

Approximately Counting Integral Flows and Cell-Bounded Contingency Tables*

Mary Cryan[†], Martin Dyer[‡] and Dana Randall[§]

Abstract

We consider the problem of approximately counting integral flows in a network. We show that there is an *fpras* based on volume estimation if all capacities are sufficiently large, generalising a result of Dyer, Kannan and Mount (1997). We apply this to approximating the number of contingency tables with prescribed *cell bounds* when the number of rows is constant, but the row sums, column sums and cell bounds may be arbitrary. We provide an *fpras* for this problem via a combination of dynamic programming and volume estimation. This generalises an algorithm of Cryan and Dyer (2002) for standard contingency tables, but the analysis here is considerably more intricate.

1 Introduction

In this paper we consider two related counting problems. First we consider the problem of counting *integral flows in a general capacitated network*. A special case of this problem was considered by Kannan and Vempala [21]. They gave an *fpras* (fully polynomial randomized approximation scheme) to approximately count integral $s-t$ flows in an undirected network when all edge capacities are sufficiently large. Recently Baldoni-Silva, De Loera and Vergne showed that integer-valued flows in a general capacitated network can be represented as lattice points inside a related *flow polytope* [3]. Hence they construct exact counting algorithms using Barvinok's approach [2] to counting lattice points in a fixed-dimensional polytope. Their algorithms run in polynomial-time if the dimension of the *flow polytope* is constant. Some applications are also discussed in [3]. In general, exactly counting lattice points is $\#P$ -Complete [16, 28], and only approximation is possible in polynomial-time. Jerrum, Sinclair and Vigoda [19] gave an *fpras* for the special case of *0-1 flows*, where all capacities are 0 or 1, via a reduction to the problem of evaluating the permanent. By contrast, our first contribution in this paper (in §2) is to show there is an *fpras*, based on sampling and volume estimation for convex bodies [14], whenever the minimum (tight) capacity in the network (as defined in §2 below) is sufficiently large. Interestingly, the proof relies on the properties of the *maximum spanning tree* in the network (using the capacities as weights) to show that the flow polytope is *well-rounded* [17]. We note that establishing this property is far from straightforward for general flow polytopes, whereas it follows directly for the special case considered in [21].

*This work was partly done while the authors were visiting the Isaac Newton Institute for Mathematical Sciences, Cambridge, UK.

[†]School of Informatics, University of Edinburgh, Edinburgh EH9 3JZ, UK. Supported by EPSRC grant GR/S76175/01.

[‡]School of Computing, University of Leeds, Leeds LS2 9JT, UK. Supported by EPSRC grant GR/S76151/01.

[§]College of Computing, Georgia Inst. of Technology, Atlanta GA 30332-0280. Supported in part by NSF grant CCR-0105639.

The result of §2 can be applied directly to counting *cell-bounded contingency tables*, which we now define: Let $[\ell]$ denote the set $\{1, \dots, \ell\}$. Suppose we are given a list of positive integers $r = (r_1, \dots, r_m)$ called the *row sums*, another list of positive integers $c = (c_1, \dots, c_n)$ called the *column sums*, and a *cell bound* b_{ij} , for every $i \in [m], j \in [n]$. Assume that $b_{ij} \leq \min\{r_i, c_j\}$ for all $i \in [m], j \in [n]$, and that $\sum_{i=1}^m r_i = \sum_{j=1}^n c_j$ (the *table sum*). Then the set $\Sigma_{r,c,b}$ of *cell-bounded contingency tables* is the set of all $m \times n$ non-negative integer matrices x which satisfy the row and column sums, and also satisfy $x_{ij} \leq b_{ij}$ for all $i \in [m], j \in [n]$. We will see in §3 that cell-bounded contingency tables are equivalent to integer flows in an appropriately-defined bipartite network. The definition of $\Sigma_{r,c,b}$ is a generalisation of the set $\Sigma_{r,c}$ of *standard contingency tables* (where the cell bounds are $b_{ij} = \min\{r_i, c_j\}$ for all $i \in [m], j \in [n]$).

The problem of sampling standard contingency tables from $\Sigma_{r,c}$ almost uniformly at random has important applications in practical statistics (see Diaconis and Efron [8]). Much work has been carried out on this and on the related problem of approximating the total number of tables $|\Sigma_{r,c}|$ (see, e.g., [5, 6, 7, 9, 10, 13, 16, 18, 24, 25]). The algorithm of Barvinok [2] for counting lattice points in a polytope can be used to count contingency tables *exactly* in polynomial-time when the numbers of rows and columns are both constant. This has been implemented successfully in [7]. Dyer, Kannan and Mount [16] showed that, whenever the row sums satisfy $r_i \in \Omega(n^2m)$ for all $i \in [m]$ and the column sums satisfy $c_j \in \Omega(m^2n)$ for all $j \in [n]$, there is an *fpras* for counting standard contingency tables. This was improved later by Morris [23]. Dyer and Greenhill [13] gave an *fpras* to approximately count contingency tables with two rows by proving a natural Markov chain called the 2×2 heat bath chain is rapidly mixing. Subsequently, Cryan and Dyer [5] gave an *fpras* based on a combination of dynamic programming and volume estimation to count tables when the number of rows m is a constant. Cryan et al. [6] gave an alternative algorithm for the constant m case by proving that the 2×2 heat bath Markov chain on such tables is rapidly mixing. Later, Dyer [12] gave an *fpras* based entirely on dynamic programming, improving substantially on the running time of the algorithms in [5] and [6]. However, the existence of an *fpras* for counting standard contingency tables in the general case remains a notorious open problem.

