choose as third point the partner of b. This defines a hyperedge $\{a\}\cup\mathbf{b}$, which we add to our hypergraph H, and an (undirected) auxiliary graph edge $\{a,\mathbf{b}\}$, which we add to a set Ψ. Semi-random offers not of the specified form are ignored. (Recall what this means from the beginning of Section 2.)

Note that each $(a,\mathbf{b})\in\Psi$ is uniformly random in $A_0\times B_0$, with replacement. Given ϵ_D and k, let C be the corresponding value in Lemma 6. Each semirandom offer is of the desired form w.p. (with probability) $\frac{1}{3}\cdot\frac{2}{3}+\frac{2}{3}\cdot\frac{1}{3}+O(\epsilon_D) = 4/9+O(\epsilon_D) > 1/3$ for ϵ_D sufficiently small, so in $3Cn$ semirandom offers, a.a.s. there are at least Cn of the stipulated form, exceeding the lemma's CN (since by (4.1), N is about $\frac{2}{3}n$).

Should this a.a.s. statement fail, we consider the entire construction to have failed. We will take the same approach throughout, proceeding on the assumption that all the a.a.s. statements hold.

It follows from Lemma 6 that from this set Ψ of undirected edges we can a.a.s. construct a bipartite digraph D_1 on parts $A_1\subseteq A_0$ and $B_1\subseteq B_0$, with $|A_1|\ge(1-\epsilon_D)|A_0|$ and $|B_1|\ge (1-\epsilon_D)|B_0|$, and D_1 is uniformly k-out on these node sets.

By construction, H contains the hyperedges corresponding to the edges of D_1. We will discard (or ignore) all other hyperedges in H so that there is a one-to-one correspondence between the edges in D_1 and the hyperedges in H: they are two representations of the same thing.

Let $X_1 = (A_0 \setminus A_1) \cup (B_0 \setminus B_1)$ be the set of "failed" nodes missing from D_1. By Lemma 6 and (4.1), we have

$$|A_1| \geq (1 - \epsilon_D)\,|A_0| > \frac{n}{3} + 3\epsilon_D n \tag{4.2}$$

$$|B_1| \geq (1 - \epsilon_D)\,|B_0| > \frac{n}{3} - 3\epsilon_D n \tag{4.3}$$

$$|X_1| \leq \epsilon_D(|A_0| + |B_0|) < \epsilon_D n. \tag{4.4}$$

Let $\dot{X}_1$ be the set of all the points (hypergraph vertices) contained in X_1, so $\dot{X}_1 = (A_0 \cap X_1) \cup \bigcup_{\mathbf{b} \in (B_0 \cap X_1)} \mathbf{b}$. Letting $a_x = |A_0 \cap X_1|$ and $b_x = |B_0 \cap X_1|$, note that

$$|\dot{X}_1| = a_x + 2b_x. \tag{4.5}$$

4.2. **Phase 2: Matching exceptional points.** This phase will place every point of $\dot{X}_1$ into a hyperedge to be included in the matching $\mathcal{M}$, while keeping D robust enough that it will contain a matching.

At the beginning of this phase, $A = A_1$, $B = B_1$, $D = D_1$, $X = X_1$, and $\dot{X} = \dot{X}_1$. Along the way, $A, B, D, X, \dot{X}$ will change; e.g., points will be removed from $\dot{X}$ as they are placed in hyperedges of $\mathcal{M}$.

We first consider the operations we will perform in this phase, deferring the question of how many semirandom offers are required.

4.2.1. *Phase 2a.* Here we treat all the points in X_1, from apexes and base pairs alike. We will be more precise in a moment, but roughly speaking, any time we get a semirandom offer consisting of two points $a, a' \in A$, we choose an arbitrary $x \in \dot{X}$, defining a hyperedge $e = \{a, a', x\}$, and add e irrevocably to $\mathcal{M}$. The failed vertex x is now in a hyperedge in $\mathcal{M}$, but a and a' must be deleted from A and D: neither is any longer available to be combined with a base pair.

As said in the Outline (Section 3), we will ensure that, as we delete nodes from D, every node in D retains outdegree at least $k - 1$. To do this, we "block" some nodes of D from future deletion; we let Q denote the set of blocked nodes. To avoid that any single node's deletion decreases an inneighbour's outdegree by 2 or more, we intialise Q to be the set of nodes with parallel inedges. To avoid that the outdegree of any node u is decreased more than once, if ever an outneighbour v of u is deleted, we block all other outneighbours of u from future deletion. That is, whenever a node v is deleted from D, we add $N^+(N^-(v))$ to Q. All that remains is to ensure that in any single action, we do not delete two nodes with a common inneighbour.

With this out of the way, we can be precise. For any semirandom offer consisting of two points $a, a' \in A_0$, we choose an arbitrary $x \in \dot{X}$, defining a hyperedge $e = \{a, a', x\}$. We discard e unless the following conditions both hold:

- (M2a-C1) $a, a' \in A \setminus Q$.
- (M2a-C2) $N^-(a)$ and $N^-(a')$ are disjoint.

(Read the label as M for matching, 2a for the phase, and C for "condition", and in the next part, A for "action".)

If the conditions hold, we take the following action:

- (M2a-A1) Add e, irrevocably, to $\mathcal{M}$.
- (M2a-A2) Delete x from $\dot{X}$ and a and a' from A (and hence also from D).
- (M2a-A3) Add $N^+(N^-(\{a, a'\}))$ to Q, blocking the outneighbours of the inneighbours of the deleted nodes.

Each such action resolves one point in $\dot{X}$, so by (4.5), a_x+2b_x actions are required. Each action deletes 2 nodes of part A and 0 nodes of part B, so at the end of this phase, $|B|=b_0-b_x$ and $|A|=(a_0-a_x)-2(a_x+2b_x)=a_0-3a_x-4b_x$. Also,

$$|A|-|B| = a_0-b_0-3a_x-3b_x \;=\; a_0-b_0-3\,|X_1| \tag{4.6}$$

$$\geq a_0-b_0-3\epsilon_D n \;=3\epsilon_D n>0 \qquad \text{(by (4.4) and (4.1))}. \tag{4.7}$$

Thus, A is larger than B at the end of Phase 2a.

4.2.2. *Phase 2b.* This phase makes the sizes of A and B exactly equal so that they can be matched in D. Recall that at the end of Phase 2a, $|A|>|B|$. To even them up, we will generate hyperedges each consisting of three vertices from A, deleting the points from A and adding the hyperedge to $\mathcal{M}$.

