

On the Diaconis-Gangolli Markov Chain for Sampling Contingency Tables with Cell-Bounded Entries

Ivona Bezáková*

Nayantara Bhatnagar†

Dana Randall‡

August 1, 2009

Abstract

The problems of uniformly sampling and approximately counting contingency tables have been widely studied, but efficient solutions are only known in special cases. One appealing approach is the Diaconis and Gangolli Markov chain which updates the entries of a random 2×2 submatrix. This chain is known to be rapidly mixing for cell-bounded tables only when the cell bounds are all 1 and the row and column sums are regular. We demonstrate that the chain can require exponential time to mix in the cell-bounded case, even if we restrict to instances for which the state space is connected. Moreover, we show the chain can be slowly mixing even if we restrict to natural classes of problem instances, including regular instances with cell bounds of either 0 or 1 everywhere, and dense instances where at least a linear number of cells in each row or column have non-zero cell-bounds.

1 Introduction

We consider a popular approach for sampling standard and cell-bounded contingency tables based on a local Markov chain. Standard contingency tables are non-negative matrices consistent with prescribed row and column sums, while *cell-bounded* contingency tables additionally satisfy given cell bounds. More precisely, suppose that we are given a list of positive integers $r = (r_1, \dots, r_m)$ called the *row sums* and another list of positive integers $c = (c_1, \dots, c_n)$ called the *column sums* such that $\sum_{i=1}^m r_i = \sum_{j=1}^n c_j$. The set of *standard contingency tables* $\Sigma_{r,c}$ is the set of $m \times n$ non-negative integer matrices T that satisfy the given row and column sums. In the more general case of *cell-bounded contingency tables*, we are given row sums r , column sums c , and *cell-bounds* b_{ij} with $1 \leq i \leq m$ and $1 \leq j \leq n$ satisfying $b_{ij} \leq \min(r_i, c_j)$, for all i, j . Then the cell-bounded contingency tables $\Sigma_{r,c,b}$ are the $m \times n$ non-negative matrices T such that $\sum_{j=1}^n t_{i,j} = r_i$, $\sum_{i=1}^m t_{i,j} = c_j$, and $0 \leq t_{i,j} \leq b_{i,j}$ for all i, j .

The problem of sampling contingency tables almost uniformly at random was first studied for its applications in statistics [5]. Subsequently, there has been a long line of work focused on sampling and approximately counting contingency tables (see, e.g., [1, 2, 3, 7, 8, 12, 15]). Despite this extensive amount of work, these problems have been solved only in special cases. Dyer, Kannan and Mount [8] and Morris [14] showed that there is an *fpras*, a fully polynomial randomized approximation

*Department of Computer Science, Rochester Institute of Technology, Rochester, NY. Email: ib@cs.rit.edu.

†Department of Statistics, University of California, Berkeley, CA. Email: nayan@stat.berkeley.edu. Supported by DOD ONR grant N0014-07-1-05-06 and DMS 0528488.

‡College of Computing, Georgia Institute of Technology, Atlanta, GA. Email: randall@cc.gatech.edu. Supported in part by NSF grants CCF-0830367 and DMS-0505505.

scheme (see [16]), for approximately counting standard contingency tables when the row and column sums are all sufficiently large using methods for estimating the volume of a convex body. There are also a variety of proofs showing there is an fpras for approximately counting standard tables when there are only a constant number of rows: by Cryan and Dyer [2] using dynamic programming and volume estimation, by Cryan et al. [3] using a local Markov chain, and by Dyer [7] using dynamic programming. Cryan, Dyer and Randall [4] extended these results to the more general setting with cell bounds by showing that there is also an fpras for cell-bounded tables in the same two cases: when the row and column sums are large or when the number of rows is a constant. Their proofs used dynamic programming and volume estimation. In addition to generalizing previous results, this work is significant because cell-bounded tables are *self-reducible*, the class of problems for which there is a known polynomial-time reduction between approximate counting and sampling (see [11]), so the existence of an fpras implies we can efficiently sample cell-bounded tables in these cases as well. The standard contingency table problem is not known to have this useful property.

