



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/179182/>

Version: Accepted Version

Article:

Dyer, M, Greenhill, C, Kleer, P et al. (2021) Sampling hypergraphs with given degrees. Discrete Mathematics, 344 (11). 112566. ISSN: 0012-365X

<https://doi.org/10.1016/j.disc.2021.112566>

© 2021, Elsevier. This manuscript version is made available under the CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Sampling hypergraphs with given degrees

Martin Dyer* Catherine Greenhill† Pieter Kloor‡ James Ross†
Leen Stougie§

22 June 2020, revised 11 July 2021

Abstract

There is a well-known connection between hypergraphs and bipartite graphs, obtained by treating the incidence matrix of the hypergraph as the biadjacency matrix of a bipartite graph. We use this connection to describe and analyse a rejection sampling algorithm for sampling simple uniform hypergraphs with a given degree sequence. Our algorithm uses, as a black box, an algorithm \mathcal{A} for sampling bipartite graphs with given degrees, uniformly or nearly uniformly, in (expected) polynomial time. The expected runtime of the hypergraph sampling algorithm depends on the (expected) runtime of the bipartite graph sampling algorithm \mathcal{A} , and the probability that a uniformly random bipartite graph with given degrees corresponds to a simple hypergraph. We give some conditions on the hypergraph degree sequence which guarantee that this probability is bounded below by a positive constant.

Keywords: hypergraph, degree sequence, sampling, algorithm, Markov chain

1 Introduction

Hypergraphs are combinatorial objects which can be used to abstractly represent general dependence structures, with applications in many areas including machine learning [29] and bioinformatics [26]. We consider the problem of efficiently sampling simple, k -uniform hypergraphs with a given degree sequence, either uniformly or approximately uniformly.

More precisely, a hypergraph $H = (V, E)$ consists of a finite set $V = V(H) = \{v_1, \dots, v_n\}$ of n nodes, and a multiset $E = E(H) = \{e_1, e_2, \dots, e_m\}$ of m edges, where each edge is a nonempty multisubset of V . We say that H is *simple* if there are no repeated edges in E and no edge of E contains a repeated node (so E is a set of subsets of nodes). For any node $v_i \in V$, we define the degree of v_i by

$$d_i = \deg_H(v_i) = |\{e \in E(H) : v_i \in e\}|,$$

*School of Computing, University of Leeds, Leeds LS2 9JT, UK. m.e.dyer@leeds.ac.uk

†School of Mathematics and Statistics, UNSW Sydney, NSW 2052, Australia.
c.greenhill@unsw.edu.au, james.ross@unsw.edu.au

‡Tilburg University, Tilburg, The Netherlands. p.s.kloor@tilburguniversity.edu

§CWI, INRIA-Erable, Vrije Universiteit, Amsterdam, The Netherlands. Leen.Stougie@cwi.nl

and write $\mathbf{d} = (d_1, \dots, d_n)$ for the degree sequence of H . For a positive integer k , we say that H is k -uniform if every edge contains exactly k nodes, counting multiplicities when H is not simple. We then write $\mathcal{H}_k(\mathbf{d})$ for the set of k -uniform simple hypergraphs with degree sequence \mathbf{d} . If there is a positive integer d such that $d_i = d$ for all $i \in [n]$, then we write $\mathcal{H}_k(n, d)$ for the set of all k -uniform d -regular simple hypergraphs on n nodes, instead of $\mathcal{H}_k(\mathbf{d})$.

Recently, Deza, Levin, Meesum & Onn [13] proved that the construction problem for simple 3-uniform hypergraphs is NP-hard. That is, given \mathbf{d} it is NP-hard to decide whether there exists a 3-uniform simple hypergraph with degree sequence \mathbf{d} . This implies that it is not possible to approximate $|\mathcal{H}_k(\mathbf{d})|$ efficiently in general, since approximate counting can distinguish 0 from a positive number. Moreover, hardness of construction also directly implies that (approximate) uniform sampling is a difficult problem in general.

Arafat et al. [2] recently gave an algorithm to construct a non-simple hypergraph with given degrees and edge sizes. Chodrow [9] considered Markov Chain Monte Carlo approaches for generating such hypergraphs. We emphasize that throughout this work, we only consider simple hypergraphs. To the best of our knowledge, the only rigorously-analysed algorithm for this problem in the literature is the configuration model, see Section 2.2.

Our approach is based on the well-known connection between hypergraphs and bipartite graphs. We first fix some notation for bipartite graphs and then explain this relation. A bipartite graph $B = (X \cup Y, A)$ consists of a bipartition $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_m\}$ of nodes, and an edge set $A \subseteq X \times Y = \{\{x, y\} : x \in X, y \in Y\}$. For a pair of nonnegative integer sequences $\mathbf{d} = (d_1, \dots, d_n)$ and $\mathbf{k} = (k_1, \dots, k_m)$, let $\mathcal{B}(\mathbf{d}, \mathbf{k})$ be the set of all simple bipartite graphs B such that

$$\deg_B(x_i) = d_i \text{ for all } x_i \in X, \quad \text{and} \quad \deg_B(y_j) = k_j \text{ for all } y_j \in Y.$$

We say that (\mathbf{d}, \mathbf{k}) is the (bipartite) degree sequence of B . Note that if $\sum_{i=1}^n d_i \neq \sum_{j=1}^m k_j$, then $\mathcal{B}(\mathbf{d}, \mathbf{k}) = \emptyset$. If there is a fixed integer k such that $k_j = k$ for all $j \in [m]$, then we write $\mathcal{B}(\mathbf{d}, k)$ instead of $\mathcal{B}(\mathbf{d}, \mathbf{k})$, and we call such bipartite graphs *half-regular*. If in addition $\mathbf{d} = (d, d, \dots, d)$ is regular then we write $\mathcal{B}(n, d, k)$. For any node $v \in X \cup Y$, let $\mathcal{N}_B(v) = \{w \in X \cup Y : \{v, w\} \in A\}$ be the neighbourhood of node v in B .

Every hypergraph $H = (V, E)$ can be represented as a bipartite graph B_H , as follows. Fix a labelling of the edges of H , say $E = \{e_1, \dots, e_m\}$, then let

$$X = V, \quad Y = E \quad \text{and} \quad A = \left\{ \{v_i, e_j\} \in X \times Y \mid v_i \in e_j \right\}.$$

If $H \in \mathcal{H}_k(\mathbf{d})$ then $B(H) \in \mathcal{B}(\mathbf{d}, k)$. Conversely, every bipartite graph $B = (X \cup Y, A)$ corresponds to a hypergraph $H_B = (V, E)$, where $V = X$ and

$$E = \{\mathcal{N}_B(y) \mid y \in Y\}.$$

Furthermore, H_B is simple if and only if every node in Y has a distinct set of neighbours in B ; that is, if $\mathcal{N}_B(y_i) = \mathcal{N}_B(y_j)$ implies $i = j$. If H_B is simple then we say that the bipartite graph B is *H-simple*.

Write $\mathcal{B}^*(\mathbf{d}, k)$ for the set of all H-simple half-regular bipartite graphs, and define $\varphi : \mathcal{B}^*(\mathbf{d}, k) \rightarrow \mathcal{H}_k(\mathbf{d})$ as the canonical mapping that maps B to H_B , as described above. We can

use rejection sampling to turn any sampling algorithm for $\mathcal{B}(\mathbf{d}, k)$ into a sampling algorithm for $\mathcal{H}_k(\mathbf{d})$, as follows:

```

HYPERGRAPH SAMPLING( $\mathbf{d}, k, \mathcal{A}$ )
INPUT: Parameters  $(\mathbf{d}, k)$ ;
        algorithm  $\mathcal{A}$  for sampling from  $\mathcal{B}(\mathbf{d}, k)$ .
Begin
  repeat
    sample  $B$  from  $\mathcal{B}(\mathbf{d}, k)$  using  $\mathcal{A}$ 
  until  $B$  is H-simple
output  $\varphi(B)$ 
end.

```

Note that for all $H \in \mathcal{H}_k(\mathbf{d})$ we have $|\varphi^{-1}(H)| = m!$, as there are $m!$ distinct ways to label the edges of H when H is simple. Hence, if \mathcal{A} samples *uniformly* from $\mathcal{B}(\mathbf{d}, k)$ then the output of HYPERGRAPH SAMPLING is uniform over $\mathcal{H}_k(\mathbf{d})$.

The expected number of times HYPERGRAPH SAMPLING draws a sample from $\mathcal{B}(\mathbf{d}, k)$, using algorithm \mathcal{A} , depends on the proportion of bipartite graphs in $\mathcal{B}(\mathbf{d}, k)$ which are H-simple: that is, the ratio $|\mathcal{B}^*(\mathbf{d}, k)|/|\mathcal{B}(\mathbf{d}, k)|$. The goal of our work is to identify pairs (\mathbf{d}, k) for which $|\mathcal{B}(\mathbf{d}, k)|/|\mathcal{B}^*(\mathbf{d}, k)|$ is bounded above by a polynomial in n . For such pairs, if the output of \mathcal{A} is close to uniform then this implies that the expected number of times we run \mathcal{A} before an H-simple element of $\mathcal{B}(\mathbf{d}, k)$ is found is at most polynomial. We also identify pairs (\mathbf{d}, k) for which the ratio is bounded above by a constant, in which case only a constant number of calls to \mathcal{A} are expected. This is made more specific in the next subsection.