To be precise, each time we get a semirandom offer consisting of two vertices $a,a'\in A_0$, we define a hyperedge $e=\{a,a',a''\}$ for a uniformly random $a''\in A\setminus\{a,a'\}$. We discard the offer unless the following conditions both hold:

(M2b-C1) $a,a',a''\in A\setminus Q$.
(M2b-C2) $N^-(a)$, $N^-(a')$, and $N^-(a'')$ are disjoint.

If the conditions hold, we take the following action:

(M2b-A1) Add e, irrevocably, to $\mathcal{M}$.
(M2b-A2) Delete a,a', and a'' from A (and hence also from D).
(M2b-A3) Add $N^+(N^-(\{a,a',a''\}))$ to Q, blocking the outneighbours of the inneighbours of the deleted nodes.

From (4.6), at the beginning of Phase 2b we have $|A|-|B|\equiv a_0-b_0\equiv a_0+2b_0=n\equiv 0 \pmod 3$. Since each action deletes 3 vertices from A, we can make $|A|=|B|$ exactly. The precise number of actions needed is not critical — it is obviously $O(\epsilon_D n)$ — but from (4.6) and (4.1) it is

$$\frac{|A|-|B|}{3}=\frac{a_0-b_0}{3}-|X_1|\leq\frac{a_0-b_0}{3}=2\epsilon_D n. \tag{4.8}$$

4.2.3. *Analysis and parameter choices.* How many nodes can ever be blocked? Initially, $Q=Q_1$ consists of the nodes with double inedges. By Claim 7, $|Q_1|=O(n^{2/3})$, with failure probability exponentially small in n. As usual, we will declare failure of the construction if this fails.

Recall that when an action deletes a node v from D, it adds $N^+(N^-(v))$ to Q. Phase 2a takes a_x+2b_x actions, each deleting 2 nodes from D; note that $a_x+2b_x\leq 2|X_1|\leq 2\epsilon_D n$, by (4.4). Phase 2b takes at most $2\epsilon_D n$ actions, each deleting 3 nodes from D. Altogether, they delete at most $10\epsilon_D n$ nodes from D.

The number of inneighbours of these nodes is bounded by Lemma 9, which will be our focus for the next few paragraphs. The lemma's N is $|A_1|+|B_1|>\frac{2}{3}n$ by (4.2) and (4.3), so $10\epsilon_D n$ represents at most a $15\epsilon_D$ fraction of N. Thus we take the lemma's δ — call it δ_D — equal to $15\epsilon_D$.

The two parts A_1 and B_1 are of nearly equal size, so we may take Lemma 9's c as $\frac{1}{3}$. We will make many assertions of this sort, so let us be more precise this one time. In this case, from (4.1), (4.2), and (4.3) we have $(\frac{1}{3}+3\epsilon_D)n\leq|A_1|\leq|A_0|=(\frac{1}{3}+4\epsilon_D)n$ and $(\frac{1}{3}-3\epsilon_D)n\leq|B_1|\leq|B_0|=(\frac{1}{3}-2\epsilon_D)n$. We will choose ϵ_D small enough that it is obvious that each of A_1 and B_1 has size at least 1/3rd that of their union. More to the point, forgetting the precise multiples of ϵ_D here, we will have expressions like $|A_1|=(\frac{1}{3}+O(\epsilon_D))n$, and it will be clear that ϵ_D *can* be chosen sufficiently small to satisfy the various assertions.

Continuing, by Lemma 9, the number of inneighbours is at most $\eta_{1/3,k}(\delta_D)N$. Each inneighbour has k outneighbours, so the number of outneighbours is at most $k\,\eta_{1/3,k}(\delta_D)N$. Since $|Q_1|=o(n)=$

$o(N)$, we have that at all times $|Q| \le k\,\eta_{1/3,k}(\delta)N + o(N)$. We will contrive that $k\,\eta_{1/3,k}(\delta_D) \le \frac{1}{10}$, so that $|Q| \le \frac{N}{9}$ and $|Q| \,/\, |A_1| \le (\frac{N}{9})/(\frac{N}{3}) \le \frac{1}{3}$ and symmetrically $|Q| \,/\, |B_1| \le \frac{1}{3}$.

This dictates our parameter choices. We take $k = 10$ throughout, for purposes of Lemma 10. Then, in Lemma 9, choosing $\delta = \delta_D = 0.000\,15$ suffices, giving $k\,\eta_{1/3,k}(\delta_D) \le \frac{1}{10}$ as desired. (The value of η, defined in Lemma 9, is determined by the first term in the max; the second term, $40\,\delta$, is smaller.) Since we set $\delta_D = 15\epsilon_D$, taking $\epsilon_D = 0.000\,01$ suffices.

Recall that Phase 1 used $3Cn$ semirandom offers, with C given by Lemma 6. Now that we have fixed $\epsilon_D = 0.000\,01$, the proof of Lemma 6 fixes its $\delta = \epsilon_D/2$ (the second term in the min is much larger). In turn, to give $\Pr(\mathrm{Po}(\lambda) < k+3) = \delta/4 = \epsilon_D/8 = 0.000\,001\,25$, it fixes λ to be between 37 and 38, for $C = 2.1\lambda < 80$. Thus, $3Cn < 240n$ semirandom offers suffice for Phase 1.

Looking ahead, we will invoke Lemma 10 to prove existence of a perfect matching in D. The lemma's $\mathcal{A}$ and $\mathcal{B}$ are the parts A_1 and B_1 at the start of Phase 2, its $\mathcal{A}'$ and $\mathcal{B}'$ are the parts A and B at the end of Phase 2, and its ρ is (an upper bound on) the fraction of nodes deleted, so we take $\rho = \delta$. With the choices above, then, $\rho \le 0.049$, satisfying the hypothesis.

Having controlled the size of Q, we are ready to consider the number of semirandom offers needed in Phase 2.

In Phase 2a, that the first point offered is in A_0 occurs with probability just above 1/3. The cardinality of A is nearly that of A_0 and we have arranged that $|Q| \le \frac{1}{3}\,|A_1|$, so

$$\Pr(a \in A \setminus Q) \ge \tfrac{1}{3} \cdot (1 - O(\epsilon_D) - \tfrac{1}{3}) \;>\; \tfrac{1}{5}. \tag{4.9}$$

Independently (almost), the same holds for a', so (M2a-C1) holds with probability at least 1/25.[1] By Claim 8, the chance that (M2a-C2) fails to hold is $O(1/n)$. Thus, each offer satisfies the conditions with probability at least 1/26. Since at most $2\epsilon_D n$ actions are needed (see (4.4)), a.a.s., $27 \cdot 2\epsilon_D n$ offers suffice.