Our work here focuses specifically on the Markov chain approach introduced by Diaconis and Gangolli [6]. In each step of the chain we choose two rows and two columns at random, and then we modify the entries in this 2×2 submatrix by trying to add 1 (or in general a fixed integer) to two diagonal entries and subtracting the same amount from the other two entries, so long as the resulting matrix is non-negative (and still satisfies the cell-bounds, if they exist). In the case of standard contingency tables, the Diaconis-Gangolli chain always connects the state space and is known to be rapidly mixing when there are a constant number of rows [2]. However, there is no evidence suggesting that this chain is a poor approach to sampling in other situations.

Even less is known for cell-bounded tables. Kannan, Tetali and Vempala [12] showed that the Diaconis-Gangolli chain is rapidly mixing when the cell bounds are all 1 and the row and column sums are regular. However, if the cell bounds are restricted to being either 0 or 1, the chain might not even connect the state space. Consider, for example, the 3×3 case where $r = (1, 1, 1)$, $c = (1, 1, 1)$ and $b_{1,1} = b_{2,2} = b_{3,3} = 0$; note that it is not possible to move between the two valid tables with moves of the Markov chain.

It is natural to ask whether the Diaconis-Gangolli chain is always rapidly mixing when restricted to instances $\Sigma_{r,c,b}$ where the state space is connected. We show that in the cell-bounded setting, even when the chain connects the space, it can require exponential time to converge. While not so surprising, no such evidence exists for standard tables. This may help explain why there has been such limited success showing the chain is fast for large classes of graphs.

Further, we give additional evidence that this chain can be a bad approach to sampling cell-bounded contingency tables by considering two natural special cases of inputs. We show that the chain can require exponential time even in the dense case (the number of cells with positive cell-bounds in each row and column is linear) and the regular case (when all the r_i and c_j are equal). Our proofs of slow mixing are based on demonstrating a bottleneck in the state space that implies that the *conductance* of the Markov chain is very small.

2 Preliminaries

Let $m, n \in \mathbf{N}$ and let $r = (r_1, r_2, \dots, r_m) \in \mathbf{N}^m$ and $c = (c_1, c_2, \dots, c_n) \in \mathbf{N}^n$ be vectors with $\sum_{i=1}^m r_i = \sum_{j=1}^n c_j$. Let $b = b_{i,j} \in \mathbf{N}$ for every $i \in [m], j \in [n]$ with $b_{i,j} \leq \min(r_i, c_j)$. A non-negative integer matrix $(t_{i,j})_{m \times n}$ is a *cell-bounded contingency table* with row sums r , column sums c , and cell-bounds b , if it satisfies:

- (i) for every $i \in [m]$, we have $\sum_{j=1}^n t_{i,j} = r_i$,
- (ii) for every $j \in [n]$, we have $\sum_{i=1}^m t_{i,j} = c_j$, and
- (iii) for every $i \in [m], j \in [n]$, we have $0 \leq t_{i,j} \leq b_{i,j}$.

The row and column sums are also called the *marginals*. Denote the set of all such tables by $\Sigma_{r,c,b}$. Note that the lower bounds on $t_{i,j}$ of 0 could be generalized but this does not change the model, since any lower bound $0 < a_{i,j} \leq b_{i,j}$ can be eliminated by replacing the cell-bound by $b_{i,j} - a_{i,j}$, and making the corresponding row and column sums $r_i - a_{i,j}$ and $c_j - a_{i,j}$, respectively.

The *cell-bounded contingency table problem* is to generate a (uniformly) random table in $\Sigma_{r,c,b}$. We describe below a well-known Markov chain which performs a random walk on the space of cell-bounded tables.

The Diaconis-Gangolli (DG) Markov chain.

Given $T \in \Sigma_{r,c,b}$, the Markov chain performs a step as follows:

1. With probability 1/2 do nothing, i. e., let $T' = T$.
2. Otherwise, choose $i_1 < i_2$ uniformly at random from $[m]$ and $j_1 < j_2$ uniformly at random from $[n]$.
3. With probability 1/2 let $d = 1$; else let $d = -1$.
4. Let $D = (d_{i,j})_{m \times n}$ be a matrix with only four non-zero entries: $d_{i_1, j_1} = d_{i_2, j_2} = d$ and $d_{i_1, j_2} = d_{i_2, j_1} = -d$. If $T + D$ satisfies the cell bounds $b_{i,j}$ and is non-negative, then let $T' = T + D$.
5. Return T' as the new contingency table.