1.1 Our contributions

The total variation distance between two probability distributions σ, π on a set Ω is given by

$$d_{TV}(\sigma, \pi) = \frac{1}{2} \sum_{x \in \Omega} |\sigma(x) - \pi(x)| = \max_{S \subseteq \Omega} |\sigma(S) - \pi(S)|. \quad (1.1)$$

Suppose that the distribution of the output of algorithm \mathcal{A} on $\mathcal{B}(\mathbf{d}, k)$ is $\sigma_{\mathcal{B}}$, and let $\sigma_{\mathcal{H}}$ be the output of the algorithm HYPERGRAPH SAMPLING. Then $\sigma_{\mathcal{H}}$ is a distribution on $\mathcal{H}_k(\mathbf{d})$ which is obtained by setting $H = \varphi(B)$ where B has distribution $\sigma_{\mathcal{B}}$, conditioned on $B \in \mathcal{B}^*(\mathbf{d}, k)$. To make it clear which distribution we are using, we write \mathbb{P}_{σ} for the probability mass function of the distribution σ . Let $\pi_{\mathcal{B}}$ be the uniform distribution on $\mathcal{B}(\mathbf{d}, k)$, and let $\pi_{\mathcal{H}}$ be the uniform distribution on $\mathcal{H}_k(\mathbf{d})$.

A *fully-polynomial almost uniform sampler (FPAUS)* for sampling from a set Ω is an algorithm that, with probability at least $\frac{3}{4}$, outputs an element of Ω in time polynomial in $\log |\Omega|$ and $\log(1/\varepsilon)$, such that the output distribution is ε -close to the uniform distribution π on Ω in *total variation distance*: that is, $d_{TV}(\sigma, \pi) \leq \varepsilon$. If $\Omega = \mathcal{B}(\mathbf{d}, k)$ or $\Omega = \mathcal{H}_k(\mathbf{d})$ then $\log |\Omega| = O(M \log M)$ for $M = \sum_{i=1}^n d_i$: this follows from [18, Theorem 1.3], restated below as Theorem 4.1. So an FPAUS for $\mathcal{H}_k(\mathbf{d})$ or $\mathcal{B}(\mathbf{d}, k)$ must have running time bounded above by a polynomial in M and $\log(1/\varepsilon)$.

Denote by $\rho(\mathbf{d}, k)$ the runtime of an algorithm for testing H-simplicity, when run on inputs from $\mathcal{B}(\mathbf{d}, k)$. Given $B \in \mathcal{B}(\mathbf{d}, k)$, we can test whether $\varphi(B)$ is H-simple by creating the edges of $\varphi(B)$ and sorting them lexicographically, requiring $O(M \log M)$ time and $O(M)$ space. (Better implementations may be possible.) Hence we can assume that $\rho(\mathbf{d}, k) = O(M \log M)$, and that the H-simplicity test also creates the hypergraph $\varphi(B)$. Our first result summarises the properties of the algorithm `HYPERGRAPH SAMPLING` $(\mathbf{d}, k, \mathcal{A})$ in terms of the output distribution and runtime of \mathcal{A} , and the runtime $\rho(\mathbf{d}, k)$ of the H-simplicity test. The proof of Theorem 1.1, which is fairly standard, is presented in Section 3.

Theorem 1.1. *Suppose that n is a positive integer, $\mathbf{d} = (d_1, \dots, d_n)$ is a sequence of positive integers, and k is a positive integer such that $k \geq 3$ and $\mathcal{B}^*(\mathbf{d}, k)$ is non-empty.*

(i) *The output distribution $\sigma_{\mathcal{H}}$ of `HYPERGRAPH SAMPLING` satisfies*

$$d_{TV}(\sigma_{\mathcal{H}}, \pi_{\mathcal{H}}) \leq \frac{3}{2} \cdot \frac{d_{TV}(\sigma_{\mathcal{B}}, \pi_{\mathcal{B}})}{\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))}.$$

(ii) *The expected runtime of `HYPERGRAPH SAMPLING` $(\mathbf{d}, k, \mathcal{A})$ is*

$$q(\mathbf{d}, k) \left(\tau(\mathbf{d}, k) + \rho(\mathbf{d}, k) \right)$$

where $q(\mathbf{d}, k)^{-1} = \mathbb{P}_{\sigma_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))$ is the probability that a sampled bipartite graph is H-simple and $\tau(\mathbf{d}, k)$ is the (expected) runtime of algorithm \mathcal{A} on $\mathcal{B}(\mathbf{d}, k)$. Furthermore, the probability that `HYPERGRAPH SAMPLING` needs more than $t q(\mathbf{d}, k)$ iterations of \mathcal{A} before finding an element of $\mathcal{B}^*(\mathbf{d}, k)$ is at most $\exp(-t)$ for any $t > 0$.

(iii) *Suppose that $d_{TV}(\sigma_{\mathcal{B}}, \pi_{\mathcal{B}}) \leq \varepsilon$ and $\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k)) \geq 1 - c_0$ for some $\varepsilon \in (0, 1)$ and $c_0 \in (0, 1 - \varepsilon)$. Then*

$$d_{TV}(\sigma_{\mathcal{H}}, \pi_{\mathcal{H}}) \leq \frac{3\varepsilon}{2(1 - c_0)}$$

and the expected runtime of `HYPERGRAPH SAMPLING` $(\mathbf{d}, k, \mathcal{A})$ is at most

$$(1 - c_0 - \varepsilon)^{-1} \left(\tau(\mathbf{d}, k) + \rho(\mathbf{d}, k) \right).$$

(iv) *If \mathcal{A} is an FPAUS for $\mathcal{B}(\mathbf{d}, k)$ and the assumptions of (iii) hold, then we can transform `HYPERGRAPH SAMPLING` $(\mathbf{d}, k, \mathcal{A})$ into an FPAUS for $\mathcal{H}_k(\mathbf{d})$ by terminating after $\lceil 2(1 - c_0 - \varepsilon)^{-1} \rceil$ calls to \mathcal{A} and reporting FAIL.*

Remark 1.2. For Theorem 1.1(iv) to provide an FPAUS with an *explicit* upper bound on the runtime, explicit bounds are needed on ε , c_0 and $\rho(\mathbf{d}, k)$. However, for Theorem 1.1(iii) it is enough to know that sufficiently small values of c_0 , ε exist.

Remark 1.3. We require $O(M)$ space to store a bipartite graph $B \in \mathcal{B}(\mathbf{d}, k)$, or the corresponding hypergraph $\varphi(B)$. It follows that the space complexity of `HYPERGRAPH SAMPLING` is bounded above by the sum of $O(M)$ and the space complexity of the chosen algorithm \mathcal{A} .

We see from Theorem 1.1 that $d_{TV}(\sigma_{\mathcal{B}}, \pi_{\mathcal{B}})$ and $\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))$ are the two crucial quantities which control both the expected runtime of $\text{HYPERGRAPH_SAMPLING}(\mathbf{d}, k, \mathcal{A})$, and how far the output varies from uniform. The first of these quantities is determined by the choice of algorithm \mathcal{A} . In our next two theorems, we provide a lower bound on $\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))$ when $\mathbf{d} = (d, d, \dots, d)$ is regular, and give an asymptotic lower bound on this probability which holds when \mathbf{d} is irregular but sparse.

Theorem 1.4. *Let $n \in \mathbb{N}$, $d \in \mathbb{N}$ and $k \geq 3$ so that $m = nd/k \in \mathbb{N}$ and $\binom{m}{2} < \binom{n}{k}$. Then*

$$\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(n, d, k)) \geq 1 - \binom{m}{2} \binom{n}{k}^{-1}.$$

Hence Theorem 1.1(iii) applies when $\binom{m}{2} \leq c_0 \binom{n}{k}$ for some $c_0 \in (0, 1 - \varepsilon)$, where $d_{TV}(\sigma_{\mathcal{B}}, \pi_{\mathcal{B}}) \leq \varepsilon$.

Remark 1.5. When $k \geq 3$ is a fixed constant, the lower bound in Theorem 1.4 is $1 - o(1)$ if $d = o(n^{k/2-1})$.

Theorem 1.6. *Assume that for each positive integer n we have an integer $k = k(n) \geq 3$ and a sequence $\mathbf{d} = \mathbf{d}(n) = (d_1, \dots, d_n)$ of positive integers such that $M = \sum_{i=1}^n d_i$ tends to infinity with n . Suppose that k divides M for infinitely many values of n . Assume that $k^2 d_{\max}^2 = o(M)$ and let $\pi_{\mathcal{B}}$ be the uniform distribution on $\mathcal{B}(\mathbf{d}, k)$. Then*

$$\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k)) \geq 1 - \frac{n^k d_{\max}^k}{M^k} \cdot \binom{M/k}{2} \binom{n}{k}^{-1} \cdot (1 + o(1)).$$

Writing $m = M/k$ and $d = M/n$, we see that Theorem 1.1(iii) applies when

$$\left(\frac{d_{\max}}{d}\right)^k \binom{m}{2} \leq c_0 \binom{n}{k} \tag{1.2}$$

for some $c_0 \in (0, 1 - \varepsilon)$, where $d_{TV}(\sigma_{\mathcal{B}}, \pi_{\mathcal{B}}) \leq \varepsilon$.

In the hypergraph setting, m is the number of edges and d is the average degree of any hypergraph in $\mathcal{H}_k(\mathbf{d})$.

Remark 1.7. When $k \geq 3$ is a fixed constant, the lower bound in Theorem 1.6 is $1 - o(1)$ if $d_{\max} = o(M^{1-2/k})$. Similarly, if $k \geq 3$ is a fixed constant and $d_{\max} = O(d)$ then the lower bound in Theorem 1.6 is $1 - o(1)$ whenever $d^2 = o(n^{k-2})$, as in the regular case.

There are several algorithms \mathcal{A} in the literature for efficiently sampling from $\mathcal{B}(\mathbf{d}, k)$, either uniformly or almost uniformly, under various conditions on \mathbf{d} and k . These will be reviewed in Section 2.1, together with the properties of the resulting algorithm $\text{HYPERGRAPH_SAMPLING}(\mathbf{d}, k, \mathcal{A})$. In Section 2.2 we discuss the configuration model for hypergraphs, which can be used as an expected polynomial-time sampling algorithm when $k d_{\max} = O(\log n)$.