Similarly, in Phase 2b, an offer plus the added point a'' satisfies condition (M2b-C1) with probability more than $(1/5)^3$ (even if we chose a'' in $[n]$ rather than A), and again (M2b-C2) fails with probability $O(1/n)$, so each offer results in action with probability at least 1/126. By (4.8), the number of actions needed is at most $2\epsilon_D n$, so, a.a.s., $127 \cdot 2\epsilon_D n$ offers suffice.

It is clear that, with control of $|Q|$, each offer in Phases 2a and 2b results in action with probability $\Theta(1)$, and thus $O(n)$ offers suffice. Numerically, Phases 2a and 2b together require at most $(27 + 127) \cdot 2\epsilon_D < 0.01n$ offers, so $241n$ offers suffice for Phases 1 and 2 (thus for the whole algorithm). We have made no effort to optimise the constants.

We use A_2, B_2, D_2 to denote A, B, D at the end of Phase 2. At this point, all points outside of A_2, B_2 appear in hyperedges of $\mathcal{M}$. Since no nodes were deleted from B during Phase 2, and applying (4.3), we have:

$$|A_2| = |B_2| = |B_1| > \frac{n}{3} - 3\epsilon_D n. \tag{4.10}$$

4.3. **Phase 3: Matching the bulk of the points.** We complete the hypergraph matching $\mathcal{M}$ with hyperedges corresponding to the edges of a perfect matching in D_2; Lemma 10 proves that a.a.s. D_2 has a perfect matching.

We apply the lemma with $\mathcal{D} = D_1$ and $\mathcal{D}' = D_2$, and accordingly $\mathcal{A} = A_1$, $\mathcal{A}' = A_2$, and $\mathcal{B} = \mathcal{B}' = B_1 = B_2$.

[1] We would have independence if the offered a' were independent of a, allowing the possibility that $a' = a$ (in which case we would reject the offer). We can simulate this independence by tossing a coin and with probability $1/n$ replacing the offered a' with a. In this simulated model, (M2a-C1) and distinctness hold with probability at least $\Pr(a \in A \setminus Q)^2 - 1/n$, here simply 1/25 by absorbing the $1/n$ into slack in (4.9). Since the simulated substitution of a' for a is never advantageous, (M2a-C1) holds with probability at least 1/25 in the original model.

We already noted in Section 4.2.3 (and confirmed in (4.10)) that hypothesis (i) is satisfied: the fractional size difference between A_1 and A_2 is $O(\epsilon_D)$ and far less than $1/100$. Hypothesis (ii), that the two parts are of equal cardinality, is satisfied by construction; see just before (4.8). Hypothesis (iii), that every node degree is k or $k-1$, is also satisfied by construction, specifically the various action conditions based on the blocked set Q. Thus, a.a.s. D_2 contains a perfect matching.

This concludes the proof of Theorem 1.

4.4. **Extension to s-uniform hypergraphs.** We now establish Theorem 2. As already explained, constructing a perfect matching in the 1-offer s-uniform model can trivially be simulated in the 2-offer s-uniform model, simulating a 2nd "offer" vertex with the extra "chosen" vertex.

Constructing a matching in the 2-offer s-uniform model cannot be done using the 2-offer 3-uniform model as a black box, but can be done in exactly the same way. In short, in Phase 0, take a bit more than n/s vertices as apexes A_0, arbitrarily partitioning the rest into a bit fewer than n/s "base" groups each of size $s-1$, comprising B_0.

4.4.1. *Phase 1.* In Phase 1, when offered an apex a and a point b from a base $\mathbf{b}$, propose the hyperedge $a \cup \mathbf{b}$ and correspondingly an auxiliary-graph edge $\{a, \mathbf{b}\}$, directed one way or the other. As before, with such proposals we can construct a directed auxiliary graph $D = D_1$, bipartite with parts A_1 and B_1 comprising most of A_0 and B_0 respectively, with D_1 a uniformly random k-out digraph on A_1 and B_1. Let $X = X_1 = (A_0 \cup B_0) \setminus (A_1 \cup B_1)$ be the nodes missing from D, and $\dot{X}$ the set of points in X.

4.4.2. *Phase 2.* In Phase 2a, until $\dot{X}$ is empty, when offered two points in A_0, propose a hyperedge consisting of them and any $s-2$ points in $\dot{X}$. If in a final step $0 < |\dot{X}| < s-2$, supplement the points of $\dot{X}$ with distinct unblocked points from A. Correspondingly, propose the action of deleting the two apexes from part A of the auxiliary graph D.

As we did for matchings, we block $N^+(N^-(v))$ for any node v deleted from D. We accept a proposed action (and accept the corresponding hyperedge for $\mathcal{M}$) if the offered vertices lie in the current set A and are not blocked. By analogy with the argument around (4.9), by controlling the size of $\dot{X}_1$ we can ensure that the number of nodes deleted and blocked is small, ensuring in turn that each proposed action is accepted w.p. at least $(1/s \cdot 0.98)^{s-2} = \Omega(1)$. Also as for matchings, we choose the sizes $|A_0|$ and $|B_0|$ so that, a.a.s., when Phase 2a ends, $|A| \geq |B|$.

In Phase 2b, when offered 2 vertices $a, a' \in A_0$, propose a hyperedge consisting of them and $s-2$ distinct points in $A \setminus \{a, a'\}$ and correspondingly propose the action of deleting these s points from part A of D. Accept the proposal if the two offered points were in $A \setminus Q$; in this case, delete these s apexes from D, and accept the hyperedge for $\mathcal{M}$.

Again, each proposed action is accepted w.p. at least $(1/s \cdot 0.98)^{s-2} = \Omega(1)$. Stop when the number of apexes is equal to the number of bases, $|A| = |B|$.

For the stopping condition, since at the end of Phase 2a, $|A| \geq |B|$, and Phase 2b decreases the size of A without changing B, it is clear that at some time we get $|A| \leq |B|$. At this time, we get equality if $|A| \equiv |B| \pmod{s}$, and this is so. Let $V = [n]$ be the set of all points, with $|V| \equiv 0 \pmod{s}$. Let $V(\dot{D})$ be the set of points in D. At this step, the points of $V(\dot{D})$ are all in hyperedges in $\mathcal{M}$, so $\left|V \setminus V(\dot{D})\right| \equiv 0 \pmod{s}$. Thus, at this step,

$$|A| - |B| \equiv |A| + (s-1)\,|B| = \left|V(\dot{D})\right| = |V| - \left|V \setminus V(\dot{D})\right| \equiv 0 \pmod{s}. \tag{4.11}$$

4.4.3. *Phase 3.* Finally, construct a matching in D (ignoring the edge directions). The corresponding hypergraph edges complete a perfect matching in the hypergraph.