Since this Markov chain is symmetric, if the chain is connected, its stationary distribution π is uniform over $\Sigma_{r,c,b}$. However, it is possible that for certain cell-bounds the Markov chain is not irreducible. In this work we focus on input instances for which the Markov chain is irreducible (and therefore also ergodic).

It is fairly standard to use the mixing time to bound the convergence time of a Markov chain. Let P denote the transition matrix of a Markov chain M with state space Ω . Thus, $P^t(x, \cdot)$ denotes the distribution after t steps of the chain, with starting state x . The *mixing time* $\tau_x(\varepsilon)$ of M starting at state $x \in \Omega$ is

$$\tau_x(\varepsilon) = \min\{t \geq 0 \mid d_{tv}(P^t(x, \cdot), \pi) \leq \varepsilon\},$$

where $d_{tv}(\mu, \nu) = \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \nu(x)|$ is the *total variation distance*. The mixing time of M is $\tau(\varepsilon) = \max_{x \in \Omega} \tau_x(\varepsilon)$.

The *conductance* Φ_M of an ergodic Markov chain $M = (\Omega, P)$ with stationary distribution π is defined as follows:

$$\Phi_M = \min_{S \subseteq \Omega, \pi(S) \leq 1/2} \frac{\sum_{s_1 \in S, s_2 \in \Omega \setminus S} \pi(s_1) P(s_1, s_2)}{\pi(S)}.$$

The following theorem relates conductance to the mixing time.

Theorem 1 ([9, 13]). *Let M be a Markov chain on Ω such that $M(u, u) \geq \frac{1}{2}$ for every $u \in \Omega$ and let $\pi_{\min} = \min_{x \in \Omega} \pi(x)$. Then,*

$$\frac{1}{2} \left(\frac{1}{2\Phi_M} - 1 \right) \log \left(\frac{1}{2\varepsilon} \right) \leq \tau(\varepsilon) \leq \frac{2}{\Phi_M^2} \log \left(\frac{1}{\pi_{\min}\varepsilon} \right).$$

3 Slow Mixing of the Diaconis-Gangolli Chain

Our main theorems demonstrate that the Diaconis-Gangolli chain can have exponentially large mixing time in the cell-bounded case. Before stating our results, it will be useful to present a graph interpretation of cell-bounded tables when all the cell bounds $b_{i,j}$ are 0 or 1. Let $G = (V_1, V_2, E)$ be a bipartite graph with partition sizes m, n and adjacency matrix $B = \{b_{i,j}\}$, i. e., $V_1 = \{v_{1,1}, \dots, v_{m,1}\}$, $V_2 = \{v_{1,2}, \dots, v_{n,2}\}$, and $E = \{(v_{i,1}, v_{j,2}) \mid b_{i,j} > 0\}$. Then a cell-bounded contingency table T satisfying row sums r_i , column sums c_j and bounds $b_{i,j}$ can be interpreted as a subgraph of G with degree requirements $\text{deg}_T(v_{i,1}) = r_i$ for every $i \in [m]$ and $\text{deg}_T(v_{j,2}) = c_j$ for every $j \in [n]$.

3.1 Dense Instances

We show there is a family of *dense* instances, i.e., where at least a linear number of cell bounds in each row and column are non-zero, for which the Diaconis-Gangolli chain connects the space but mixes exponentially slowly. Dense instances are significant because they are close to the graphical case with all $b_{i,j} = 1$.

Theorem 2. *There exists a family of $n \times n$ instances of the cell-bounded contingency table problem such that*

- (i) *the upper-bound on each cell is either 0 or 1,*
- (ii) *the Diaconis-Gangolli chain connects the state space,*
- (iii) *the number of cells with non-zero upper-bound is at least $n/4$ in each row and each column, and*
- (iv) *the mixing time of the chain is greater than an exponential in n .*

Proof. Since all the upper-bounds for all cells are at most 1, we can use the graph representation of the problem. Before we describe the family of $n \times n$ instances, we introduce an important building block for each instance, a complete bipartite graph on $k + k$ vertices: $H_k = (W_{k,1}, W_{k,2}, F_k)$ where $W_{k,\ell} = \{w_{1,\ell}, w_{2,\ell}, \dots, w_{k,\ell}\}$ for $\ell \in \{1, 2\}$, and $F_k = W_{k,1} \times W_{k,2}$.