Then in Section 3 we provide a general framework which we use to analyse the algorithm $\text{HYPERGRAPH_SAMPLING}$. In Section 4 we fill in the details for the regular regime (Theorem 1.4) and the irregular, sparse regime (Theorem 1.6).

2 Related work

2.1 Various bipartite sampling algorithms, and implications

In this section we describe several algorithms for efficient sampling from $\mathcal{B}(\mathbf{d}, k)$, uniformly or almost uniformly, under various conditions on \mathbf{d}, k . We also apply Theorem 1.1 to describe when $\text{HYPERGRAPH SAMPLING}(\mathbf{d}, k, \mathcal{A})$ is an efficient algorithm for sampling from $\mathcal{H}_k(\mathbf{d})$ (uniformly or near-uniformly), for each choice of \mathcal{A} . We focus on the time complexity of these algorithms, as most authors have not stated the space complexity of their algorithms.

The first two algorithms mentioned below perform *exactly uniform* sampling from $\mathcal{B}(\mathbf{d}, k)$.

- (I) If $kd_{\max} \leq C \log n$ then the bipartite configuration model gives rise to an algorithm for sampling (exactly) uniformly from $\mathcal{B}(\mathbf{d}, k)$. But the bipartite configuration model is equivalent to the configuration model for hypergraphs, described in Section 2.2, and so there is no advantage to working in the bipartite graph setting when $kd_{\max} \leq C \log n$. (See Lemma 2.3).
- (II) Next suppose that $(d_{\max} + k)^4 = O(M)$. Building on the work of [17, 24], Arman, Gao and Wormald [3, Theorem 4] describe an algorithm which samples uniformly from $\mathcal{B}(\mathbf{d}, k)$ with expected runtime $O(M)$ and space complexity $O(nM)$. Note that $d_{TV}(\sigma_{\mathcal{B}}, \pi_{\mathcal{B}}) = 0$ as the output is exactly uniform. Using this algorithm as \mathcal{A} and applying Theorem 1.1, we see that $\text{HYPERGRAPH SAMPLING}(\mathbf{d}, k, \mathcal{A})$ performs exact sampling from $\mathcal{H}_k(\mathbf{d})$ with expected runtime $O(M + \rho(\mathbf{d}, k))$ whenever (1.2) holds for some constant $c_0 \in (0, 1)$. In particular, this holds whenever (\mathbf{d}, k) are as described in Remark 1.7.

The next algorithm produces output which is *asymptotically uniform*, meaning that the output distribution is only $o(1)$ from uniform in total variation distance.

- (III) If $d_{\max} + k = O(M^{1/4-\tau})$ for some positive constant τ then the algorithm from [3] can be applied, as described in (II). An alternative is to use the sampling algorithm of Bayati, Kim and Saberi [4, Theorem 1], which has expected runtime $O(d_{\max}M)$ (see the proof of [4, Theorem 1]). The output of this algorithm satisfies $d_{TV}(\sigma_{\mathcal{B}}, \pi_{\mathcal{B}}) = o(1)$, where this vanishing term depends only on n and cannot be made smaller by increasing the runtime of the algorithm. Hence we can take $\varepsilon = o(1)$ in Theorem 1.1, and conclude that for this choice of \mathcal{A} , the algorithm $\text{HYPERGRAPH SAMPLING}(\mathbf{d}, r, \mathcal{A})$ has expected runtime $O(d_{\max}M + \rho(\mathbf{d}, k))$ whenever (1.2) holds for some constant $c_0 \in (0, 1)$, and the distribution of the output is within $o(1)$ of uniform: that is, $d_{TV}(\sigma_{\mathcal{H}}, \pi_{\mathcal{H}}) = o(1)$.

Remark 2.1. Although the Arman, Gao and Wormald algorithm applies for a slightly wider range of values of (\mathbf{d}, k) , has a better expected runtime bound and performs exactly uniform sampling, the Bayati, Kim and Saberi algorithm has one advantage: it is much easier to implement. Indeed, Bayati, Kim and Saberi [4, Theorem 3] used sequential importance sampling to give an algorithm which is close to an FPAUS, except that the runtime is polynomial in $1/\varepsilon$, while in an FPAUS the dependence should be on $\log(1/\varepsilon)$. However, this algorithm

is valid only when $d_{\max} = O(M^{1/4-\tau})$ for some $\tau > 0$ and no longer has the advantage of simplicity, and so it is surpassed by the fast, precisely uniform sampling algorithm of Arman, Gao and Wormald [3], described in (II) above. (Other authors, such as Chen et al. [8], have used sequential importance sampling to sample bipartite graphs with given degrees, but without rigorous analysis.)

Now we survey some algorithms which are FPAUSs for $\mathcal{B}(\mathbf{d}, k)$. Each can be used as \mathcal{A} to give an FPAUS for $\mathcal{H}_k(\mathbf{d})$, using Theorem 1.1(iv), so long as (1.2) holds (or $\binom{m}{2} \leq c_0 \binom{n}{k}$ when \mathbf{d} is regular) for some $c_0 \in (0, 1 - \varepsilon)$. The runtime of each of these algorithms is the mixing time of their underlying Markov chain multiplied by the cost of performing a single step of the chain, the latter being an implementation-dependent cost. In all cases, the polynomial bound on the mixing time is quite a high-degree polynomial and is not believed to be tight. We do not always state the mixing time.

- (IV) Jerrum, Sinclair and Vigoda [21] described and analysed a simulated annealing algorithm which gives an FPAUS for sampling perfect matchings from a given bipartite graph. As a corollary, they obtained an FPAUS for sampling bipartite graphs for given degrees [21, Corollary 8.1]. Bezáková, Bhatnagar and Vigoda [5] adapted the algorithm from [21] to provide a simplified FPAUS for $\mathcal{B}(\mathbf{d}, k)$, valid for any (\mathbf{d}, k) , with running time

$$O((nm)^2 M^3 \Delta \log^4(nm/\varepsilon)),$$

where n and m are the number of nodes in each part of the bipartition, and $\Delta = \max\{d_{\max}, k\}$.

- (V) Another well-studied Markov chain for sampling (bipartite) graphs with given degrees is the switch Markov chain. It is the simplest Markov chain which walks on the set of all (bipartite) graphs with a given degree sequence, as it deletes and replaces only two edges at a time. The chain has been used in many contexts, including contingency tables [14], and was first applied to bipartite graphs by Kannan, Tetali and Vempala [22]. The mixing time of the switch chain has been shown to be polynomial for various bipartite and general degree sequences, see for example [1, 10, 16, 19, 25, 27]. If a Markov chain leads to an FPAUS then we say that the Markov chain is rapidly mixing. In particular

- Tikhomirov and Youssef [27] considered the switch Markov chain on regular bipartite graphs (in which $d = k$), and proved a sharp Poincaré inequality and an optimal upper bound on the log-Sobolev constant for the switch chain. With these results, they demonstrated that for some constant $c > 0$, when $3 \leq d \leq n^c$ the mixing time of the switch Markov chain on d -regular bipartite graphs is at most

$$O(dn(dn \log dn + \log(2\varepsilon^{-1}))).$$

This is a significant improvement on the previously known bounds in other settings.

- Miklos, Erdős and Soukup [25] showed that the switch Markov chain is rapidly mixing for half-regular bipartite degree sequences (in which one part of the bipartition has regular degrees). An explicit polynomial bound is not clearly stated.

- Cooper, Dyer and Greenhill [10] considered regular graphs and showed that the switch chain is rapidly mixing on the set of all d -regular graphs, for any $d = d(n)$. Their analysis was extended by Greenhill and Sfragara [19] who adapted the proof to sparse, irregular degree sequences. Neither of these works explicitly treated bipartite graphs, though the arguments in both papers are simpler when restricted to bipartite graphs. In Corollary A.3 we state an upper bound on the mixing time of the switch chain on $\mathcal{B}(\mathbf{d}, k)$ which arises from the arguments of [10, 19] when $3 \leq d_{\max} \leq \frac{1}{3}\sqrt{M}$. Specifically we show that in this range, the switch chain has mixing time

$$\Delta^{10} M^7 \left(\frac{1}{2} M \log(M) + \log(\varepsilon^{-1}) \right)$$

where $\Delta = \max\{d_{\max}, k\}$.

- Jerrum and Sinclair [20] defined a notion of P-stability for degree sequences. Roughly speaking, a degree sequence \mathbf{d} is P-stable if small perturbations to \mathbf{d} only change the number of realisations (graphs with degree sequence \mathbf{d}) by a small amount. The notion of P-stable can be adapted to bipartite graphs. Amanatidis and Kler [1] defined a possibly stronger notion, *strong stability*, and showed that the switch chain is rapidly mixing for any strongly stable degree sequence, and for any strongly stable bipartite degree sequence. Erdős et al. [16] proved that the switch chain is rapidly mixing for any P-stable class of bipartite degree sequences.

(VI) The Curveball chain [28] is another Markov chain for sampling bipartite graphs with given degrees, in which multiple switches are performed simultaneously. Carstens and Kler [7] showed that the Curveball chain is rapidly mixing whenever the switch chain is rapidly mixing.

Remark 2.2. We have focussed on uniform hypergraphs, but our approach can be adapted to non-uniform hypergraphs. Given a vector $\mathbf{k} = (k_1, \dots, k_m)$ which stores the desired edge sizes, let m_ℓ be the number of edges of size ℓ , that is,

$$m_\ell = |\{j \in [m] : k_j = \ell\}|$$

for $\ell \geq 2$. Each hypergraph H on $[n]$ with edge sizes given by \mathbf{k} and with degree sequence \mathbf{d} corresponds to exactly $\prod_{\ell=2}^n m_\ell!$ bipartite graphs from $\mathcal{B}(\mathbf{d}, \mathbf{k})$, as now we must restrict to edge labellings e_1, \dots, e_m so that $|e_j| = k_j$ for $j = 1, \dots, m$. All of the bipartite graph sampling algorithms mentioned in this section generalise to $\mathcal{B}(\mathbf{d}, \mathbf{k})$, with the exception of the result by Miklos, Erdős and Soukup regarding the switch chain for half-regular bipartite graphs [25].