As for matchings, it is enough to ensure that when constructing D_1, $|X| \leq \epsilon_D n$ for some sufficiently small constant ϵ_D. The number of Phase 2 steps is $O(\epsilon_D n)$ (the $O(\cdot)$ expression hiding some constant independent of ϵ_D), so the number of nodes deleted from D is also $O(\epsilon_D n)$. By Lemma 9, Q remains of size at most $k\,\eta_{c,k}(\delta)(O(\epsilon_D))n$ (plus $O(n^{2/3})$), and this can be made an arbitrarily small fraction of n by choosing ϵ_D sufficiently small. If follows that each action is accepted w.p. $\Omega(1)$.

5. Loose Hamilton cycles

Now we turn to the proof of Theorem 3. We will demonstrate a strategy that, within a linear number of steps, a.a.s. produces a 3-uniform hypergraph with a loose Hamilton cycle. The vertex set is $[n]$ for some even n. (The number of vertices must be even if a loose Hamilton cycle exists.) The odd points will occur twice each in the loose cycle, the even points just once, in the form

$$(1, a_1, 3),\ (3, a_2, 5),\ \ldots,\ (n-1, a_{n/2}, 1), \tag{5.1}$$

where we write each edge as an ordered triple to indicate the roles of the three points, and $a_1, a_2, \ldots, a_{n/2}$ is a permutation of $2, 4, \ldots, n$. We will not be able to get exactly such a cycle, with the odd numbers appearing in order, but will obtain a Hamilton cycle that differs from it on very few edges.

The approach is similar to that for generating a perfect matching in Section 4. We define apex and base sets, each of cardinality exactly $n/2$:

$$A_0 = \{2, 4, \ldots, n\} \qquad\qquad B_0 = \{\{1,3\}, \{3,5\}, \ldots \{n-1, 1\}\}. \tag{5.2}$$

Here the bases are overlapping, in order to build a cycle like that of (5.1), where for hypergraph matchings they were disjoint.

The next steps are structurally the same as for matchings in Section 4. We will construct a directed bipartite graph on A_0 and B_0, but it will not have a perfect matching because a small number of nodes fail to have sufficiently high degree (and cause other nodes to "fail" as well). We deal with the failed nodes first, and in doing so add certain hyperedges, irrevocably, to our cycle-in-the-making $\mathcal{C}$. ($\mathcal{C}$ plays the role that $\mathcal{M}$ did for hypergraph matchings.) When the failed nodes are dealt with, most of our bipartite graph will remain, and what remains will have a perfect matching. The matching edges correspond to hyperedges in $\mathcal{H}$, and they complete $\mathcal{C}$ to a loose Hamilton cycle.

5.1. **Phase 1.** In the first phase, we build a graph D in the same manner as we did when building a perfect matching in Section 4.

For each offer of the form $\{a, b\}$, where $a \in A$ and $b \in \mathbf{b} \in B$ is a member of a base pair, choose as third point the partner of b and add $\{a, \mathbf{b}\}$ to a set Ψ. (We ignore any offers that are not of this form.)

Ψ forms a set of edges of a bipartite graph on $A_0 \times B_0$. Note that a perfect matching on this graph would yield a loose Hamilton cycle of the form in (5.1).

For any $\epsilon_D > 0$, Lemma 6 ensures, exactly as in Phase 1 of Section 4, that within $O(n)$ steps we can construct a directed bipartite graph $D = D_1$ with parts $A_1 \subseteq A_0$ and $B_1 \subseteq B_0$, such that D_1 is a uniformly random k-out bipartite multigraph on $A_1 \times B_1$, and a.a.s. $|A_1| \geq (1-\epsilon_D)\,|A_0|$ and $|B_1| \geq (1-\epsilon_D)\,|B_0|$. We refer to the apexes in A_1 as *good* and those in $A_X = A_0 \setminus A_1$ as *failed*, and likewise for bases. We have

$$a_x = |A_X| \leq \epsilon_D n/2 \qquad\qquad b_x = |B_X| \leq \epsilon_D n/2. \tag{5.3}$$

5.2. **Phase 2.** We introduce an additional structure, $\mathcal{P}$, central to the analysis. $\mathcal{P}$ is a graph on the odd points, consisting of vertex-disjoint paths. In particular, a path in $\mathcal{P}$ can be an isolated vertex. Initially, the edges in $\mathcal{P}$ are precisely the base pairs B_1. Eventually we will make $\mathcal{P}$ a single path and then a Hamilton cycle. Always, each edge in $\mathcal{P}$ is one of two types: it consists of the two odd points in a hyperedge $e = (b_1, a, b_2) \in \mathcal{H}$ (and all the apexes used for such edges will be distinct); or, as in the initial $\mathcal{P}$, it consists of a base pair in B (and later it will be possible to perfectly match all these base pairs and the remaining apexes, in D). When $\mathcal{P}$ is a Hamilton cycle, this ensures that there is a corresponding loose hypergraph cycle $\mathcal{C}$ as desired.

In this phase we will again use a set Q of nodes of D "blocked" from use. Initially we will set $Q = Q_0$ to consist of any nodes whose inedges include any double edge; recall from Claim 7 that $|Q_0| = o(n)$. Initially, $\mathcal{P}$ has b_x components. Let us describe the phase just enough to count the number of nodes it will delete from each part of D, both to bound $|Q|$ and to confirm that, going in to Phase 3, $|A| = |B|$.

In Phase 2a, each action (see Figure 1) will assign an apex in A_X to a hyperedge (along with two base points from different base pairs) and place it in $\mathcal{C}$, increasing the number of paths in $\mathcal{P}$ by 1, and deleting 2 base pairs from B. Phase 2a takes a_x actions to resolve all the failed apexes, resulting in a total of $2a_x$ deletions from the set of base pairs in B, so that $|B_0 \setminus B| = b_x + 2a_x$, and $b_x + a_x$ components in $\mathcal{P}$. The part A is unchanged, with $|A_0 \setminus A| = a_x$.