Counting and sampling cell-bounded contingency tables also has natural applications in statistics, in the case where each cell has a known maximum. According to Diaconis and Gangolli [9], this is a “practical class of problems”. However, there appears to have been little work on this problem except in the case where some cell upper bounds may be zero (so-called *structural zeros*). The Markov chain used in [6, 13], which updates values in a 2×2 submatrix during each step of the simulation, is no longer viable in the cell-bounded case since the state space might not even be connected; for example, a 3×3 table with structural zeros on the diagonal has no allowable moves. Aoki [1] and Rapollo [26] recently considered the design of alternative chains for sampling such tables using the “Markov basis” approach of Diaconis and Sturmfels [11]. The required sampling distribution will be uniform when the conditional volume test of [8] is employed. However, the main focus of these papers is on sampling from the hypergeometric distribution, which is easier than the general case. These Markov chains take small steps, so cannot lead to polynomial time sampling unless all numbers are given in unary. More recently, Bezáková, Bhatnagar and Vigoda [4] considered binary contingency tables, which are cell-bounded contingency tables with all cell bounds equal to 1. Binary contingency tables can be represented as 0-1 flows, so an *fpras* is known to exist [19]. Bezáková et al. gave an improved *fpras* for binary contingency tables, by bounding the mixing time of a Markov chain on the set of binary contingency tables and “near-contingency tables”.

There are general theoretical reasons for studying cell-bounded tables. It is well-known that

for any *self-reducible* relation, the problem of obtaining an *fpras* for approximate counting is equivalent to the problem of finding an *fpaus* (fully polynomial almost uniform sampler) [20]. Standard contingency tables are not known to be self-reducible, which is unusual. However there is a simple self-reducible characterisation of cell-bounded contingency tables. Another interesting fact about cell-bounded contingency tables is that they generalise the concept of perfect matchings in a bipartite graph (where all cell-bounds are 0 or 1). Approximating the number of perfect matchings in a bipartite graph (the *0-1 permanent*) was an important open problem for several years, until Jerrum, Sinclair and Vigoda [19] finally established the existence of an *fpras*.

The second overall contribution of our paper is to put the problem of approximately counting *cell-bounded* contingency tables on precisely the same footing as that of approximately counting *standard* contingency tables. One immediate consequence of the main result of §2 is that, if all cell bounds are “sufficiently large” (in a sense we will define later), then we can use sampling and volume estimation for the flow polytope to obtain an *fpras* to count approximately the number of cell-bounded tables. This result does not depend on the number of rows or columns. Therefore it can be seen as a direct generalisation to the cell-bounded case of the result of Dyer, Kannan and Mount [16] for standard contingency tables.

Following the other thread of results for contingency tables, in §3 we assume that the number of rows is constant but make no assumptions about the size of the cell bounds. We show that we can combine dynamic programming and volume estimation to design an *fpras* in this case. This *fpras* is broadly similar to that of Cryan and Dyer [5], which was also a combination of dynamic programming and volume estimation, but the structure of cell-bounded tables is considerably more intricate than that of standard tables. However, since cell-bounded contingency tables may be viewed as integral flows in a bipartite network, we are able to extend the approach of §2. Our proof relies even more strongly on properties of the maximum spanning tree.

These are the first results that provide provably efficient counting and sampling algorithms for non-trivial classes of cell-bounded contingency tables. Moreover, they demonstrate that these seemingly broader classes of problems might not be harder than counting and sampling standard contingency tables without cell bounds. The question of existence of an *fpras* for the general problem of counting cell-bounded contingency tables remains open, as it does for standard contingency tables. However, since an *fpras* for general cell-bounded tables would include approximating the 0-1 permanent [19] as a special case, it may prove elusive. On the other hand, the results of this paper are some indication that standard contingency tables have no exploitable structure beyond that which exists in the cell-bounded case. Consequently, there is no obvious reason to expect that an *fpras* for standard contingency tables may be found any more easily than one for the general cell-bounded case.