2.2 Sampling hypergraphs using the configuration model

To the best of our knowledge, the only rigorously-analysed algorithm for sampling hypergraphs with given degrees is the *configuration model*. The analogue of the configuration

model for hypergraphs has been used by various authors, for example, in the study of random hypergraphs [12] and for asymptotic enumeration [15]. In the configuration model corresponding to $\mathcal{H}_k(\mathbf{d})$ there are n objects, called cells, and the i th cell contains d_i (labelled) points. A *configuration* is a partition of the $M = \sum_{i=1}^n d_i$ points into M/k parts, each containing k points. A random configuration can be chosen in $O(M)$ time. Shrinking each cell to a node gives an k -uniform hypergraph which may contain loops (that is, an edge containing a node more than once) or repeated edges. If the resulting hypergraph is not simple then the configuration is rejected and a new random configuration is sampled. We say that a configuration is simple if the corresponding hypergraph is simple.

Hence, the configuration model can be used for efficient sampling when the probability that a randomly chosen configuration is simple is bounded below by the inverse of some polynomial in n . This implies that the expected number of trials before a simple configuration is found is at most polynomial.

It follows from asymptotic results of Dudek, Frieze, Ruciński and Šileikis [15] that when $\mathbf{d} = (d, d, \dots, d)$ is regular and $k \geq 3$ is constant, the expected number of trials before a simple configuration is sampled is

$$\exp\left(\frac{1}{2}(k-1)(d-1) + o(1)\right)$$

assuming that $k = 3$ and $d = d(n) = o(n^{1/2})$, or $k \geq 4$ and $d = d(n) = o(n)$. (Asymptotics are as $n \rightarrow \infty$, restricted to values of n such that dn is divisible by k .) For irregular degrees, let $M_2 = M_2(\mathbf{d}) = \sum_{j=1}^n d_j(d_j - 1)$. It follows from [6, Corollary 2.3] that the expected number of trials before a simple configuration is sampled is

$$\exp\left(\frac{(k-1)M_2}{2M} + o(1)\right) \leq \exp\left(\frac{1}{2}(k-1)(d_{\max} - 1) + o(1)\right)$$

whenever $k^4 d_{\max}^3 = o(M)$. Here $k = k(n)$ and $\mathbf{d} = \mathbf{d}(n)$ are such that k divides M for infinitely many values of n . We collect these facts together into the following lemma.

Lemma 2.3. *The configuration model gives an efficient algorithm for sampling uniformly from $\mathcal{H}_k(\mathbf{d})$ whenever $kd_{\max} = O(\log n)$. If $kd_{\max} \leq C \log n$ for some constant $C > 0$ then the expected runtime of this algorithm for $\mathcal{H}_k(\mathbf{d})$ is $O(n^C M) = O(d_{\max} n^{C+1})$. If $kd = o(\log n)$ then the expected runtime of this algorithm is $O(M) = O(d_{\max} n)$. Note $d_{TV}(\sigma_{\mathcal{H}}, \pi_{\mathcal{H}}) = 0$ as the output is exactly uniform.*

Gao and Wormald [17] built on earlier work of McKay and Wormald [24] to give a fast algorithm for exactly uniform sampling of d -regular graphs. Using a recent improvement of Arman, Gao and Wormald [3], a uniformly random d -regular graph on n vertices can be generated in expected runtime $O(dn + d^4)$ whenever $d = o(n^{1/2})$, and a random graph with degree sequence \mathbf{d} can be generated in runtime $O(M)$ whenever $d_{\max}^4 = O(M)$. It is likely that this approach could be adapted to the problem of sampling hypergraphs uniformly.

3 Analysis of HypergraphSampling

First we prove Theorem 1.1.

Proof of Theorem 1.1. To prove (i), observe that by definition,

$$\sigma_{\mathcal{H}}(H) = \sum_{B \in \varphi^{-1}(H)} \mathbb{P}_{\sigma_{\mathcal{B}}}(B \mid B \in \mathcal{B}^*(\mathbf{d}, k)) = \frac{1}{\mathbb{P}_{\sigma_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))} \cdot \sum_{B \in \varphi^{-1}(H)} \sigma_{\mathcal{B}}(B).$$

This equality also holds with $\sigma_{\mathcal{H}}, \sigma_{\mathcal{B}}$ replaced by $\pi_{\mathcal{H}}, \pi_{\mathcal{B}}$, respectively. Since the set of all preimages $\{\varphi^{-1}(H) : H \in \mathcal{H}_k(\mathbf{d})\}$ forms a partition of $\mathcal{B}^*(\mathbf{d}, k)$, and using the triangle inequality, we have

$$\begin{aligned} d_{TV}(\sigma_{\mathcal{H}}, \pi_{\mathcal{H}}) &= \frac{1}{2} \sum_{H \in \mathcal{H}_k(\mathbf{d})} |\sigma_{\mathcal{H}}(H) - \pi_{\mathcal{H}}(H)| \\ &\leq \frac{1}{2} \sum_{H \in \mathcal{H}_k(\mathbf{d})} \sum_{B \in \varphi^{-1}(H)} \left| \frac{\sigma_{\mathcal{B}}(B)}{\mathbb{P}_{\sigma_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))} - \frac{\pi_{\mathcal{B}}(B)}{\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))} \right| \\ &\leq \frac{1}{2} \sum_{B \in \mathcal{B}^*(\mathbf{d}, k)} \left| \frac{\sigma_{\mathcal{B}}(B)}{\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))} - \frac{\pi_{\mathcal{B}}(B)}{\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))} \right| \\ &\quad + \frac{1}{2} \sum_{B \in \mathcal{B}^*(\mathbf{d}, k)} \left| \frac{\sigma_{\mathcal{B}}(B)}{\mathbb{P}_{\sigma_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))} - \frac{\sigma_{\mathcal{B}}(B)}{\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))} \right| \\ &\leq \frac{d_{TV}(\sigma_{\mathcal{B}}, \pi_{\mathcal{B}})}{\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))} + \frac{1}{2} \cdot \left| \frac{1}{\mathbb{P}_{\sigma_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))} - \frac{1}{\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))} \right| \sum_{B \in \mathcal{B}^*(\mathbf{d}, k)} \sigma_{\mathcal{B}}(B) \\ &\leq \frac{3}{2} \cdot \frac{d_{TV}(\sigma_{\mathcal{B}}, \pi_{\mathcal{B}})}{\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))}. \end{aligned}$$

The final inequality follows from applying (1.1) with $S = \mathcal{B}^*(\mathbf{d}, k)$.

Next, (ii) is immediate as the number of times that HYPERGRAPH SAMPLING calls \mathcal{A} has a geometric distribution with mean $q(\mathbf{d}, k)$, and we test whether B is H-simple once for every call of \mathcal{A} . We assume here that the test for H-simplicity also creates the hypergraph $\varphi(B)$, so there is no additional cost for the final (output) step of the algorithm. Then, (1.1) implies that

$$q(\mathbf{d}, k)^{-1} = \mathbb{P}_{\sigma_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k)) \geq \mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k)) - \varepsilon,$$

and (iii) follows.

Finally, suppose that \mathcal{A} is an FPAUS for $\mathcal{B}(\mathbf{d}, k)$ and (1.2) holds for some $c_0 \in (0, 1 - \varepsilon)$. It follows from (ii) and (iii) that the probability that HYPERGRAPH SAMPLING($\mathbf{d}, k, \mathcal{A}$) performs more than $\lceil 2(1 - c_0 - \varepsilon)^{-1} \rceil$ iterations of \mathcal{A} is at most $e^{-2} \leq \frac{1}{4}$. Therefore, terminating HYPERGRAPH SAMPLING($\mathbf{d}, k, \mathcal{A}$) after this many calls to \mathcal{A} gives an FPAUS for $\mathcal{H}_k(\mathbf{d})$. To achieve a total variation of ε from HYPERGRAPH SAMPLING($\mathbf{d}, k, \mathcal{A}$), the algorithm \mathcal{A} should be given input $\varepsilon' = 2\varepsilon(1 - c_0)/3$, by (i). \square

A general approach for bounding $\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k))$ is given by the following lemma. The constant c_1 in Lemma 3.1 captures the maximum edge probability relative to the uniform case (in which every neighborhood is equally likely). If c_1 is large then some neighbourhoods are much more likely under $\sigma_{\mathcal{B}}$ than they would be under the uniform distribution. The extent to which the events “ $\mathcal{N}_B(y) = \mathcal{W}$ ” are negatively correlated, as y varies over Y for

fixed $\mathcal{W} \subseteq X$, is described by the constant c_2 . Intuitively, if the degree sequence is close to regular then we expect both c_1 and c_2 to be close to one.