In Phase 2b, each action will decrease the number of components of $\mathcal{P}$ by 1 (see Figure 2), deleting 2 apexes from D (committing them to hyperedges) and deleting 1 base pair from B. This phase takes $a_x + b_x$ actions to make $\mathcal{P}$ a cycle, in the process making $|A_0 \setminus A| = a_x + 2(a_x + b_x)$ and $|B_0 \setminus B| = (b_x + 2a_x) + (a_x + b_x)$, both equal to $3a_x + 2b_x$.

In the end, then, we have equal numbers of apexes and base pairs not in D, and thus also equal numbers remaining in D.

The number of apex nodes deleted from D by the two phases is $0 + 2(a_x + b_x) = 2a_x + 2b_x$, and the number of base nodes deleted is $2a_x + (a_x + b_x) = 3a_x + b_x$; by (5.3) each is at most $2\epsilon_D n$.

We can make the blocked set Q arbitrarily small by choosing ϵ_D sufficiently small, as we now show. For any given $\epsilon_Q > 0$, let Lemma 9's η equal ϵ_Q/k, and take the lemma's corresponding δ to determine $4\epsilon_D$ (choosing ϵ_D smaller if required elsewhere). Let S be the set of deleted nodes. Then, Lemma 9 ensures that the deleted set's inneighbourhood has size $|N^-(S)| \le \eta N \le (\epsilon_Q/k)n$. Since D is k-out, the inneighbourhood's outneighbourhood — which is to say, the rest of Q beyond Q_0 — has size

$$|Q \setminus Q_0| = \left|N^+(N^-(S))\right| \le k(\epsilon_Q/k)n = \epsilon_Q n.$$

5.2.1. *Phase 2a.* First we treat the failed apex nodes, those in A_X, doing something analogous to Phase 2a for matchings. In Phase 2b we will treat the failed bases, which has some extra complexity.

Upon offer of odd points b_1 and b_2, we propose a hyperedge $e = (b_1, a, b_2)$, where a is an arbitrary apex in A_X. If b_1 and b_2 belong to a common path $P \in \mathcal{P}$ then, on the induced path between them, let b_1' be the neighbour of b_1 and b_2' that of b_2. If not, let $b_1' = b_1 + 2$ and $b_2' = b_2 + 2$.

We discard the hyperedge unless the following conditions all hold:

- (H2a-C1) (b_1, b_2) is not in $\mathcal{P}$ (therefore not in B).
- (H2a-C2) (b_1, b_1') and (b_2, b_2') are in $B \setminus Q$.
- (H2a-C3) $N^-(b_1) \cap N^-(b_2) = \emptyset$.

(The labelling here is H to connote Hamilton cycle, 2a the phase, and C a condition, with A for action in the next group.) Several of these conditions are unnecessary, but it is probable that they all hold (as will be shown),

If the conditions all hold, take the following action (see Figure 1).

- (H2a-A1) Add (b_1, a, b_2) irrevocably to $\mathcal{C}$, and add (b_1, b_2) to $\mathcal{P}$.
- (H2a-A2) Delete (b_1, b_1') and (b_2, b_2') from $\mathcal{P}$, B, and D. Also, delete a from A_X.
- (H2a-A3) Add $N^+(N^-(\{(b_1, b_1'), (b_2, b_2')\}))$ to Q, blocking the outneighbours of the inneighbours of the deleted nodes.

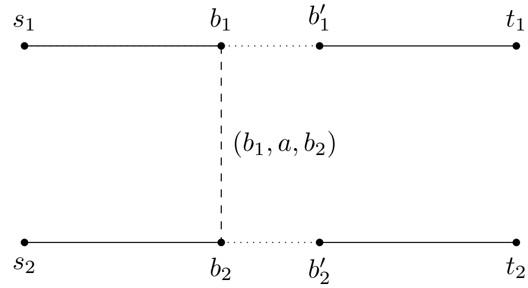

FIGURE 1. Salvaging a "failed" apex a. Upon offer of suitable odd points b_1 and b_2, hyperedge (b_1, a, b_2) is introduced (dashed line) and added to $\mathcal{C}$. The good base pairs $\{b_1, b'_1\}$ and $\{b_2, b'_2\}$ (dotted lines) are deleted from D. (As a failed apex, a was already absent from D.) This converts two paths in $\mathcal{P}$ (the figure's upper and lower lines) into three, or one path into two if b_1 and b_2 lie on a common path (if in the figure, $t_1 = t_2$).

The action incorporates the failed apex a into a hyperedge, at the expense of removing two good base pairs (b_1, b'_1) and (b_2, b'_2) from D, and turning two paths in $\mathcal{P}$ into three (as shown in the figure, when b_1 and b_2 are in distinct paths) or one into two (when b_1 and b_2 are in a common path). The action decreases the outdegree of any node in D by at most 1: it is only (b_1, b'_1) and (b_2, b'_2) that are deleted, by condition (H2a-C2) neither was in Q therefore neither has a double inedge, and by condition (H2a-C3) they have no common inneighbour.

In any trial, the conditions are likely to be satisfied. For (H2a-C1), there is probability $O(1/n)$ that b_1 and b_2 are path neighbours. For (H2a-C2), $\mathcal{P}$ has $(1/2 - O(\epsilon_D))n$ edges on its $n/2$ points, so most points have two edges in $\mathcal{P}$ (so it is likely b_1 and b_2 do) while only $O(\epsilon_Q)n$ points are in Q (so it is unlikely that b_1 or b_2 are). Finally, (H2a-C3) follows from Claim 8. Since all of the handful of failure events have probability $O(\epsilon_D + \epsilon_Q)$, so does their union.

5.2.2. *Phase 2b.* We now turn to the $O(\epsilon_D)n$ failed bases, which has some extra complexity. Details will follow but the basic idea is illustrated in Figure 2. If there were a previous semirandom offer of (b, a) accepted as hyperedge (b, a, s_1), with s_1 an endpoint of a path $P_1 \in \mathcal{P}$ whose other endpoint is t_1, then if there is a new offer (b', a'), where (b, b') is an edge in a path P in $\mathcal{P}$, accepting the new offer as (b', a', t_1) allows the paths P and P_1 to be merged. In this case, the hyperedges (b, a, s_1) and (b', a', t_1) are both added irrevocably to $\mathcal{C}$ and (b, b') is deleted from P, B, and D.

The operation works equally well whether b' lies between b and t, as shown, or between b and s. If $\mathcal{P}$ consists of a single path, $s, \ldots, b, b', \ldots, t$, then a similar operation, using hyperedges (b, a, t) and (b', a', s), turns $\mathcal{P}$ into a cycle.