2 Counting integral flows

Suppose that we have a flow network $\mathcal{N} = (\mathbf{V}, \mathbf{A})$ with *capacities* $\mathbf{b}(\mathbf{a}) \in \mathbb{Z}^+ \cup \{\infty\}$ ($\mathbf{a} \in \mathbf{A}$) and *supplies* $\rho(\mathbf{v}) \in \mathbb{Z}$ ($\mathbf{v} \in \mathbf{V}$) such that $\sum_{\mathbf{v} \in \mathbf{V}} \rho(\mathbf{v}) = 0$. Negative supplies are called *demands*. We will write $\mathbf{n} = |\mathbf{V}|$, $\mathbf{m} = |\mathbf{A}|$ and $\mathbf{d} = \mathbf{m} - \mathbf{n} + 1$. A flow \mathbf{x} must satisfy

$$\sum_{\mathbf{w}: \mathbf{a}=(\mathbf{v}, \mathbf{w})} \mathbf{x}(\mathbf{a}) - \sum_{\mathbf{w}: \mathbf{a}=(\mathbf{w}, \mathbf{v})} \mathbf{x}(\mathbf{a}) = \rho(\mathbf{v}) \quad (\mathbf{v} \in \mathbf{V}), \quad (1)$$

$$0 \leq \mathbf{x}(\mathbf{a}) \leq \mathbf{b}(\mathbf{a}) \quad (\mathbf{a} \in \mathbf{A}). \quad (2)$$

The solution set of (1)–(2), the set of *flows*, is a convex polyhedron. We are interested in estimating the number of integral flows. Any capacitated flow problem can be put in the form (1)–(2) without changing the size of the solution set. We allow \mathcal{N} to contain parallel and antiparallel arcs between any two vertices since replacing these with a single edge might alter the number of integer flows satisfying (1)–(2). We may dispose of the case in which the number of solutions is infinite, corresponding to the existence of a directed cycle in the set $\{\mathbf{a} \in \mathbf{A} : \mathbf{b}(\mathbf{a}) = \infty\}$. If no such cycle exists, the system (1)–(2) determines a *flow polytope* \mathcal{P} . The integer solutions are the *lattice points* inside \mathcal{P} , and we use $\mathcal{I}(\mathcal{P})$ to denote the set of all such lattice points.

Our goal in this section is to develop an *fpras* for counting the number of integral flows when the minimum capacity is sufficiently large. That is, given an error tolerance $\epsilon \in (0, \frac{1}{2})$, we will design an algorithm that runs in time polynomial in n, ϵ^{-1} , and $\log(\max_{\mathbf{a}} \mathbf{b}(\mathbf{a}))$ and produces an estimate $|\mathcal{I}'|$ of $|\mathcal{I}(\mathcal{P})|$ that satisfies

$$(1 - \epsilon)|\mathcal{I}(\mathcal{P})| \leq |\mathcal{I}'| \leq (1 + \epsilon)|\mathcal{I}(\mathcal{P})|$$

with high probability. We accomplish this by relating $|\mathcal{I}(\mathcal{P})|$ to the volume of \mathcal{P} .

We assume that \mathcal{N} has *tight capacities*, defined as follows.

Definition 1 *A network \mathcal{N} has tight capacities if for each $\mathbf{a} \in \mathbf{A}$, there exist flows $f_{\mathbf{a}}^{-}, f_{\mathbf{a}}^{+}$ satisfying $f_{\mathbf{a}}^{-}(\mathbf{a}) = 0$ and $f_{\mathbf{a}}^{+}(\mathbf{a}) = \mathbf{b}(\mathbf{a})$ (and $\mathbf{b}(\mathbf{a}) > 0$).*

If the network does not satisfy this condition, we can make a polynomial-time transformation to define a new network with the same number of integral flows. For each arc \mathbf{a} , we find $\max_{\mathbf{x}} \mathbf{x}(\mathbf{a})$, the maximum value of $\mathbf{x}(\mathbf{a})$ over all flows, by solving a minimum cost flow problem. (See, for example, Schrijver [27, §12]). Similarly, for every $\mathbf{a} \in \mathbf{A}$, we find $\min_{\mathbf{x}} \mathbf{x}(\mathbf{a})$ by solving a minimum cost flow problem. We let the new capacities be $(\max_{\mathbf{x}} \mathbf{x}(\mathbf{a}) - \min_{\mathbf{x}} \mathbf{x}(\mathbf{a}))$ ($\mathbf{a} \in \mathbf{A}$), and let the new supplies be

$$\rho(\mathbf{u}) - \sum_{(\mathbf{u}, \mathbf{v}) \in \mathbf{A}} \min_{\mathbf{x}} \mathbf{x}(\mathbf{u}, \mathbf{v}) + \sum_{(\mathbf{w}, \mathbf{u}) \in \mathbf{A}} \min_{\mathbf{x}} \mathbf{x}(\mathbf{w}, \mathbf{u}) \quad (\forall \mathbf{u} \in \mathbf{V}).$$

It is straightforward to check that there is a bijection between flows in the two networks. The modified network can be constructed in polynomial time by solving at most $2m$ minimum cost flow problems. We also may assume that $\mathbf{b}(\mathbf{a}) > 0$ for all $\mathbf{a} \in \mathbf{A}$, since otherwise arc \mathbf{a} can be deleted from \mathbf{A} .

Using the convexity of the flow polytope, we can now define an *internal flow*.