An $n \times n$ instance G_n is a bipartite graph consisting of four copies of H_k , let us call them $H_k^{(1)}$, $H_k^{(2)}$, $H_k^{(3)}$, and $H_k^{(4)}$, where all the vertices are distinct except for $w_{1,2}^{(j)} = w_{1,1}^{(j+1)}$ for $j \in [3]$ and $w_{1,2}^{(4)} = w_{1,1}^{(1)}$ (and the vertices $w_{1,1}^{(j)}$, $j \in [4]$, are four distinct vertices), see Figure 1. We also need to specify the row and column sums, or, alternatively for the graph representation, the required degrees (see Figure 1): for $\ell \in [2]$ and $j \in [4]$,

$$\text{deg}(w_{i,\ell}^{(j)}) = \begin{cases} i - 1 & \text{for } i > 1, \\ 1 & \text{for } i = 1. \end{cases}$$

Notice that part (i) of the theorem follows from the construction and since $n = 4k - 2$, part (iii) is also immediate. Before we prove the remaining two parts, let us show the following claim:

Claim 3. *Any subgraph G' of G_n that satisfies the required degree sequence is exactly one of the following two types:*

- a) *both $w_{1,1}^{(1)}$ and $w_{1,2}^{(1)}$ are connected to vertices in $H_k^{(1)}$ in G' and both $w_{1,1}^{(3)}$ and $w_{1,2}^{(3)}$ are connected to vertices in $H_k^{(3)}$ in G' , or*




Figure 1: The graph representation of the input instances from Theorem 2.

b) both $w_{1,1}^{(2)}$ and $w_{1,2}^{(2)}$ are connected to vertices in $H_k^{(2)}$ in G' and both $w_{1,1}^{(4)}$ and $w_{1,2}^{(4)}$ are connected to vertices in $H_k^{(4)}$ in G' .

Proof. We know that $w_{1,1}^{(1)} = w_{1,2}^{(4)}$ has exactly one neighbor in G' and by the definition of G_n , this neighbor is either in $H_k^{(1)}$ or $H_k^{(4)}$. Suppose first that the neighbor is in $H_k^{(1)}$. Then we claim that the (only) neighbor of $w_{1,2}^{(1)}$ must also be in $H_k^{(1)}$. Suppose for contradiction that $w_{1,2}^{(1)}$ is connected to a vertex outside of $H_k^{(1)}$ (i. e., it is connected to a vertex in $H_k^{(2)}$). Let us consider the part $G'^{(1)}$ of G' that is a subgraph of $H_k^{(1)}$. Then the degrees (in $G'^{(1)}$) of the vertices in $W_{k,1}^{(1)}$ are $(1, 1, 2, \dots, k-1)$ whereas the degrees of the vertices in $W_{k,2}^{(1)}$ are $(0, 1, 2, \dots, k-1)$. However, this is not possible since the sum of the degrees of the vertices in $W_{k,1}^{(1)}$ must be equal to the sum of the degrees of the vertices in $W_{k,2}^{(1)}$ (and it equals the number of edges of $G'^{(1)}$). Thus, $w_{1,1}^{(1)}$ is connected to a vertex in $H_k^{(1)}$ if and only if $w_{1,2}^{(1)}$ is connected to a vertex in $H_k^{(1)}$. The statement of the claim follows by symmetry of $H_k^{(1)}, \dots, H_k^{(4)}$. \square

Part (ii) of the theorem follows from a proof by Kannan et al. [12] who show that the Diaconis-Gangolli chain connects the state space for every input instance consisting of the complete graph and all cell-bounds equal to 1. In our case the graph is not the complete graph; however, the H_k 's are complete graphs. Now suppose we want to move from a graph G'_1 to a graph G'_2 using Diaconis-Gangolli (DG) moves. Without loss of generality assume that the vertex $w_{1,1}^{(1)}$ is connected to a vertex in $H_k^{(1)}$. Then, by the above claim, G'_1 satisfies a) while G'_2 satisfies either a), or b). We will consider both cases: if G'_2 satisfies b), then we will use DG moves to perform the following modifications:

1. Consider the part G''_1 of G'_1 that is a subset of $H_k^{(1)}$ and the part G''_2 of G'_2 that is a subset of $H_k^{(1)}$:

notice that the vertices $w_{1,1}^{(1)}$ and $w_{1,2}^{(1)}$ have degree 0 in G_2'' but all the other degrees are equal. Let G_2^* be identical to G_2'' with the edge $(w_{1,1}^{(1)}, w_{1,2}^{(1)})$ included. We use Kannan et al. to modify G_1'' into G_2^* .