Lemma 3.1. *Suppose that $k \geq 3$ is an integer, \mathbf{d} is a sequence of nonnegative integers such that $\mathcal{B}(\mathbf{d}, k)$ is non-empty, and that $B = (X \cup Y, A) \in \mathcal{B}(\mathbf{d}, k)$ is a random bipartite graph according to the uniform distribution $\pi_{\mathcal{B}}$. Then, suppose that there are constants c_1 and c_2 such that for any $y, y' \in Y$ and any subset $\mathcal{W} \subseteq X$ of size k ,*

$$\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{N}_B(y) = \mathcal{W}) \leq c_1 \cdot \binom{n}{k}^{-1}$$

and

$$\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{N}_B(y') = \mathcal{W} \mid \mathcal{N}_B(y) = \mathcal{W}) \leq c_2 \cdot \mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{N}_B(y') = \mathcal{W}).$$

Then

$$\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k)) \geq 1 - c_1 c_2 \binom{m}{2} \binom{n}{k}^{-1}.$$

Proof. Let B be an element from the set $\mathcal{B}(\mathbf{d}, k)$ drawn uniformly at random, and let $X^{\{k\}}$ be the set of all k -subsets of X . We omit the subscript $\pi_{\mathcal{B}}$ on all following probabilities. For $1 \leq s < t \leq m$, we define the random variable

$$Z_{st} = \begin{cases} 1 & \text{if } \mathcal{N}_B(y_s) = \mathcal{N}_B(y_t) \\ 0 & \text{otherwise,} \end{cases}$$

for $y_s, y_t \in Y$, and we let

$$Z = \sum_{1 \leq s < t \leq m} Z_{st}$$

be the random variable denoting the number of pairs of nodes (y_s, y_t) that have the same neighborhood in B . Note that

$$\mathbb{P}(B \in \mathcal{B}^*(\mathbf{d}, k)) = 1 - \mathbb{P}(Z \geq 1). \quad (3.1)$$

Using the union bound over all possible pairs of nodes (y_s, y_t) ,

$$\mathbb{P}(Z \geq 1) \leq \sum_{1 \leq s < t \leq m} \mathbb{P}(Z_{st} = 1).$$

Likewise, for a fixed pair of nodes (y_s, y_t) , applying the union bound over $X^{\{k\}}$ shows us that

$$\mathbb{P}(Z_{st} = 1) \leq \sum_{\mathcal{W} \in X^{\{k\}}} \mathbb{P}(\mathcal{N}_B(y_s) = \mathcal{N}_B(y_t) = \mathcal{W}).$$

Now, using the law of total probability,

$$\begin{aligned} \mathbb{P}(Z \geq 1) &\leq \sum_{1 \leq s < t \leq m} \sum_{\mathcal{W} \in X^{\{k\}}} \mathbb{P}(\mathcal{N}_B(y_s) = \mathcal{N}_B(y_t) = \mathcal{W}) \\ &= \sum_{1 \leq s < t \leq m} \sum_{\mathcal{W} \in X^{\{k\}}} \mathbb{P}(\mathcal{N}_B(y_s) = \mathcal{W} \mid \mathcal{N}_B(y_t) = \mathcal{W}) \cdot \mathbb{P}(\mathcal{N}_B(y_t) = \mathcal{W}) \\ &\leq c_1 c_2 \binom{m}{2} \binom{n}{k}^{-1}. \end{aligned} \quad (3.2)$$

The proof is completed by combining (3.1) and (3.2). \square

4 Probability of H-simplicity

The expected running time of HYPERGRAPH SAMPLING is governed by the runtime of algorithm \mathcal{A} and the probability that a randomly chosen element of $\mathcal{B}(\mathbf{d}, k)$ is H-simple. In Section 4.1 we provide an asymptotic estimate which holds when \mathbf{d} is irregular and sparse. In Section 4.2 we give a combinatorial argument for the case of d -regular k -uniform hypergraphs. In particular, these sections yield Theorem 1.6 and Theorem 1.4, respectively.

4.1 Irregular, sparse degrees

In this section we prove a lower bound on the probability that a uniformly random graph from $\mathcal{B}(\mathbf{d}, k)$ is H-simple, using an asymptotic formula for irregular, sparse bipartite graphs. Given a bipartite degree sequence (\mathbf{d}, \mathbf{k}) , for any positive integer r define $M_r = \sum_{i=1}^n (d_i)_r$ and $L_r = \sum_{j=1}^m (k_j)_r$, where $(a)_b = a(a-1)\cdots(a-b+1)$ denotes the falling factorial. Let $L = L_1$ and $M = M_1$, and note that $L = M$ for any graphical bipartite degree sequence. We also let $d_{\max} = \max_i d_i$ and $k_{\max} = \max_j k_j$.

The following asymptotic enumeration result is a simpler, but weaker restatement of the main theorem from [18], which is slightly stronger than that of McKay [23]. (It follows from Theorem 4.1 that $\mathcal{B}(\mathbf{d}, \mathbf{k}) \neq \emptyset$ when $k_{\max} d_{\max} = o(M^{2/3})$.)

Theorem 4.1. ([18, Theorem 1.3]) *Suppose that $M \rightarrow \infty$, and that $\mathbf{d} = (d_1, \dots, d_n)$, $\mathbf{k} = (k_1, \dots, k_m)$ are sequences of nonnegative integer functions of M which both sum to M . If $k_{\max} d_{\max} = o(M^{2/3})$ then*

$$|\mathcal{B}(\mathbf{d}, \mathbf{k})| = \frac{M!}{\prod_{i=1}^n d_i! \prod_{j=1}^m k_j!} \exp \left(-\frac{M_2 L_2}{2M^2} + O \left(\frac{d_{\max}^2 k_{\max}^2}{M} \right) \right).$$

Using this enumeration result, we can prove the main result of this section. First some useful identities which will be used in the proof without further comment: if $|\eta| < 1$ then $\exp(\eta) = 1 + O(\eta)$ and $(1 + \eta)^{-1} = 1 + O(\eta)$. Also observe that if $k^2 = o(M)$ then

$$\frac{M^k}{(M)_k} \leq \frac{M^k}{(M - k)^k} = (1 - k/M)^{-k} = \exp(O(k^2/M)) = 1 + O(k^2/M). \quad (4.1)$$

Theorem 4.2. *Assume that for each positive integer n we have an integer $k = k(n) \geq 3$ and a sequence $\mathbf{d} = \mathbf{d}(n) = (d_1, \dots, d_n)$ of positive integers such that $M = \sum_{i=1}^n d_i$ tends to infinity with n . Assume that $k^2 d_{\max}^2 = o(M)$ and let $\pi_{\mathcal{B}}$ be the uniform distribution on $\mathcal{B}(\mathbf{d}, k)$. Then*

$$\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k)) \geq 1 - \frac{n^k d_{\max}^k}{M^k} \cdot \binom{M/k}{2} \binom{n}{k}^{-1} \cdot (1 + o(1)).$$

Proof. Let $m = M/k$, which by assumption is an integer. Then, suppose that we have some neighbourhood $\mathcal{W} \in X^{\{k\}}$, and two integers $s, t \in [m]$. We will prove the desired result by conditioning on the neighbourhoods of y_s and y_t being equal to \mathcal{W} , at which point we can

apply Lemma 3.1. To do this, we first define two bipartite degree sequences $(\mathbf{d}', \mathbf{k}')$ and $(\mathbf{d}'', \mathbf{k}'')$, as follows:

$$d'_i = \begin{cases} d_i - 1 & \text{if } x_i \in \mathcal{W}, \\ d_i & \text{if } x_i \in X \setminus \mathcal{W}, \end{cases} \quad \text{and} \quad k'_j = \begin{cases} 0 & \text{if } j = s, \\ k & \text{if } j \in [m] \setminus \{s\}, \end{cases}$$

and

$$d''_i = \begin{cases} d_i - 2 & \text{if } x_i \in \mathcal{W}, \\ d_i & \text{if } x_i \in X \setminus \mathcal{W}, \end{cases} \quad \text{and} \quad k''_j = \begin{cases} 0 & \text{if } j \in \{s, t\}, \\ k & \text{if } j \in [m] \setminus \{s, t\}. \end{cases}$$

We also extend the notation for M and M_2 to \mathbf{d}' and \mathbf{d}'' by appending one or two dashes, and likewise for L and L_2 . By assumption, $k^2 d_{\max}^2 = o(M)$, which implies that $kd_{\max} = o(M^{1/2})$. Hence, the conditions of Theorem 4.1 are satisfied, and we can approximate both $|\mathcal{B}(\mathbf{d}, \mathbf{k})|$ and $|\mathcal{B}(\mathbf{d}', \mathbf{k}')|$. Considering the ratio of these, since $d'_i = d_i$ whenever $x_i \notin \mathcal{W}$, and $k'_j = k_j$ whenever $j \neq s$, many factors cancel, leading to

$$\frac{|\mathcal{B}(\mathbf{d}', \mathbf{k}')|}{|\mathcal{B}(\mathbf{d}, \mathbf{k})|} = \frac{k!}{(M)_k} \cdot \prod_{x_i \in \mathcal{W}} d_i \cdot \exp\left(\frac{M_2 L_2}{2M^2} - \frac{M'_2 L'_2}{2(M')^2} + O(k^2 d_{\max}^2 / M)\right). \quad (4.2)$$

Next, by definition of M'_2 , we see that

$$M'_2 = M_2 - \sum_{x_i \in \mathcal{W}} (d_i)_2 + \sum_{x_i \in \mathcal{W}} (d_i - 1)_2 = M_2 - 2 \sum_{x_i \in \mathcal{W}} (d_i - 1) = M_2 (1 - O(kd_{\max} M_2^{-1})).$$

Similarly,

$$M' = M - k = M (1 - O(kM^{-1})), \quad L'_2 = (k - 1) M' = L_2 (1 - O(kM^{-1})).$$

Then

$$\begin{aligned} \frac{M_2 L_2}{2M^2} - \frac{M'_2 L'_2}{2(M')^2} &\leq \frac{M_2 L_2}{2M^2} - \frac{M_2 L_2}{2M^2} \cdot (1 - O(kd_{\max} M_2^{-1} + kM^{-1})) \\ &= O(kd_{\max} L_2 M^{-2} + kM_2 L_2 M^{-3}) \\ &= O\left(\frac{k^2 d_{\max}}{M}\right). \end{aligned}$$