This merging phase proceeds in rounds, each starting with $\mathcal{P}$ having ℓ components and reducing that to $\lfloor (9/10)\ell \rfloor$ components, until the final round where $\mathcal{P}$ goes from a Hamilton path to a Hamilton cycle. If any round fails, we declare failure of the whole algorithm; we will show this to be unlikely.

We partition the $n/2$ odd points into $\lfloor n/4 \rfloor$ disjoint "cell" pairs $\{1, 3\}, \{5, 7\}, \{9, 11\}, \ldots$. (If $n \equiv 2 \pmod 4$ there is also an odd-man-out singleton $\{n - 1\}$; it will never be used, is merely a pesky detail, and will only be mentioned once again.) These cells are fixed for all rounds. Hyperedges added in a round are used only within that round; except for those used in actions, they are ignored in future rounds and can be thought of as deleted.

Algorithm for a round. We consider a round starting when $\mathcal{P}$ has ℓ components. For each path P_i in $\mathcal{P}$, designate one endpoint as the start s_i and the other as the terminal t_i. (If P_i is an isolated

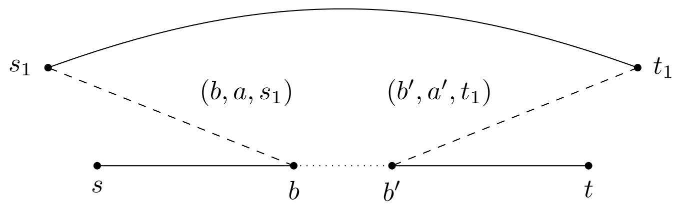

FIGURE 2. Patching together base-pair paths. Suppose a previous offer included an apex a and base point b, to which we added some path end s_1 to propose hyperedge (b, a, s_1). If by coincidence the current offer includes base point $b+2$ and an apex a', then, for some the opposite path end t_1, propose hyperedge $(b+2, a', t_1)$. Accepting this pair of proposals means adding the (dashed) path edges $\{b, s_1\}$ and $\{b+2, t_1\}$ and deleting the (dotted) edge $\{b, b+2\}$, turning two $\mathcal{P}$ paths into one. Correspondingly, we accept the two hyperedges into $\mathcal{C}$, and delete the base pair $\{b, b+2\}$ from B.

point, $s_i = t_i$.) The round will consider up to $10f$ semirandom offers, with

$$f = f(\ell) = n^{3/4}\ell^{1/4}. \tag{5.4}$$

The round terminates with success as soon as $\mathcal{P}$ has $\lfloor (9/10)\ell \rfloor$ components, or with failure after $10f$ semirandom offers or after f^2/n actions have been attempted (see below), whichever comes first.

Consider only offers of the form (b, a), b odd (and not the pesky odd-one-out singleton, if any) and a even, discarding others. Also discard any offer where b was offered earlier in this round. Assume then that this is the first time b is offered. If the cell partner b' of b also has not been offered, complete this offer to (b, a, t) where t is a random terminal, and add (b, a, t) to $\mathcal{H}$. Consider the cell "seeded". If the cell partner b' was previously seen (the cell was seeded), then some hyperedge (b', a', t_i) is in $\mathcal{H}$, in which case complete the current offer to (b, a, s_i). Consider the cell of b "filled", and "attempt action". Attempting action means taking action (see below) if the following conditions are satisfied:

- (H2b-C1) The path P_i with endpoints s_i and t_i has not already been merged into another path: s_i and t_i remain path endpoints. (P_i may have subsumed another path; that is fine.)
- (H2b-C2) $(b, b') \notin P_i$. (This condition does not apply in the final round; see below.)
- (H2b-C3) (b, b') is a good base pair in B (and therefore an edge in $\mathcal{P}$).
- (H2b-C4) a and a' are distinct apexes in A.
- (H2b-C5) None of (b, b'), a, nor a' is in Q (blocked).
- (H2b-C6) $N^-(a) \cap N^-(a') = \emptyset$.

If these conditions are satisfied, the action is taken as follows. The action decreases the number of components of $\mathcal{P}$ by 1; see Figure 2.

- (H2b-A1) Add the hyperedges (b, a, s_i) and (b', a', t_i) to $\mathcal{C}$, and add the edges (b, s_i) and (b', t_i) to $\mathcal{P}$.
- (H2b-A2) Delete (b, b') from $\mathcal{P}$, B, and D, and delete a and a' from A and D.
- (H2b-A3) Add $N^+(N^-(\{(b, b'), a, a'\}))$ to Q, blocking the outneighbours of the inneighbours of the deleted nodes.

The last round is a special case. Any round starting with $\ell \geq 10$ paths ends with $\left\lfloor \frac{9}{10}\ell \right\rfloor \geq 9$. The rounds starting with $2 \leq \ell \leq 9$ paths each decrease the number of paths by 1, so the last round starts with $\ell = 1$. Here, if b and b' share a cell in the single path $s, \ldots, b, b', \ldots, t$, then if offered (b, a) we make hyperedge (b, a, t), and if offered (b', a') we make hyperedge (b', a', s). In this case, an action makes $\mathcal{P}$ a Hamilton cycle. At that point, we proceed to Phase 3.

Analysis. The conditions ensure that no node in D has its outdegree reduced by more than 1.

Also, Phase 2b completes within $O(n)$ offers. Round 0 starts with $\ell_0 = O(n)$ so round i starts with $\ell_i \le \ell_0(9/10)^i$, and over all rounds the total number of offers is at most

$$\sum 10f(\ell_i) = \sum 10n^{3/4}((9/10)^i\ell_0)^{1/4} = O(n)\sum((9/10)^{1/4})^i = O(n).$$

It remains only to show that a.a.s., in every round, at least $(1/10)\ell$ actions are taken.

Note that for every i, with $f_i = f(\ell_i)$,

$$f_i \ge f(1) = \Omega(n^{3/4}) \qquad \text{and} \qquad f_i \le f(\ell_0) = n\epsilon_D^{1/4} \le n/100, \tag{5.5}$$

the latter depending on having ϵ_D sufficiently small.

We now focus on a single round and thus refer simply to ℓ and $f = f(\ell)$. We will use the phrase "overwhelmingly unlikely" for probabilities of order $\exp(-\Omega(n^p))$ for some constant $p > 0$, and "with overwhelming probability" for the complement.