2. Analogously, modify the part of G_1' that is a subset of $H_k^{(3)}$ into the part of G_2' that is a subset of $H_k^{(3)}$, with edge $(w_{1,1}^{(3)}, w_{1,2}^{(3)})$ included.
3. Using a single DG move, swap edges $(w_{1,1}^{(1)}, w_{1,2}^{(1)})$ and $(w_{1,1}^{(3)}, w_{1,2}^{(3)})$ for $(w_{1,1}^{(2)}, w_{1,2}^{(2)})$ and $(w_{1,1}^{(4)}, w_{1,2}^{(4)})$.
4. Modify the part of G_1' that is a subset of $H_k^{(2)}$, with edge $(w_{1,1}^{(2)}, w_{1,2}^{(2)})$ included, into the part of G_2' that is a subset of $H_k^{(2)}$.
5. Finally, modify the part of G_1' that is a subset of $H_k^{(4)}$, with edge $(w_{1,1}^{(4)}, w_{1,2}^{(4)})$ included, into the part of G_2' that is a subset of $H_k^{(4)}$.

If G_2' satisfies a), then the situation is even easier:

1. Using the result of Kannan et al., modify the part of G_1' that is a subset of $H_k^{(1)}$ into the part of G_2' that is subset of $H_k^{(1)}$.
2. Similarly, modify parts of G_1', G_2' that are subsets of $H_k^{(3)}$.
3. Notice that the part of G_1' that is a subset of $H_k^{(2)} \setminus \{w_{1,1}^{(2)}, w_{1,2}^{(2)}\}$ is uniquely given (its degree sequence is $(1, 2, \dots, k-1), (1, 2, \dots, k-1)$ and there is a single graph satisfying this degree sequence); the same holds for the corresponding part of G_2' . Thus, they are identical and no modification is necessary.
4. Again, no modification is necessary for the parts of G_1', G_2' that are subsets of $H_k^{(4)} \setminus \{w_{1,1}^{(4)}, w_{1,2}^{(4)}\}$.

It remains to prove part (iv) of the theorem which will follow from an upper bound on the conductance of the chain. Define the set S to be the set of graphs satisfying part a) of the above claim. We can define a bijection between S and $\Sigma_{r,c,b} \setminus S$ by symmetry, so $\pi(S) = 1/2$. We need to compute $\sum_{s_1 \in S, s_2 \in \Sigma_{r,c,b} \setminus S} \pi(s_1)P(s_1, s_2)$. Suppose that $s_1 \in S$ and there exists $s_2 \in \Sigma_{r,c,b} \setminus S$ such that $P(s_1, s_2) > 0$. Then, s_1 is a graph containing the edges $(w_{1,1}^{(1)}, w_{1,2}^{(1)})$ and $(w_{1,1}^{(3)}, w_{1,2}^{(3)})$. Notice there is a single graph satisfying this condition; moreover, there is a single $s_2 \in \Sigma_{r,c,b} \setminus S$ such that $P(s_1, s_2) > 0$ (more precisely, $P(s_1, s_2) = \frac{1}{(n(n-1))^2}$). The last quantity to bound is $\pi(s_1)$.

Since $\pi(s_1) = 1/|\Sigma_{r,c,b}|$, we need to estimate $|\Sigma_{r,c,b}|$. Let $T(k)$ be the number of subgraphs of H_k with the degree sequence $(1, 1, 2, \dots, k-1), (1, 1, 2, \dots, k-1)$. It is not difficult to see that $T(k) \geq 4T(k-2)$ and $T(2) = 2, T(1) = 1$. (The inequality $T(k) \geq 4T(k-2)$ follows from the fact that the vertices with required degree $k-1$ are connected to all but one of their neighbors. Suppose that the omitted neighbors would be the vertices of required degree 1. Without loss of generality, let the omitted vertices be $w_{1,1}$ and $w_{1,2}$. After updating the required degree, we get that the new required degree sequence is $(0, 1, 1, 2, 3, \dots, k-3), (0, 1, 1, 2, 3, \dots, k-3)$. Thus, since we have two choices for the vertex of degree 1 for both W_1 and W_2 , we get that $T(k) \geq 2^2T(k-2)$.) Therefore, $T(k) \geq 2^{k-1}$.