The final equality follows as $M_2 \leq d_{\max} M$ and $L_2 = (k - 1)M$. Therefore, combining the above identities with (4.1) and (4.2) implies that

$$\frac{|\mathcal{B}(\mathbf{d}', \mathbf{k}')|}{|\mathcal{B}(\mathbf{d}, \mathbf{k})|} = \frac{k!}{M^k} \cdot \prod_{x_i \in \mathcal{W}} d_i \cdot (1 + O(k^2 d_{\max}^2 / M)). \quad (4.3)$$

Next, observe that there is a bijective relationship between bipartite graphs $B_0 \in \mathcal{B}(\mathbf{d}', \mathbf{k}')$, and bipartite graphs $B \in \mathcal{B}(\mathbf{d}, \mathbf{k})$ such that $\mathcal{N}_B(y_s) = \mathcal{W}$, using the map $B \mapsto B_0 = B \setminus \{y_s\}$ which deletes vertex y_s and reduces the degrees of each neighbour of y_s by 1. Hence,

$$\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{N}_B(y_s) = \mathcal{W}) = \frac{|\mathcal{B}(\mathbf{d}', \mathbf{k}')|}{|\mathcal{B}(\mathbf{d}, \mathbf{k})|}. \quad (4.4)$$

By assumption, $k^2 d_{\max}^2 = o(M)$, so using (4.3) we find that

$$\binom{n}{k} \cdot \mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{N}_B(y_s) = \mathcal{W}) \leq \frac{n^k d_{\max}^k}{M^k} \cdot (1 + o(1)). \quad (4.5)$$

We can use a similar argument to prove a bound on the conditional edge probability needed for Lemma 3.1. Observe that if \mathbf{d} and \mathbf{k} satisfy the conditions of Theorem 4.1 then so do \mathbf{d}'' and \mathbf{k}'' . Hence, we can repeat the argument that produces (4.4) followed by the substitutions used above to see that

$$\begin{aligned} \mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{N}_B(y_t) = \mathcal{W} \mid \mathcal{N}_B(y_s) = \mathcal{W}) &= \frac{|\mathcal{B}(\mathbf{d}'', \mathbf{k}'')|}{|\mathcal{B}(\mathbf{d}', \mathbf{k}')|} \\ &= \frac{k!}{(M-k)^k} \cdot \prod_{x_i \in \mathcal{W}} (d_i - 1) \cdot (1 + O(k^2 d_{\max}^2 / M)). \end{aligned}$$

We will divide this expression by the one given in (4.3) to obtain, using (4.4),

$$\begin{aligned} \frac{\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{N}_B(y_t) = \mathcal{W} \mid \mathcal{N}_B(y_s) = \mathcal{W})}{\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{N}_B(y_t) = \mathcal{W})} &= \frac{(M)_k}{(M-k)_k} \cdot \prod_{x_i \in \mathcal{W}} \left(1 - \frac{1}{d_i}\right) \cdot (1 + O(k^2 d_{\max}^2 / M)) \\ &\leq \frac{(M)_k}{(M-k)_k} \cdot (1 + O(k^2 d_{\max}^2 / M)) \\ &= 1 + o(1), \end{aligned}$$

using (4.1) and the assumption that $k^2 d_{\max}^2 = o(M)$. From the above inequality and (4.5), we can apply Lemma 3.1 with

$$c_1 = \frac{n^k d_{\max}^k}{M^k} \cdot (1 + o(1)) \quad \text{and} \quad c_2 = 1 + o(1),$$

to complete the proof. □

4.2 Regular degrees

In this section we present a combinatorial argument to establish a lower bound on the probability that a uniformly random graph from $\mathcal{B}(\mathbf{d}, k)$ is H-simple, when $\mathbf{d} = (d, d, \dots, d)$ is regular. We first prove a ‘sensitivity result’ for the set of all bipartite graphs with given degrees. We show that adjusting the degrees on one side of the bipartition, to make them closer to regular, can only increase the number of bipartite graphs. It is possible that this result is known, though we could not find it in the literature. We give a proof in Section 4.2.1 for completeness.

Proposition 4.3. *Let $n, m \in \mathbb{N}$ and let (\mathbf{d}, \mathbf{k}) be a bipartite degree sequence for the bipartition $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_m\}$. Suppose that we have integers $g, h \in [n]$ such that $d_g \geq d_h + 2$ and define \mathbf{d}' by*

$$d'_i = \begin{cases} d_i - 1 & \text{if } i = g \\ d_i + 1 & \text{if } i = h \\ d_i & \text{if } i \in [n] \setminus \{g, h\}. \end{cases}$$

Then

$$|\mathcal{B}(\mathbf{d}, \mathbf{k})| \leq |\mathcal{B}(\mathbf{d}', \mathbf{k})|.$$

Using this proposition, we now prove Theorem 1.4.

Proof of Theorem 1.4. Let $B = (X \cup Y, A) \in \mathcal{B}(n, d, k)$ so that all nodes in X are d -regular and all nodes in Y are k -regular. Throughout the proof we let $\mathcal{W} \in X^{\{k\}}$ be a fixed neighbourhood of size k , and we consider two fixed nodes $y_s, y_t \in Y$.

We will prove the desired result by conditioning on the neighbourhoods of y_s and y_t being equal to \mathcal{W} , at which point we can apply Lemma 3.1. To do this, let $\mathcal{U} \in X^{\{k\}}$ be any k -subset of X . We will analyse

$$\mathbb{P}_{\pi_B}(\mathcal{N}_B(y_s) = \mathcal{U} \mid \mathcal{N}_B(y_t) = \mathcal{W}). \quad (4.6)$$

Our goal will be to show that (4.6) achieves a minimum at $\mathcal{U} = \mathcal{W}$. Let Δ denote the symmetric difference operator. Given $\mathcal{W} = \mathcal{N}_B(y_t)$, for any subset $\mathcal{U} \subseteq X^{\{k\}}$ of size k , we define a new bipartite degree sequence $(\mathbf{d}^{\mathcal{U}}, \mathbf{k}^{\mathcal{U}})$ by

$$d_i^{\mathcal{U}} = \begin{cases} d_i - 2 & \text{if } x_i \in \mathcal{U} \cap \mathcal{W} \\ d_i - 1 & \text{if } x_i \in \mathcal{U} \Delta \mathcal{W} \\ d_i & \text{if } x_i \in X \setminus (\mathcal{U} \cup \mathcal{W}) \end{cases} \quad \text{and} \quad k_j^{\mathcal{U}} = \begin{cases} 0 & \text{if } y_j \in \{y_s, y_t\} \\ k & \text{if } y_j \in Y \setminus \{y_s, y_t\} \end{cases}.$$

Now, when $\mathcal{U} \in X^{\{k\}}$ and $|\mathcal{U} \Delta \mathcal{W}| > 0$, we can select a node $x_g \in \mathcal{U} \setminus \mathcal{W}$ and $x_h \in \mathcal{W} \setminus \mathcal{U}$, and create a new neighbourhood $\mathcal{U}' = (\mathcal{U} \cup \{x_h\}) \setminus \{x_g\}$. Then, we see that $d^{\mathcal{U}'}$ is equal to $d^{\mathcal{U}}$ whenever $i \notin \{g, h\}$, and that $d^{\mathcal{U}'}$ is a more locally balanced degree sequence than $d^{\mathcal{U}}$. Hence, $d^{\mathcal{U}'}$ and $d^{\mathcal{U}}$ satisfy the conditions of Proposition 4.3, and applying Proposition 4.3 we conclude that

$$|\mathcal{B}(\mathbf{d}^{\mathcal{U}'}, \mathbf{k}^{\mathcal{U}'})| \leq |\mathcal{B}(\mathbf{d}^{\mathcal{U}}, \mathbf{k}^{\mathcal{U}})|.$$

Since $|\mathcal{U} \Delta \mathcal{W}| > 0$ is at most $2k$, we can repeat the above process a finite number of times to conclude that for any $\mathcal{U} \in X^{\{k\}}$,

$$|\mathcal{B}(\mathbf{d}^{\mathcal{W}}, \mathbf{k}^{\mathcal{W}})| \leq |\mathcal{B}(\mathbf{d}^{\mathcal{U}}, \mathbf{k}^{\mathcal{U}})|. \quad (4.7)$$

Since

$$\frac{|\mathcal{B}(\mathbf{d}^{\mathcal{U}}, \mathbf{k}^{\mathcal{U}})|}{|\mathcal{B}(\mathbf{d}, \mathbf{k})|} = \mathbb{P}_{\pi_B}(\mathcal{N}_B(y_t) = \mathcal{W}) \cdot \mathbb{P}_{\pi_B}(\mathcal{N}_B(y_s) = \mathcal{U} \mid \mathcal{N}_B(y_t) = \mathcal{W},)$$

the inequality in (4.7) implies that for any $\mathcal{U}, \mathcal{W} \in X^{\{k\}}$, we have

$$\mathbb{P}_{\pi_B}(\mathcal{N}_B(y_s) = \mathcal{W} \mid \mathcal{N}_B(y_t) = \mathcal{W}) \leq \mathbb{P}_{\pi_B}(\mathcal{N}_B(y_s) = \mathcal{U} \mid \mathcal{N}_B(y_t) = \mathcal{W}). \quad (4.8)$$

But $X^{\{k\}}$ has size $\binom{n}{k}$, so summing over all possible choices for $\mathcal{U} \in X^{\{k\}}$ in (4.8) shows us that (since the probabilities on the right must sum to unity)

$$\mathbb{P}_{\pi_B}(\mathcal{N}_B(y_s) = \mathcal{W} \mid \mathcal{N}_B(y_t) = \mathcal{W}) \leq \left(\frac{n}{k}\right)^{-1}.$$

Since \mathbf{d} and \mathbf{k} are both regular, by symmetry every possible $\mathcal{W} \in X^{\{k\}}$ is equally likely; hence

$$\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{N}_B(y_t) = \mathcal{W}) = \binom{n}{k}^{-1}.$$

Thus we can apply Lemma 3.1 with $c_1 = c_2 = 1$ to conclude that

$$\mathbb{P}_{\pi_{\mathcal{B}}}(\mathcal{B}^*(\mathbf{d}, k)) \geq 1 - \binom{m}{2} \binom{n}{k}^{-1}.$$

This completes the proof. \square

4.2.1 Sensitivity result for bipartite degree sequences

In this section we prove Proposition 4.3.