Definition 2 *Given a flow polytope \mathcal{P} with tight capacities, we define an internal flow*

$$\mathbf{f} \stackrel{\text{def}}{=} \frac{1}{2m} \sum_{\mathbf{a} \in \mathbf{A}} (f_{\mathbf{a}}^{-} + f_{\mathbf{a}}^{+}) \quad (3)$$

In general \mathbf{f} will not be an integral flow. We will use *rational flow* for functions like \mathbf{f} whereas, without qualification, flow will mean *integral flow*. Note that $\frac{1}{2m} \mathbf{b}(\mathbf{a}) \leq \mathbf{f}(\mathbf{a}) \leq \mathbf{b}(\mathbf{a}) - \frac{1}{2m} \mathbf{b}(\mathbf{a})$ for all $\mathbf{a} \in \mathbf{A}$, using the tightness of the capacities. For any flow \mathbf{x} , define the *slack* for \mathbf{x} on arc \mathbf{a} as $\min\{\mathbf{x}(\mathbf{a}), \mathbf{b}(\mathbf{a}) - \mathbf{x}(\mathbf{a})\}$. Thus our internal flow \mathbf{f} has slack at least $\mathbf{b}(\mathbf{a})/2m$ on every arc \mathbf{a} .

We now define the concept of a *maximum spanning tree* of a network. We may assume \mathcal{N} is connected, since otherwise the system (1), (2) decomposes and we can consider each component separately and take the product of the number of solutions for each component.

Definition 3 Consider the (connected) undirected multigraph $G = (V, E)$ underlying \mathcal{N} . We will abuse notation and refer to an arc $\mathbf{a} \in A$ as an edge $\mathbf{a} \in E$, forgetting its direction. Let E have edge weights $\mathbf{b}(\mathbf{a})$. Then a maximum spanning tree for \mathcal{N} is any maximum weight spanning tree T in G .

It follows from standard network flow theory [27, §13] that we can eliminate the variables $x(\mathbf{a})$ ($\mathbf{a} \in T$) from the system of equations (1)–(2) to give a system of $2m$ inequalities in $d = (m - n + 1)$ bounded variables $x(\mathbf{a})$ ($\mathbf{a} \in A' = A \setminus T$).

We now show, using the spanning tree T , that there is an ellipsoid, and also a ball, centred at f , lying entirely inside \mathcal{P} . The approach of using a spanning tree of a network to define a full-dimensional representation of the flow polytope has been used before [21], but it is the idea of using the *maximum-weight* spanning tree which drives our result.

Theorem 1 Let $\mathcal{N} = (V, A)$ be a connected network with tight capacities $\mathbf{b}(\cdot)$. Let T be any maximum spanning tree for \mathcal{N} . Let $\mathbf{b}_{\min} = \min_{\mathbf{a} \in A \setminus T} \mathbf{b}(\mathbf{a})$, and let $\delta = \mathbf{b}_{\min}/2m\sqrt{d}$ (where $d = m - n + 1$). Then the flow polytope \mathcal{P} contains the ball $\mathcal{B}(f, \delta)$.

Proof: For any $\mathbf{a} \in A' = A \setminus T$, consider the rational flow $g_{\mathbf{a}}^+$ defined by

$$\begin{aligned} g_{\mathbf{a}}^+(\mathbf{a}) &= f(\mathbf{a}) + \frac{1}{2m}\mathbf{b}(\mathbf{a}) & \text{and} \\ g_{\mathbf{a}}^+(\mathbf{a}') &= f(\mathbf{a}') & (\mathbf{a}' \neq \mathbf{a}, \mathbf{a}' \in A'). \end{aligned}$$

Clearly $0 \leq g^+(\mathbf{a}') \leq \mathbf{b}(\mathbf{a}')$ for all $\mathbf{a}' \in A'$. To define a rational flow, we must complete this with feasible values of $g_{\mathbf{a}}^+(\mathbf{a}')$ ($\mathbf{a}' \in T$). First, we follow the unique circuit $C_{\mathbf{a}} \subseteq T \cup \{\mathbf{a}\}$ in the direction of \mathbf{a} . For each edge $\mathbf{a}' \neq \mathbf{a}$ traversed, if the direction of \mathbf{a}' is the same as that of \mathbf{a} , let $g_{\mathbf{a}}^+(\mathbf{a}') = f(\mathbf{a}') + \mathbf{b}(\mathbf{a})/2m$; alternatively, if \mathbf{a}' has the opposite direction to \mathbf{a} , $g_{\mathbf{a}}^+(\mathbf{a}') = f(\mathbf{a}') - \mathbf{b}(\mathbf{a})/2m$. We set $g_{\mathbf{a}}^+(\mathbf{a}') = f(\mathbf{a}')$ for all $\mathbf{a}' \in T \setminus C_{\mathbf{a}}$. This ensures that $g_{\mathbf{a}}^+$ still satisfies all the supplies $\rho(\cdot)$. It follows that for all $\mathbf{a}' \in C_{\mathbf{a}}$ we have $\mathbf{b}(\mathbf{a}) \leq \mathbf{b}(\mathbf{a}')$ and

$$g_{\mathbf{a}}^+(\mathbf{a}') \in f(\mathbf{a}') \pm \frac{1}{2m}\mathbf{b}(\mathbf{a}) \in f(\mathbf{a}') \pm \frac{1}{2m}\mathbf{b}(\mathbf{a}') \in [0, \mathbf{b}(\mathbf{a}')]]$$