Now we can estimate $|\Sigma_{r,c,b}|$: if a $G' \in \Sigma_{r,c,b}$ satisfies part a) of the above claim, then the intersection of G' and $H_k^{(2)}$ is fixed, similarly the intersection of G' and $H_k^{(4)}$. The intersection of G' with $H_k^{(1)}$

is simply any subgraph of $H_k^{(1)}$ satisfying the degree sequence $(1, 1, 2, \dots, k-1), (1, 1, 2, \dots, k-1)$; similarly for the intersection of G' with $H_k^{(3)}$. Thus, the graphs satisfying part a) of the above claim contribute at least $(2^{k-1})^2 = 2^{n/2-1}$ to $|\Sigma_{r,c,b}|$. We get the same estimate for graphs satisfying part b), therefore $|\Sigma_{r,c,b}| \geq 2^{n/2}$.

Finally, we bound the conductance as follows:

$$\Phi_M \leq \frac{\sum_{s \in S, t \notin S} \pi(s)P(s,t)}{\pi(S)} = \frac{\frac{1}{|\Sigma_{r,c,b}|} \cdot \frac{1}{n^2(n-1)^2}}{\frac{1}{2}} = \frac{2}{2^{n/2} \cdot n^2(n-1)^2} \leq \frac{1}{2^{n/2}}.$$

Therefore, part (iv) of the theorem follows from Theorem 1. \square

3.2 Regular Instances

Kannan, et al. [12] showed that if the cell-bounds are all 1, then the Diaconis-Gangolli chain mixes rapidly for all regular instances, i. e., the r_i 's and the c_j 's are all identical. What happens if the instance is regular but there are cell bounds? It is not difficult to find instances for which the Diaconis-Gangolli chain does not connect the state space. However, even in those cases when the state space is connected, is it always rapidly mixing? We answer this question negatively.

Theorem 4. *There exists a family of $n \times n$ instances of the cell-bounded contingency table problem such that*

- (i) *the upper-bound on each cell is either 0 or 1,*
- (ii) *the Diaconis-Gangolli chain connects the state space,*
- (iii) *the marginals are regular with $r_i = 1$ and $c_j = 1$ for every $i, j \in [n]$, and*
- (iv) *the mixing time of the chain is at least exponential.*

Proof. Similarly to the instances from Theorem 2, the instances here consist of four identical building blocks. Each building block $H_k = (W_{k,1}, W_{k,2}, F_k)$ is a bipartite graph on $(2k+2) + (2k+2)$ vertices, where $W_{k,\ell} = \{w_{1,\ell}, w_{2,\ell}, \dots, w_{2k+2,\ell}\}$ for $\ell \in \{1, 2\}$, and F_k is defined as follows, see also Figure 2:

$$\begin{aligned} F_k = & \{(w_{1,1}, w_{1,2}), (w_{1,1}, w_{2k+2,2}), (w_{2k+2,1}, w_{1,2}), (w_{2,1}, w_{2k+2,2})\} \cup \\ & \{(w_{1,1}, w_{2i+1,2}) \mid i \in [k]\} \cup \{(w_{2(i+1),1}, w_{2i+1,2}) \mid i \in [k]\} \cup \\ & \{(w_{2i,1}, w_{2i,2}), (w_{2i,1}, w_{2i+1,2}), (w_{2i+1,1}, w_{2i+1,2}), (w_{2i+1,1}, w_{2i,2}) \mid i \in [k]\}. \end{aligned}$$

As before, the graph G_n consists of four copies of H_k , let's call them $H_k^{(j)}$, $j \in [4]$, where all the vertices are distinct except for $w_{1,2}^{(j)} = w_{1,1}^{(j+1)}$ for $j \in [3]$ and $w_{1,2}^{(4)} = w_{1,1}^{(1)}$. The construction is depicted on Figure 3.

Parts (i) and (iii) of the theorem follow immediately from the construction. We will now prove part (ii). First let us observe that Claim 3 from Theorem 2 holds for the graph G_n from this proof. We will prove (ii) by first showing that it is possible to use DG moves to move between any two subgraphs of G_n satisfying part a) of the claim (and, analogously, the same holds for two subgraphs satisfying part b)), and then we show that there is a subgraph satisfying part a) and a subgraph satisfying part b) such that it is possible to move between them.