Proof of Proposition 4.3. Let $n, m \in \mathbb{N}$ and let (\mathbf{d}, \mathbf{k}) be a bipartite degree sequence for the bipartition X, Y . Suppose that $g, h \in [n]$ are chosen such that $d_g \geq d_h + 2$. We emphasise that g and h are fixed throughout the proof. Let Δ denote the symmetric difference operator. For any $B = (V, E)$ such that $B \in \mathcal{B}(\mathbf{d}, \mathbf{k})$, let

$$R_2(B) = \mathcal{N}_B(x_g) \cap \mathcal{N}_B(x_h), \quad R_1(B) = \mathcal{N}_B(x_g) \Delta \mathcal{N}_B(x_h),$$

and let $S(B)$ be defined by

$$S(B) = (\{x_g, x_h\} \times R_2(B)) \cup \{e \in E : e \cap \{x_g, x_h\} = \emptyset\}.$$

We see that $S(B)$ is the set of all edges from B which contain vertices $y \in Y$ that are adjacent to both x_g and x_h , or adjacent to neither of these vertices. Next, let \mathcal{S} be the collection of possible values for $S(B)$:

$$\mathcal{S}(\mathbf{d}, \mathbf{k}) = \{S(B) : B \in \mathcal{B}(\mathbf{d}, \mathbf{k})\}.$$

Observe that $\mathcal{S}(\mathbf{d}, \mathbf{k}) \subseteq \mathcal{S}(\mathbf{d}', \mathbf{k})$, since in any $B \in \mathcal{B}(\mathbf{d}, \mathbf{k})$ with $S(B) = S_0$, we may create a bipartite graph $B' \in \mathcal{B}(\mathbf{d}', \mathbf{k})$ with $S(B') = S_0$ by deleting the edge $x_g y$ and replacing it with $x_h y$, for some $y \in \mathcal{N}_B(x_g)$.

For any $S_0 \in \mathcal{S}(\mathbf{d}, \mathbf{k})$, we let $\mathcal{T}(S_0; \mathbf{d}, \mathbf{k})$ be the subset of $\mathcal{B}(\mathbf{d}, \mathbf{k})$ defined by

$$\mathcal{T}(S_0; \mathbf{d}, \mathbf{k}) = \{B \in \mathcal{B}(\mathbf{d}, \mathbf{k}) : S(B) = S_0\}.$$

Then the set $\{\mathcal{T}(S_0; \mathbf{d}, \mathbf{k}) : S_0 \in \mathcal{S}(\mathbf{d}, \mathbf{k})\}$ forms a partition of $\mathcal{B}(\mathbf{d}, \mathbf{k})$. Similarly, the set $\{\mathcal{T}(S_0; \mathbf{d}', \mathbf{k}) : S_0 \in \mathcal{S}(\mathbf{d}, \mathbf{k})\}$ forms a partition of $\mathcal{B}(\mathbf{d}', \mathbf{k})$.

We claim that $|\mathcal{T}(S_0; \mathbf{d}, \mathbf{k})| \leq |\mathcal{T}(S_0; \mathbf{d}', \mathbf{k})|$ for every $S_0 \in \mathcal{S}(\mathbf{d}, \mathbf{k})$. To see this, first note that for any $B \in \mathcal{B}(\mathbf{d}, \mathbf{k})$, the sets $R_1(B)$ and $R_2(B)$ depend only on $S(B)$. Hence we will instead write $R_1(S_0)$, $R_2(S_0)$ for these sets, where $S_0 = S(B)$. Further, for every $B \in \mathcal{T}(S_0; \mathbf{d}, \mathbf{k})$ the degree of vertex $x_i \in X$ in $E \setminus S_0$ is

$$\begin{cases} d_i - |R_2(S_0)| & \text{if } i \in g, h, \\ 0 & \text{otherwise,} \end{cases}$$

and vertex $y_j \in Y$ has degree 1 if $j \in R_1(S_0)$, or 0 otherwise. Next, let S'_0 be any set of bipartite edges satisfying this degree sequence. Then $S_0 \cap S'_0 = \emptyset$ by construction, and the combined bipartite degree sequence of $S_0 \cup S'_0$ is (\mathbf{d}, \mathbf{k}) . Hence, the bipartite graph with edges $S_0 \cup S'_0$ is an element of $\mathcal{T}(S_0; \mathbf{d}, \mathbf{k})$, and every element of $\mathcal{T}(S_0; \mathbf{d}, \mathbf{k})$ is created exactly once in this way. It follows that for all $S_0 \in \mathcal{S}(\mathbf{d}, \mathbf{k})$,

$$|\mathcal{T}(S_0; \mathbf{d}, \mathbf{k})| = \binom{d_g + d_h - 2|R_2(S_0)|}{d_g - |R_2(S_0)|} \leq \binom{d'_g + d'_h - 2|R_2(S_0)|}{d'_g - |R_2(S_0)|} = |\mathcal{T}(S_0; \mathbf{d}', \mathbf{k})|,$$

since $d'_g = d_g - 1 \geq d_h + 1 = d'_h$. This verifies the claimed inequality. Finally, we see that

$$|\mathcal{B}(\mathbf{d}, \mathbf{k})| = \sum_{S_0 \in \mathcal{S}(\mathbf{d}, \mathbf{k})} |\mathcal{T}(S_0; \mathbf{d}, \mathbf{k})| \leq \sum_{S_0 \in \mathcal{S}(\mathbf{d}, \mathbf{k})} |\mathcal{T}(S_0; \mathbf{d}', \mathbf{k})| \leq |\mathcal{B}(\mathbf{d}', \mathbf{k})|,$$

which completes the proof of Proposition 4.3. □

Acknowledgements

We are grateful to the referees for their helpful comments.

Research was supported by the Netherlands Organisation for Scientific Research (NWO) through Gravitation Programme Networks 024.002.003. Martin Dyer is supported by the EPSRC research grant EP/S016562/1 “Sampling in hereditary classes”. Catherine Greenhill is supported by the Australian Research Council Discovery Project DP190100977.

Pieter Kler acknowledges that part of this work was carried out while he was a postdoctoral fellow at the Max Planck Institute for Informatics in Saarbrücken, Germany.

References

- [1] G. Amanatidis and P. Kler, Rapid mixing of the switch Markov chain for strongly stable degree sequences, *Random Structures & Algorithms*, **57(3)** (2020), 637–657.
- [2] N. A. Arafat, D. Basu, L. Deceusefond and S. Bressan, Construction and random generation of hypergraphs with prescribed degree and dimension sequences, in *Database and Expert Systems Applications (DEXA 2020)*, Lecture Notes in Computer Science vol. 12392, Springer, Berlin, 2021, pp. 130–145.
- [3] A. Arman, P. Gao and N. Wormald, Fast Uniform Generation of Random Graphs with Given Degree Sequences, *Proceedings of IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS 2019)*, IEEE, 2019, pp. 1371–1379.
- [4] M. Bayati, J.-H. Kim and A. Saberi, A sequential algorithm for generating random graphs, *Algorithmica* **58(4)** (2010), 860–910.
- [5] I. Bezáková, N. Bhatnagar, E. Vigoda, Sampling binary contingency tables with a greedy start, *Random Structures & Algorithms* **30(1-2)** (2007), 168–205.

- [6] V. Blinovskiy and C. Greenhill, Asymptotic enumeration of sparse uniform linear hypergraphs with given degrees, *Electronic Journal of Combinatorics* **23(3)** (2016), #P3.17.
- [7] C. J. Carstens and P. Kleer, Speeding up switch Markov chains for sampling bipartite graphs with given degree sequence, in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018*, vol. 116 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018, pp. 36:1–36:18.
- [8] Y. Chen, P. Diaconis, S. P. Holmes and J. S. Liu. Sequential Monte Carlo methods for statistical analysis of tables, *Journal of the American Statistical Association* **100(469)** (2005), 109–120.
- [9] P. S. Chodrow, Configuration models of random hypergraphs, *Journal of Complex Networks* **8(3)** (2019), cnaa018
- [10] C. Cooper, M. Dyer and C. Greenhill. Sampling regular graphs and a peer-to-peer network. *Combinatorics, Probability and Computing* **16(4)** (2007), 557–593.
- [11] C. Cooper, M. E. Dyer and C. Greenhill, Corrigendum: Sampling regular graphs and a peer-to-peer network. [arXiv:1203.6111](https://arxiv.org/abs/1203.6111)
- [12] C. Cooper, A. Frieze, M. Molloy and B. Reed, Perfect matchings in random r -regular, s -uniform hypergraphs, *Combinatorics, Probability and Computing* **5(1)** (1996), 1–14.
- [13] A. Deza, A. Levin, S. M. Meesum and S. Onn, Hypergraphic degree sequences are hard, *Bulletin of the EATCS* **127** (2019).
- [14] P. Diaconis and A. Gangolli, Rectangular arrays with fixed margins, in D. Aldous, P. Diaconis, J. Spencer and J. M. Steele (eds.), *Discrete Probability and Algorithms*, Springer–Verlag, New York, (1995), pp. 15–41.
- [15] A. Dudek, A. Frieze, A. Ruciński, M. Šileikis, Approximate counting of regular hypergraphs, *Information Processing Letters* **113(19-21)** (2013), 785–788.
- [16] P. L. Erdős, C. Greenhill, T. R. Mezei, I. Miklos, D. Soltész, L. Soukup, The mixing time of the swap (switch) Markov chains: a unified approach. Preprint (2019). [arXiv:1903.06600](https://arxiv.org/abs/1903.06600)
- [17] P. Gao and N. Wormald, Uniform generation of random regular graphs, *SIAM Journal on Computing* **46(4)** (2017), 1395–1427.
- [18] C. Greenhill, B. D. McKay and X. Wang, Asymptotic enumeration of sparse 0-1 matrices with irregular row and column sums, *Journal of Combinatorial Theory (Series A)* **113(2)** (2006), 291–324.
- [19] C. Greenhill and M. Sfragara, The switch Markov chain for sampling irregular graphs and digraphs, *Theoretical Computer Science* **719** (2018), 1–20.