Claim 11. *With probability* $1 - \exp(-\Omega(n^{1/4}))$, $10f$ *offers fill at least* f^2/n *cells, enabling* f^2/n *attempted actions.*

Proof. Although the offers arrive successively (and the round is terminated once $(1/10)\ell$ actions are taken), consider the round's full $10f$ offers. Each offer has probability asymptotically $1/2$ of being of the requisite form (b, a). Each one of that form has equal probability of b being a "left" or "right" element of a cell, i.e., of $b \equiv 1 \pmod 4$ or $b \equiv 3 \pmod 4$. In $10f$ offers, with probability $1 - \exp(-\Omega(n))$, there will be at least $2.4f$ each lefts and rights; we assume henceforth that this is so.

We have $\ell \le \ell_0 \le O(\epsilon_D)n$ and thus (for ϵ_D sufficiently small) $2.4f(\ell) \le n/10$ and thus each new left falls into a new cell with probability at least $9/10$, so the expected number of left-occupied cells is at least $2.2f$. Each of these left "balls" fell in a random cell, and changing the location of one ball changes the number of left-occupied cells by at most 1, so by the Azuma-Hoeffding inequality the probability of a deviation of the order of the mean is overwhelmingly unlikely: of order

$$\exp(-\Omega(f^2/(f \cdot 1))) = \exp(-\Omega(f)).$$

We presume henceforth that there are at least $2.1f$ left-occupied cells.

Likewise, the right balls are overwhelmingly likely to fall into $2.1f$ distinct cells, and we presume henceforth that they do. The locations of these right-occupied cells are uniformly random. Each has probability at least $(2.1f)/(n/4)$ of also being left-occupied, so the expected number of filled cells is at least $17.6f^2/n$. Changing the location of one right-occupied cell changes the number of full cells by at most 1, so by the Azuma-Hoeffding inequality the probability of a deviation on the order of the mean is overwhelmingly unlikely, of order

$$\exp(-\Omega((f^2/n)^2/(f \cdot 1))) = \exp(-\Omega(f^3/n^2)) = \exp(-\Omega(n^{1/4})).$$

This finishes the proof of the claim. □

Claim 12. *When a new cell is filled, the conditions for action are satisfied with probability at least* $2/10$*, assuming that the number of paths in* $\mathcal{P}$ *is at most* $(9/10)\ell$.

Proof. The location of the new cell is uniformly random over all cells, conditioned on not being one of the cells previously filled. For a uniformly random cell X, let F be the event that X was not previously filled, and let E be the event that X satisfies the conditions for action. We are interested in $\Pr(E \mid F)$; consider first just $\Pr(E)$.

For a random cell X, condition (H2b-C1) fails w.p. at most 1/10, or the round would already have ended successfully. Condition (H2b-C2) fails only if the randomly selected terminal is the terminal of the path containing b, and the probability of this is $1/\ell \leq 1/2$. Conditions (H2b-C3) and (H2b-C4) fail with probability $O(\epsilon_D)$, because A and B contain all but an $O(\epsilon_D)$ fraction of A_0 and B_0 respectively. Condition (H2b-C5) fails w.p. $O(\epsilon_Q)$ because $|Q| = O(\epsilon_Q)n$ throughout. Finally, (H2b-C6) fails w.p. $O(1/n)$ by Claim 8. In all,

$$\Pr(E) \geq 1 - (1/10 + 1/2 + O(\epsilon_D + \epsilon_Q) + O(1/n)) \geq 3/10$$

for appropriate choices of ϵ_D and ϵ_Q.

The algorithm terminates if the number of attempted actions — cells filled — exceeds f^2/n. Thus, the probability that a random cell X was previously filled is $\Pr(\overline{F}) \leq (f^2/n)/(n/4) = 4(f/n)^2 \leq 2/100^2$ by (5.5). So, $\Pr(F) \geq 0.99$. It follows that $\Pr(E \mid F) \geq \Pr(E \cap F) \geq \Pr(E) - \Pr(\bar{F}) \geq 2/10$. □

From Claim 11, with overwhelming probability at least f^2/n actions can be attempted, and if the hypothesis of Claim 12 is satisfied, the number of these succeeding is distributed as $B(f^2/n, 2/10)$, and thus is overwhelmingly likely to exceed $\frac{1}{10}f^2/n$: using (5.5) again, the failure probability is of order $\exp(-\Omega(f^2/n)) = \exp(-\Omega(n^{1/2}))$.

This number of successful actions would far exceed ℓ: their ratio is

$$\frac{\frac{1}{10}f^2/n}{\ell} = \frac{n^{3/2}\ell^{1/2}}{10n\ell} = \frac{1}{10}\,(n/\ell)^{1/2} \geq \frac{1}{10}(1/\epsilon_D)^{1/2} > 1,$$

for ϵ_D sufficiently small.

It is impossible that the number of successful actions exceeds ℓ, so we conclude that, with overwhelming probability, the hypothesis of Claim 12 must at some point fail. That is, at some point $\mathcal{P}$ must fall below $(9/10)\ell$, upon which the round terminates with success.

Since the failures were all overwhelmingly unlikely (of order $\exp(-\Omega(n^p))$ for some constant $p > 0$), by the union bound, failure remains overwhelmingly unlikely even over the $O(\ln n)$ rounds.

Assuming success, Phase 2 terminates with $\mathcal{P}$ consisting of a single cycle. Some of its edges correspond to hyperedges committed to $\mathcal{C}$; its other edges consist of base pairs comprising one part of D, whose other part consists of the apexes not yet committed to $\mathcal{C}$. Thus D is an induced subgraph of D_1, in which every node's degree is at most 1 smaller than it is in D_1.

5.3. **Phase 3.** As previously stated, in this final phase we take a perfect matching between A and B in the graph D with directions removed from its edges. And, as noted earlier, the cardinalities of A and B match. Lemma 10 implies that this matching exists, just as in Subsection 4.3.

For each $\{b, b+2\} \in B$ and its matching partner $a \in A$, we place the hyperedge $\{b, a, b+2\}$ into $\mathcal{C}$. That hyperedge was formed in Phase 1, when the corresponding edge was added to D. That completes the Hamilton cycle.

5.4. **Extension to s-uniform hypergraphs.** The linear-time 2-offer strategy to construct a loose Hamilton cycle in a 3-uniform hypergraph extends easily to s-uniform hypergraphs.

To avoid excessive notation, consider $s = 4$ for example. Take base pairs $\{1, 4\}, \{4, 7\}, \{7, 10\}, \ldots$, and apex nodes (no longer single points) $\{2, 3\}, \{5, 6\}, \{8, 9\}, \ldots$. The generalisation to other s is clear. We will see that it is easy to keep each apex node set intact, so that it behaves just like an apex point in the $s = 3$ case.