We make one more observation: if G' is a subgraph of G_n satisfying part a) of the claim, then the intersection of G' with both $H_k^{(2)}$ and $H_k^{(4)}$ is uniquely defined. This follows from the fact that there




Figure 2: The graph H_k from Theorem 4.




Figure 3: The graph representation of the input instances from Theorem 4.




Figure 4: A possible subgraph of the graph H_k with all-one degree sequence.

is a unique subgraph of H_k such that the vertices $w_{1,1}$ and $w_{1,2}$ have degree 0 while all other vertices have degree 1. (This subgraph contains exactly the edges $(w_{2i+1,1}, w_{2i,2})$ and $(w_{2i+2,1}, w_{2i+1,2})$ for every $i \in [k]$, and the edge $(w_{2,1}, w_{2k+2,2})$.)

Now we construct a subgraph G' satisfying part a) and a subgraph G'' satisfying part b) such that we can move from G' to G'' in a single DG move: let G' contain the edges $(w_{1,1}^{(1)}, w_{1,2}^{(1)})$ and $(w_{1,1}^{(3)}, w_{1,2}^{(3)})$ (notice that by the above argument all the other edges of G' are uniquely defined); similarly, let G'' contain the edges $(w_{1,1}^{(2)}, w_{1,2}^{(2)})$ and $(w_{1,1}^{(4)}, w_{1,2}^{(4)})$. It follows immediately that G' and G'' are connected by a single DG move.

Suppose we have two subgraphs G' and G'' both satisfying part a) of the claim. To show we can move from G' to G'' using DG moves, we first show we can move from G' to the subgraph of H_k with all-one degree sequence that includes the edge $(w_{1,1}, w_{1,2})$. Let us call this graph \tilde{G}' . Suppose G' does not include this edge (or we are done). Therefore, it must connect $w_{1,1}$ to some other vertex, let it be $w_{2i+1,2}$ where $i \in [k]$ (if $w_{1,1}$ is connected to $w_{2k+2,2}$, analogous arguments hold). Then, the part of G' involving vertices $w_{j,\ell}$ for $j \in [2i+1] \cup \{2k+2\}$ and $\ell \in [2]$ is uniquely determined, and the rest can be chosen from exactly 2^{k-i} possibilities (see Figure 4). We can do the following DG moves:

1. For j going from $i+1$ to k do:
 2. If the graph G' contains the edges $(w_{2j,1}, w_{2j,2})$ and $(w_{2j+1,1}, w_{2j+1,2})$, swap them for $(w_{2j,1}, w_{2j+1,2})$ and $(w_{2j+1,1}, w_{2j,2})$.
 3. Swap the edges $(w_{1,1}, w_{2j-1,2})$ and $(w_{2j,1}, w_{2j+1,2})$ for $(w_{2j,1}, w_{2j-1,2})$ and $(w_{1,1}, w_{2j+1,2})$.
 4. Finally, swap the edges $(w_{1,1}, w_{2k+1,2})$ and $(w_{2k+2,1}, w_{1,2})$ for the edges $(w_{1,1}, w_{1,2})$ and $(w_{2k+2,1}, w_{2k+1,2})$.

We end up with the graph \tilde{G}' , as promised.

Last, we show part (iv) of the theorem. As in the proof of Theorem 2, we bound the conductance of the chain: let S be the set of all subgraphs of G_n that satisfy part a) of the claim. Then, $\pi(S) = 1/2$ and

$$\sum_{s_1 \in S, s_2 \in \Sigma_{r,c,b} \setminus S} \pi(s_1) P(s_1, s_2) = \frac{1}{|\Sigma_{r,c,b}|} \frac{1}{n^2(n-1)^2},$$

since there is a unique pair of graphs $s_1 \in S_1$, $s_2 \in S_2$ such that $P(s_1, s_2) > 0$, namely $P(s_1, s_2) = \frac{1}{4(n(n-1)/2)^2}$. To estimate $|\Sigma_{r,c,b}|$, recall that there are exactly 2^{k-i} subgraphs of H_k that satisfy the all-one degree sequence and contain the edge $(w_{1,1}, w_{2i+1,2})$. Therefore, there are $\sum_{i=1}^k 2^{k-i} + 2^k + 1 = 2^{k+1}$ subgraphs of H_k with all-one degree sequence (2^k of those that contain the edge $(w_{1,1}, w_{2k+2,2})$ and 1 that contains the edge $(w_{1,1}, w_{1,2})$). Thus, $|\Sigma_{r,c,b}| = 2(2^{k+1})^2 = 2^{2k+3} = 2^{n/4+3/2}$ and the conductance is bounded by $\frac{1}{n^2(n-1)2^{n/4}}$, an inverse exponential function. This concludes the proof of part (iv) of the theorem. \square