- [20] M. Jerrum and A. Sinclair, Fast uniform generation of regular graphs, *Theoretical Computer Science* **73(1)** (1990), 91–100.
- [21] M. Jerrum, A. Sinclair and E. Vigoda, A polynomial-time approximation algorithm for the permanent of a matrix with non-negative entries, *Journal of the ACM* **51(4)** (2004), 671–697.
- [22] R. Kannan, P. Tetali and S. Vempala, Simple Markov chain algorithms for generating random bipartite graphs and tournaments, *Random Structures & Algorithms* **14(4)** (1999), 293–308.
- [23] B. D. McKay, Asymptotics for 0-1 matrices with prescribed line sums, in *Enumeration and Design*, Academic Press, Toronto, 1984, pp. 225–238.
- [24] B. D. McKay and N. C. Wormald, Uniform generation of random regular graphs of moderate degree, *Journal of Algorithms* **11(1)** (1990), 52–67.
- [25] I. Miklos, P. L. Erdős and L. Soukup, Towards random uniform sampling of bipartite graphs with given degree sequence, *Electronic Journal of Combinatorics* **20(1)** (2013), #P16.
- [26] R. Patro, C. Kingsford, Predicting protein interactions via parsimonious network history inference, *Bioinformatics* **29(13)** (2013), i237–i246.
- [27] K. Tikhomirov, P. Youssef, Sharp Poincaré and log-Sobolev inequalities for the switch chain on regular bipartite graphs. Preprint (2020). [arXiv:2007.02729](https://arxiv.org/abs/2007.02729)
- [28] N. D. Verhelst, An efficient MCMC algorithm to sample binary matrices with fixed marginals, *Psychometrika* **73(4)** (2008), 705.
- [29] D. Zhou, J. Huang, B. Schölkopf, Learning with hypergraphs: Clustering, classification, and embedding, in *Advances in Neural Information Processing Systems* **19**, 2006, pp. 1601–1608.

A Mixing time bounds for the switch Markov chain

The switch Markov chain was analysed by Cooper, Dyer and Greenhill [10, 11] for regular graphs, and by Greenhill and Sfragara [19] for irregular graphs that are relatively sparse. Some situations which lead to additional factors in the mixing time bounds from [10, 11, 19] do not arise in bipartite graphs, and so it is possible to improve the mixing time bounds for bipartite graphs. To the best of our knowledge, these bounds have not been presented elsewhere, so we write them down here. The proofs from [10, 11, 19] use the multicommodity flow method, and are quite long and technical. We do not give full details, but rather explain how the proofs from [10, 11, 19] can be adapted to the bipartite setting, and give the resulting mixing time bounds. For all notation that is not defined, and all other missing details, we refer the reader to [10, 11, 19].

We begin with regular bipartite graphs, where there are n nodes in each side of the bipartition and all nodes have degree d . (We stress that this case is not particularly relevant for the problem of sampling hypergraphs, unless the hypergraph is d -regular and d -uniform.) For regular bipartite graphs, the bound on the bipartite switch chain is a factor of $\frac{1}{32} d^6 n^2$ smaller than the general (not-necessarily-bipartite) case. (The constant factor 32 in Theorem A.1 below arises since graphs in $\mathcal{B}(n, d, d)$ have $2n$ nodes.)

Theorem A.1. *Let $\mathcal{B}(n, d, d)$ be the set of d -regular bipartite graphs with $2n$ nodes and a given bipartition. Then the mixing time of the bipartite switch chain on $\mathcal{B}(n, d, d)$ satisfies*

$$\tau(\varepsilon) \leq 32d^{17}n^6(2dn \log(2dn) + \log(\varepsilon^{-1})).$$

Proof. Temporarily, write $N = 2n$, for ease of comparison with [10, 11]. The flow can be defined in exactly the same way as in [10], though a shortcut edge will never be needed. (Every circuit decomposes into 1-circuits.) Hence there will be at most 3 defect edges, two labelled -1 and one labelled 2 , all incident with the “start vertex” x_0 of the 1-circuit. At most two switches are needed to transform any encoding into a graph, and [10, Lemma 4] becomes $|\mathcal{L}(Z)| \leq 2d^4 N^3 |\mathcal{B}(n, d, d)|$. (This is smaller by a factor of $d^2 N^2$ than in the non-bipartite case, essentially because we save one (-1) -switch in the worst case, which costs $d^2 N^2$.)

In the bipartite case, we save a factor of d^4 compared with [11, Lemma 1] (which is a correction of [10, Lemma 5]). This is because there can be at most 10 bad pairs in the yellow-green colouring, not 14. Recall that each bad pair contributes a factor of d . (There are at most 3 bad pairs from each defect edge, plus at most one additional bad pair from wrapping around at x_0 . Alternatively, we no longer have a shortcut edge, which in [10, 11] was responsible for up to 4 bad pairs, so 14 goes down to 10.)

Combining these effects, we save a factor of $d^6 N^2$ compared to the mixing time from [11, Theorem 1]. Note that we have $\ell(f) \leq dN/2$ and, as there are dn edges in any element of $\mathcal{B}(n, d, d)$,

$$1/Q(e) \leq 4 \binom{nd}{2} |\mathcal{B}(n, d, d)| \leq \frac{1}{2} d^2 N^2 |\mathcal{B}(n, d, d)|,$$

which saves an additional factor of $\frac{1}{2}$ compared with [10, 11]. \square

Let $\Delta = \max\{d_{\max}, k_{\max}\}$. By adapting the analysis from [19], we can prove a bound in the irregular case which is a factor of $\Delta^4 M^2$ smaller than in the general (not-necessarily-bipartite) case.

Theorem A.2. *Let $\mathcal{B}(\mathbf{d}, \mathbf{k})$ be the set of all bipartite graphs with a given node bipartition, degrees \mathbf{d} on the left and degrees \mathbf{k} on the right. Suppose that all degrees are at least 1 and that $3 \leq d_{\max}, k_{\max} \leq \frac{1}{3}\sqrt{M}$. Then the mixing time of the bipartite switch chain satisfies*

$$\tau(\varepsilon) \leq \Delta^{10} M^7 \left(\frac{1}{2} M \log(M) + \log(\varepsilon^{-1}) \right).$$

Proof. We have $1/Q(e) \leq M^2 |\mathcal{B}(\mathbf{d}, \mathbf{k})|$ and $\ell(f) \leq M/2$, as in [19, Theorem 1.1]. Arguing as above, the number of bad pairs is at most 10, not 14, saving a factor of Δ^4 . The main

thing is the critical lemma [19, Lemma 2.5], where we give an upper bound on the number of encodings. We claim that

$$|\mathcal{L}^*(Z)| \leq 2M^4 |\mathcal{B}(\mathbf{d}, \mathbf{k})|$$

so long as $\Delta \leq \frac{1}{3}\sqrt{M}$, say.

In [19, Lemma 2.5], we only performed a $(-1, 2)$ -switching if we had four defect edges. This was to ensure that we definitely had a (-1) -defect edge incident with a 2-defect edge: but when there is no shortcut edge, this is already guaranteed when we have three defect edges. Therefore, letting

$$a = 2\Delta^2 M, \quad b = 2\Delta^2, \quad c = \frac{9}{8}M^2$$

be the upper bounds that we proved in [19, Lemma 2.5] on the various ratios, we obtain

$$|\mathcal{L}^*(Z)| \leq (1 + b + c + bc + c^2 + ac) |\mathcal{B}(\mathbf{d}, \mathbf{k})|.$$

(The saving here is replacing bc^2 by ac , and in omitting the terms involving b^2 or abc or ac^2 , which were needed in [19] to deal with the shortcut edge.) Using the bounds $3 \leq \Delta$ and $\Delta^2 \leq M/9$, we see that

$$|\mathcal{L}^*(Z)| \leq 2M^4.$$

This is a factor of M^2 smaller than the corresponding bound given in [19, Lemma 2.5], again because (in the worst case) we save a (-1) -switch, which gives a ratio of $\frac{9}{8}M^2$. Combining this with the earlier saving of Δ^4 , we obtain the stated bound. \square

The following corollary is most relevant to sampling uniform hypergraphs with given degrees. It follows directly from Theorem A.2 by considering a regular sequence \mathbf{k} .

Corollary A.3. *Let $\mathcal{B}(\mathbf{d}, k)$ be the set of all bipartite graphs with a given node bipartition, degrees $\mathbf{d} = (d_1, \dots, d_n)$ on one side and with m nodes of degree k on the other. Let $M = km = \sum_{j=1}^n d_j$. Suppose that all degrees are at least 1 and that $3 \leq d_{\max}, k \leq \frac{1}{3}\sqrt{M}$. Then the mixing time of the bipartite switch chain satisfies*

$$\tau(\varepsilon) \leq \Delta^{10} M^7 \left(\frac{1}{2} M \log(M) + \log(\varepsilon^{-1}) \right)$$

where $\Delta = \max\{d_{\max}, k\}$.