since T is a *maximum* spanning tree. Similarly we can find $g_{\mathbf{a}}^-$ such that $g_{\mathbf{a}}^-(\mathbf{a}) = f(\mathbf{a}) - \frac{1}{2m}\mathbf{b}(\mathbf{a})$ and $g_{\mathbf{a}}^-(\mathbf{a}') = f(\mathbf{a}')$ ($\mathbf{a}' \neq \mathbf{a}, \mathbf{a}' \in A'$). Now consider \mathcal{P} as a polytope in $\mathbb{R}^d = \mathbb{R}^{m-n+1}$, determined by $x(\mathbf{a})$ ($\mathbf{a} \in A'$). From the properties of f , \mathcal{P} is contained in the hyper-rectangle

$$-(1 - \frac{1}{2m})\mathbf{b}(\mathbf{a}) \leq x(\mathbf{a}) - f(\mathbf{a}) \leq (1 - \frac{1}{2m})\mathbf{b}(\mathbf{a}) \quad (\mathbf{a} \in A').$$

Thus \mathcal{P} is contained in the ellipsoid

$$\mathcal{E}_{\text{ext}} = \left\{ x : \sum_{\mathbf{a} \in A'} \left(\frac{x(\mathbf{a}) - f(\mathbf{a})}{\mathbf{b}(\mathbf{a})} \right)^2 \leq d \left(1 - \frac{1}{2m}\right)^2 \right\}.$$

We have shown that \mathcal{P} contains $\mathcal{H} = \text{conv}\{g_{\mathbf{a}}^+, g_{\mathbf{a}}^- : \mathbf{a} \in A'\}$. This is an axis-scaled l_1 -ball, so

$$\mathcal{H} = \left\{ x : \sum_{\mathbf{a} \in A'} \left| \frac{x(\mathbf{a}) - f(\mathbf{a})}{\mathbf{b}(\mathbf{a})} \right| \leq \frac{1}{2m} \right\}.$$

Now, using Cauchy-Schwarz, \mathcal{H} contains the ellipsoid

$$\mathcal{E}_{\text{int}} = \left\{ x : \sum_{\mathbf{a} \in A'} \left(\frac{x(\mathbf{a}) - f(\mathbf{a})}{\mathbf{b}(\mathbf{a})} \right)^2 \leq \frac{1}{4m^2d} \right\}.$$

which is \mathcal{E}_{ext} scaled by $d(2m-1)$. Let $\mathbf{b}_{\min} = \min_{\mathbf{a} \in A'} \mathbf{b}(\mathbf{a})$ and $\delta = \mathbf{b}_{\min}/2m\sqrt{d}$. Then \mathcal{E}_{int} contains the ball $\mathcal{B}(f, \delta)$. \square

We now show $|\mathcal{I}(\mathcal{P})|$ is close to $\text{vol}(\mathcal{P})$ when \mathbf{b}_{\min} is sufficiently large.

Theorem 2 *Let $\mathcal{N} = (V, A)$ be a connected network with tight capacities $\mathbf{b}(\cdot)$. Let T be a maximum spanning tree for \mathcal{N} . Let \mathcal{P} be the flow polytope in \mathbb{R}^d , with axes indexed by $A \setminus T$, where $d = m - n + 1$. If $\mathbf{b}_{\min} > 2md$ and $\epsilon \geq 2md^2/(\mathbf{b}_{\min} - 2md)$, then*

$$e^{-\epsilon} \text{vol}(\mathcal{P}) \leq |\mathcal{I}(\mathcal{P})| \leq e^{\epsilon} \text{vol}(\mathcal{P}).$$

Proof: Our proof is similar in principle to that of Theorem 3 in [5], though the polytope we consider here is more constrained. We use a full-dimensional representation of \mathcal{P} in \mathbb{R}^d , relative to the internal flow f defined in (3). Thus, we translate $z \in \mathbb{R}^d$ to z' so that $z'(\mathbf{a}') = z(\mathbf{a}') - f(\mathbf{a})$, for $\mathbf{a}' \in A' = A \setminus T$. Then, for each lattice point $z' = z - f \in \mathcal{I}(\mathcal{P})$, we associate with z' the hypercube $H(z)$ in d -dimensions, where $\mathbf{y} \in H(z)$ iff $z'(\mathbf{a}) \leq \mathbf{y}(\mathbf{a}) < z'(\mathbf{a}) + 1$ for all $\mathbf{a} \in A'$. Let $\mathcal{C} = \cup_{z \in \mathcal{I}(\mathcal{P})} H(z)$. Clearly $|\mathcal{I}(\mathcal{P})| = \text{vol}(\mathcal{C})$. We will refer to the *dilation* $\alpha\mathcal{Q}$ of a d -dimensional convex polytope \mathcal{Q} as $\{\alpha x : x \in \mathcal{Q}\}$. It is well-known that this has volume $\text{vol}(\alpha\mathcal{Q}) = \alpha^d \text{vol}(\mathcal{Q})$. We now prove the theorem in two parts.