4 Conclusions

The main question that remains open is whether the Diaconis-Gangolli chain could be slowly mixing even in the case without cell bounds. To tackle this question, we might consider several special cases. The families of instances in this paper use cell-bounds of 0 or 1. To get closer to the unbounded case, we might ask whether our results hold in the case when the cell-bounds $b_{i,j}$ are strictly positive. Also, it may be helpful to know whether the Diaconis-Gangolli chain mixes slowly for cell-bounds of the form b and $b + 1$ for larger $b > 0$. Finally, we note that for our instances it was sufficient to consider the variant of the Diaconis-Gangolli chain that always adds 1 to the entries on the chosen diagonal and subtracts 1 from the other diagonal's entries. Once we allow larger cell-bounds, adding and subtracting a larger number could possibly help to speed up the mixing time.

References

- [1] I. Bezáková, N. Bhatnagar, and E. Vigoda. Sampling binary contingency tables with a greedy start, *Random Structures and Algorithms*, **30**: 168-205, 2007.
- [2] M. Cryan and M. Dyer. A polynomial-time algorithm to approximately count contingency tables when the number of rows is constant, *Journal of Computer and System Sciences*, **67**: 291–310, 2003.
- [3] M. Cryan, M. Dyer, L. Goldberg, M. Jerrum and R. Martin. Rapidly mixing Markov chains for sampling contingency tables with a constant number of rows, in *Proc. 43rd IEEE Symposium on Foundations of Computer Science*, 711–720, 2002.
- [4] M. Cryan, M. Dyer and D. Randall. Approximately counting integral flows and cell-bounded contingency tables, in *Proc. 37th ACM Symposium on Theory of Computing*, 413-422, 2005.
- [5] P. Diaconis and B. Efron. Testing for independence in a two-way table: new interpretations of the chi-square statistic, *Annals of Statistics* **13**, 845–913, 1995.
- [6] P. Diaconis and A. Gangolli. Rectangular Arrays with Fixed Margins, *Discrete Probability and Algorithms*, eds. D. Aldous et al., Springer-Verlag, 15-41, 1995.
- [7] M. Dyer. Approximate counting by dynamic programming, in *Proc. 35th ACM Symposium on the Theory of Computing*, 693–699, 2003.
- [8] M. Dyer, R. Kannan and J. Mount. Sampling contingency tables. *Random Structures & Algorithms*, **10**, 487–506, 1997.

- [9] M. Jerrum and A. Sinclair. Approximate counting, uniform generation and rapidly mixing Markov chains, *Information and Computation*, **82**: 93-133, 1989.
- [10] M. Jerrum, A. Sinclair, and E. Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with non-negative entries, *Jour. ACM*, **51**: 671-697, 2004.
- [11] M. Jerrum, L. Valiant and V. Vazirani. Random generation of combinatorial structures from a uniform distribution, *Theoretical Comp. Sci.* **43**, pp. 169–188, 1986.
- [12] R. Kannan, P. Tetali, and S. Vempala. Simple Markov-chain algorithms for generating bipartite graphs and tournaments, *Random Structures and Algorithms*, **14**:293-308, 1999.
- [13] G. Lawler and A. Sokal. Bounds on the L2 spectrum for Markov chains and Markov processes: a generalization of Cheeger’s inequality, *Trans. Amer. Math. Soc.*, **309**: 557-580, 1988.
- [14] B. Morris. Improved bounds for sampling contingency tables, *Random Structures & Algorithms*, **21**: 135–146, 2002.
- [15] J. Mount. Application of convex sampling to optimization and contingency table generation, PhD thesis, Carnegie Mellon University, 1995. (Technical Report CMU-CS-95-152, Department of Computer Science.)
- [16] C. H. Papadimitriou. *Computational Complexity*, Addison-Wesley, 1994.