In Phase 1, on offer of a base point b and an apex point a contained in an apex set $\mathbf{a} \in A$, choose base point $b + s - 1$ (to give base pair $\mathbf{b} = \{b, b+s-1\}$), choose all the other points in $\mathbf{a}$ (to give apex $\mathbf{a}$), creating an edge between $\mathbf{b}$ and $\mathbf{a}$ in D and correspondingly creating the hyperedge $\{\mathbf{b} \cup \mathbf{a}\}$. This gives an auxiliary graph D as before.

In Phase 2a, on offer of base points $\{b_1, b_2\}$, make hyperedge $\{\{b_1, b_2\} \cup \mathbf{a}\}$ where $\mathbf{a}$ is a failed apex node. As before, this uses up the node $\mathbf{a}$ and turns two paths into three.

In Phase 2b, for an offered base point b and apex point $a \in \mathbf{a} \in A$, choose some path end point s_1 and the rest of $\mathbf{a}$ to make candidate hyperedge $\{\{b, s_1\} \cup \mathbf{a}\}$. When another such candidate hyperedge contains $b + s - 1$, just as for $s = 3$, three paths can be turned into two, using these two hyperedges and deleting the base $\{b, b + s - 1\}$.

With these extensions, the mechanics is exactly as for the $s = 3$ case. The treatment of the auxiliary graph D is identical. This gives the strategy for constructing a loose Hamilton cycle in the 2-offer model for an s-uniform hypergraph.

Of course, whatever can be done in linearly many steps in the 2-offer model can also be done in the 1-offer model. This establishes Theorem 4.

Acknowledgements

Greg is grateful to Alan Frieze for guidance on graph Hamilton cycles, to Olaf Parczyk for extensive discussions on semirandom hypergraphs, and to Warach Veeranonchai for a careful reading resulting in several improvements to the manuscript.

References

[1] Natalie Behague, Trent Marbach, Paweł Prałat, and Andrzej Ruciński. Subgraph games in the semi-random graph process and its generalization to hypergraphs. *arXiv preprint arXiv:2105.07034*, 2022.

[2] Omri Ben-Eliezer, Lior Gishboliner, Dan Hefetz, and Michael Krivelevich. Very fast construction of bounded-degree spanning graphs via the semi-random graph process. *Proceedings of the 31st Symposium on Discrete Algorithms (SODA'20)*, pages 728–737, 2020.

[3] Omri Ben-Eliezer, Dan Hefetz, Gal Kronenberg, Olaf Parczyk, Clara Shikhelman, and Miloš Stojaković. Semi-random graph process. *Random Structures & Algorithms*, 56(3):648–675, 2020.

[4] Sofiya Burova and Lyuben Lichev. The semi-random tree process. *arXiv preprint arXiv:2204.07376*, 2022.

[5] Pat Devlin and Jeff Kahn. Perfect fractional matchings in k-out hypergraphs. *The Electronic Journal of Combinatorics*, 24(3):#P3.60, 2017.

[6] Alan Frieze, Pu Gao, Calum MacRury, Paweł Prałat, and Gregory B Sorkin. Building Hamiltonian cycles in the semi-random graph process in less than $2n$ rounds. *arXiv preprint*, 2023.

[7] Alan Frieze and Gregory B Sorkin. Efficient algorithms for three-dimensional axial and planar random assignment problems. *Random Structures & Algorithms*, 46(1):160–196, 2015.

[8] Alan Frieze and Gregory B Sorkin. Hamilton cycles in a semi-random graph model. *arXiv preprint arXiv:2208.00255*, 2022.

[9] Alan M Frieze. Maximum matchings in a class of random graphs. *Journal of Combinatorial Theory, Series B*, 40(2):196–212, 1986.

[10] David Gamarnik, Mihyun Kang, and Paweł Prałat. Cliques, chromatic number, and independent sets in the semi-random process. *arXiv preprint arXiv:2303.13443*, 2023.

[11] Pu Gao, Bogumił Kamiński, Calum MacRury, and Paweł Prałat. Hamilton cycles in the semi-random graph process. *European Journal of Combinatorics*, 99:103423, 2022.

[12] Pu Gao, Calum MacRury, and Paweł Prałat. A fully adaptive strategy for hamiltonian cycles in the semi-random graph process. In Amit Chakrabarti and Chaitanya Swamy, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2022, September 19-21, 2022, University of Illinois, Urbana-Champaign, USA (Virtual Conference)*, volume 245 of *LIPIcs*, pages 29:1–29:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[13] Pu Gao, Calum MacRury, and Paweł Prałat. Perfect matchings in the semirandom graph process. *SIAM Journal on Discrete Mathematics*, 36(2):1274–1290, 2022.

[14] Shoni Gilboa and Dan Hefetz. Semi-random process without replacement. In *Extended Abstracts EuroComb 2021*, pages 129–135. Springer, 2021.

[15] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random graphs*, volume 45. John Wiley & Sons, 2011.

[16] Anders Johansson, Jeff Kahn, and Van Vu. Factors in random graphs. *Random Structures & Algorithms*, 33(1):1–28, 2008.

[17] Michal Karoński and Alan Frieze. *Introduction to Random Graphs*. Cambridge University Press, 2016.

[18] Michal Karoński, Ed Overman, and Boris Pittel. On a perfect matching in a random digraph with average out-degree below two. *Journal of Combinatorial Theory, Series B*, 143, 03 2020.
[19] Calum MacRury and Erlang Surya. Sharp thresholds in adaptive random graph processes. *arXiv preprint arXiv:2207.14469*, 2022.
[20] Colin McDiarmid. On the method of bounded differences. In J.Editor Siemons, editor, *Surveys in Combinatorics, 1989: Invited Papers at the Twelfth British Combinatorial Conference*, London Mathematical Society Lecture Note Series, pages 148–188. Cambridge University Press, 1989.
[21] Paweł Prałat and Harjas Singh. Power of k choices in the semi-random graph process. *The Electronic Journal of Combinatorics*, 31(1):#P1.11, 2024.
[22] David W Walkup. Matchings in random regular bipartite digraphs. *Discrete mathematics*, 31(1):59–64, 1980.

Department of Computer Science, University of Toronto, Toronto, ON, Canada
Email address: molloy@cs.toronto.edu

Department of Mathematics, Toronto Metropolitan University, Toronto, Canada
Email address: pralat@torontomu.ca

Department of Mathematics, The London School of Economics and Political Science, London, England
Email address: g.b.sorkin@lse.ac.uk