We first show that $\mathcal{C} \subseteq (1 + \epsilon/d)\mathcal{P}$. For every $z \in \mathcal{I}(\mathcal{P})$ and every $\mathbf{y} \in H(z)$, $\text{dist}(\mathbf{y}, z) \leq \sqrt{d}$, where $\text{dist}(\cdot, \cdot)$ denotes Euclidean distance. Suppose $\mathbf{y} \in H(z)$ for some $z \in \mathcal{P}$ but $\mathbf{y} \notin \mathcal{P}$. Clearly $\text{dist}(\mathbf{y}, \mathcal{P}) \leq \sqrt{d}$. Also, we know that there is a ball of radius δ with centre f lying inside \mathcal{P} . The dilation $(1 + \epsilon/d)\mathcal{P}$ will therefore contain all points with distance at most

$$\frac{\epsilon\delta}{d} \geq \frac{\mathbf{b}_{\min}\sqrt{d}}{(\mathbf{b}_{\min} - 2md)} > \sqrt{d}$$

from \mathcal{P} . So $H(z) \subseteq (1 + \epsilon/d)\mathcal{P}$ for every $z \in \mathcal{I}(\mathcal{P})$. Hence $\mathcal{C} \subseteq (1 + \epsilon/d)\mathcal{P}$.

Now we will show that the dilation $(1 + \epsilon/d)^{-1}\mathcal{P} \subseteq \mathcal{C}$. Let $\mathbf{y} \in (1 + \epsilon/d)^{-1}\mathcal{P}$ and let $z \in \mathbb{Z}^d$ be the lattice point with $\mathbf{y} \in H(z)$. We will show that $z \in \mathcal{I}(\mathcal{P})$. Clearly $\text{dist}(\mathbf{y}, z) \leq \sqrt{d}$. Since there is a ball of radius δ inside \mathcal{P} , there is a ball of radius $(1 + \epsilon/d)^{-1}\delta$ inside $(1 + \epsilon/d)^{-1}\mathcal{P}$. Dilating $(1 + \epsilon/d)^{-1}\mathcal{P}$ by $(1 + \epsilon/d)$ gives the original polytope \mathcal{P} . Thus the dilation \mathcal{P} of $(1 + \epsilon/d)^{-1}\mathcal{P}$ contains every point with distance from $(1 + \epsilon/d)^{-1}\mathcal{P}$ at most

$$\frac{\epsilon}{d} \cdot \frac{\delta}{1 + \epsilon/d} \geq \sqrt{d}.$$

Hence $(1 + \epsilon/d)^{-1}\mathcal{P} \subseteq \mathcal{C}$.

Observing that $(1 + \epsilon/d)^d \leq e^{\epsilon}$ completes the proof. \square

Corollary 3 *Let $\mathcal{N} = (V, A)$ be a connected network with tight capacities and let \mathcal{P} be the corresponding flow polytope in \mathbb{R}^d . If $\mathbf{b}_{\min} \in \Omega(md^2/\log m)$ then there is an *fpras* for the number of integral flows $|\mathcal{I}(\mathcal{P})|$.*

Proof: To use volume estimation to design an *fpras*, Theorem 2 is useful only if e^{ϵ} is polynomially bounded in m and d . Since $m \geq d$, this is guaranteed when $\mathbf{b}_{\min} \in \Omega(md^2/\log m)$, with the approximation error depending on ϵ . The requirement $\mathbf{b}_{\min} > 2md$ then clearly becomes void. We now show how to find an improved approximation ratio, by eliminating the apparent dependence of the approximation error on ϵ in Theorem 2. Let $\eta > 0$ be the target relative error. First we determine $\text{vol}(\mathcal{P})$ to relative error $\eta/3$, with high probability,

using an *fpras* for volume estimation (see, e.g., [14, 22]). The original volume algorithm is due to Dyer, Frieze and Kannan [14], but many improvements now exist, the most recent being that of Lovász and Vempala [22]. Similarly, we can sample a point x from \mathcal{P} using the *fpaus* of [14], or one of its successors. Suppose we sample $x \in \mathcal{P}$ almost uniformly in this way. Recall that $\mathcal{C} = \bigcup\{H(z) : z \in \mathcal{I}(\mathcal{P})\}$, where $H(z)$ is the multidimensional unit hypercube with least point z . Let $\mathcal{P}' = (1 + \epsilon/d)\mathcal{P}$, so $x' = (1 + \epsilon/d)x$ is a uniform sample from \mathcal{P}' . Note that $\text{vol}(\mathcal{P}') = (1 + \epsilon/d)^d \text{vol}(\mathcal{P})$. Now $|\mathcal{I}(\mathcal{P})| = \text{vol}(\mathcal{C})$, and $\mathcal{C} \subseteq \mathcal{P}'$. Let $p = \text{vol}(\mathcal{C})/\text{vol}(\mathcal{P}')$. We have $p \geq e^{-2\epsilon}$ from Theorem 2, so by assumption it is bounded below by an inverse polynomial. We can recognise when $x' \in \mathcal{C}$ by rounding down to find the unique $y \in \mathbb{Z}^d$ such that $x' \in H(y)$ and checking whether $y \in \mathcal{P}$. Thus, with high probability, we can obtain an approximation \hat{p} for p with relative error at most $\eta/3$ by taking $O(e^{2\epsilon}/\eta^2)$ samples. We can estimate $\text{vol}(\mathcal{C})$ as $\hat{p}(1 + \epsilon/d)^d \text{vol}(\mathcal{P})$. Now the relative error is at most η , so this procedure gives an *fpras* (and an *fpaus*) for any class of flow problems satisfying $b_{\min} \in \Omega(md^2/\log m)$. \square

Kannan and Vempala [21] have given an improvement, using randomized rounding, on the above method for approximating the number of lattice points in a polytope, which was first used in [15]. They consider a polytope \mathcal{Q} of dimension d with k facets and show that, if \mathcal{Q} contains a ball of radius $\Omega(d\sqrt{\log k})$, then there is an *fpras* for $\mathcal{I}(\mathcal{Q})$. There are at most $2m$ facets of the flow polytope \mathcal{P} , so we can employ Kannan and Vempala's result if we have a ball of radius $\Omega(d\sqrt{\log 2m})$ inside \mathcal{P} . By Theorem 1, this gives an *fpras* for the number of integral flows if $b_{\min} = \Omega(d^{3/2}m\sqrt{\log m})$, an improvement on Corollary 3.

3 Contingency tables with cell bounds

We now consider cell-bounded contingency tables with m rows and n columns, row sums $r_i > 0$ ($i \in [m]$) and column sums $c_j > 0$ ($j \in [n]$). The upper bound on cell (i, j) will be denoted by b_{ij} , and the lower bound is zero.

We employ a correspondence between cell-bounded contingency tables and integer flows in a bipartite network. For given values of r , c and b , consider a network \mathcal{N} defined by $V = [m] \uplus [n]$ and $A = [m] \times [n]$. Then $n = m+n$, $m \leq mn$. We will also define $d = m-n+1 \leq (m-1)(n-1)$. The supplies ρ are r_i ($i \in [m]$) and $-c_j$ ($j \in [n]$). The capacities are $b(i, j) = b_{ij}$ for $(i, j) \in A$. It is easy to see that there is a bijection between $\Sigma_{r,c,b}$ and the set of integer flows in \mathcal{N} .

We are interested in developing an *fpras* for counting cell-bounded contingency tables for certain classes of inputs. That is, given an error tolerance $\epsilon \in (0, \frac{1}{2})$, the goal is to design an algorithm which runs in time polynomial in n, ϵ^{-1} and $\max_{ij} \log b_{ij}$ and produces an estimate $|\Sigma'|$ of $|\Sigma_{r,c,b}|$ that satisfies

$$(1 - \epsilon)|\Sigma_{r,c,b}| \leq |\Sigma'| \leq (1 + \epsilon)|\Sigma_{r,c,b}|,$$

with high probability.

The correspondence between flows and cell-bounded tables gives us the following theorem as an immediate consequence of Corollary 3. This can be viewed as a generalisation of a result of Dyer, Kannan and Mount [16], with the additional assumption that all cell bounds, as well as row and column sums, are sufficiently large.¹ Note that, since cell-bounded contingency tables

¹When Corollary 3 is specialised to standard contingency tables, the lower bounds on row and column sums that we require are larger than in [16]. Combining our Theorem 1 with [21], we obtain $\Omega((mn)^{5/2}\sqrt{\log(mn)})$, whereas [16] has $\Omega(mn^2)$ for row totals and $\Omega(m^2n)$ for column totals. It is possible that these results could be improved by a more tailored argument or by applying the result of Morris [23], but we will not explore this further in this paper.

are self-reducible, the existence of an *fpaus* will then follow from the general results of [20].

Theorem 4 *Let $\Sigma_{r,c,b}$ be a set of cell-bounded contingency tables with row sums $r_i > 0$, ($i \in [m]$), column sums $c_j > 0$, ($j \in [n]$), and (tight) cell bounds $b_{ij} > 0$, ($i \in [m], j \in [n]$). There is an fpras for $|\Sigma_{r,c,b}|$ if $\min_{i,j} b_{ij} \in \Omega((mn)^3 / \log(mn))$.*

We note that the result of [21] referred to above would allow us to improve this, to give an fpras under the weaker condition $\min_{i,j} b_{ij} \in \Omega((mn)^{5/2} \sqrt{\log(mn)})$.

In the remainder of the paper we consider a second class of cell-bounded contingency tables where the number of rows m is a *constant*. We show the following theorem, which requires no assumptions about the size of the row and column sums. This theorem also implies the existence of an *fpaus*.

Theorem 5 *Let $\Sigma_{r,c,b}$ be a set of cell-bounded contingency tables with row sums $r_i > 0$, ($i \in [m]$), column sums $c_j > 0$, ($j \in [n]$), and cell bounds $b_{ij} > 0$, ($i \in [m], j \in [n]$), and let m be a constant. Then there is an fpras for $|\Sigma_{r,c,b}|$.*

The algorithm used to establish Theorem 5 is based on combining the approach of §2 with partial dynamic programming. A critical parameter of the algorithm will be

$$\beta = 6 + 2 \log_n(32m^6 \epsilon^{-1}).$$

The parameters p, q and r will also be important later, with $p = q + \beta/2$, $q \geq \beta$, $r \geq q + \beta$.

We define $r(\ell)$ to be the list of row sums (r_1, \dots, r_ℓ) for any $\ell \leq m$ and $c(k)$ to be the list of column sums (c_1, \dots, c_k) for any $k \leq n$. We define $r(\bar{\ell})$ to be $(r_{\ell+1}, \dots, r_m)$ for any $\ell < m$ and $c(\bar{k}) = (c_{k+1}, \dots, c_n)$ for any $k < n$.

As in §2, we will assume without loss that at the beginning of the dynamic programming phase we have *tight cell bounds*, i.e. for every $i \in [m], j \in [n]$, there is some table $\mathbf{x} \in \Sigma_{r,c,b}$ such that $x_{ij} = b_{ij}$ and some table $\mathbf{y} \in \Sigma_{r,c,b}$ such that $y_{ij} = 0$. If this is not true initially, we can transform the input into this format by solving $2mn$ minimum cost flow problems in polynomial-time (see, for example, Schrijver [27, §12]). For each $i \in [m], j \in [n]$, we evaluate $\max_{\mathbf{x}} x_{ij}$ and $\min_{\mathbf{x}} x_{ij}$. Then we modify the values of b_{ij}, r_i and c_j in the manner described in §2. It is clear that $\Sigma_{r,c,b}$ has tight cell bounds if and only if \mathcal{N} has tight capacities. If $b_{ij} = 0$, we delete (i, j) from \mathcal{A} . Later, it will become necessary to work with networks and cell-bounded tables which do not have tight capacities or bounds. Then we will use ℓ_{ij} and b'_{ij} to represent the tight lower and upper bounds for cell (i, j) .

The maximum spanning tree T of the network \mathcal{N} defined in the last section again plays an important role in our algorithm. We will assume that \mathcal{N} is connected, as in §2. In that case, the maximum spanning tree (or, if \mathcal{N} is disconnected, the maximum spanning forest) can be constructed in $O(mn \log n)$ time by a standard algorithm [27, §50]. It is straightforward to show that the maximum spanning tree T is a set of $n + m - 1$ cells. Since each column must contain an element of T , therefore by the properties of maximum spanning trees, the tree must have the following structure.

- (i) A set K of $k = |K| \leq (m - 1)$ columns containing at least two cells per column, giving $(k + m - 1) \leq 2(m - 1)$ cells in total. We denote this set of cells by B . For each $j \in K$, there is some i^* with $b_{i^*j} = \max_i b_{ij}$ such that $(i^*, j) \in \mathsf{B}$.
- (ii) In each of the remaining $(n - k)$ columns $j \in [m] \setminus K$, there is a single cell $(i^*, j) \in \mathsf{T}$ such that $b_{i^*j} = \max_i b_{ij}$, giving $(n - k)$ cells in total. Denote this set of cells by D .

Clearly $T = B \uplus D$. In the special case of *standard* contingency tables [5] T has a very special structure: B contains all cells in the column with $\max_j c_j$ and D contains all remaining cells in the row with $\max_i r_i$. In cell-bounded tables the structure of T is less predictable but is crucial for our algorithm. We will use T to partition the columns of the table into *small* and *large columns*, and the rows of the table into *small* and *large rows*. The key to this partitioning is a “jump” property for the cells of B , which results from the following observation.

Lemma 6 *For every $i \in [m]$, there exists a $j' \in [n]$ such that $(i, j') \in B$ and $b_{ij'} > r_i/mn$.*

Proof: There exists $(i, j^*) \in T$ such that $b_{ij^*} = \max_j b_{ij} \geq r_i/n$, since T is a maximum spanning tree. If $(i, j^*) \in B$, we are finished. Otherwise, suppose (i, j^*) is in D . Column j^* must contain a cell (i', j^*) with $i' \neq i$ such that $b_{i'j^*} > b_{ij^*}/m$. If not, we would have

$$\sum_{i' \neq i} b_{i'j^*} \leq (m-1)b_{ij^*}/m \leq (m-1)c_{j^*}/m.$$

Then, for every table $x \in \Sigma_{r,c,b}$, we would require $x_{ij^*} \geq c_{j^*} - \sum_{i' \neq i} b_{i'j^*} \geq c_{j^*}/m$. This contradicts the tightness of the bounds for cell (i, j^*) , since it implies a positive lower bound. Hence there must be at least one $i' \neq i$ such that $b_{i'j^*} \geq b_{ij^*}/m$, and our claim holds. Also, because $(i, j^*) \in D$, we have $(i', j^*) \notin T$. But now there is a circuit Γ in $B \cup \{(i, j^*), (i', j^*)\}$ with $(i, j^*) \in \Gamma \cap T$, $(i', j^*) \in \Gamma \setminus T$. Clearly there is some cell $(i, j') \in \Gamma \cap B$ and we must have $b_{ij'} \geq b_{i'j^*}$, since T is a maximum spanning tree. Thus $b_{ij'} \geq b_{i'j^*} > b_{ij^*}/m \geq r_i/mn$, as required. \square

We are now ready to prove the “jump” property.

Lemma 7 *Suppose that $n \geq m$ and we are given $r = (r_1, \dots, r_m), c = (c_1, \dots, c_n)$, and tight cell bounds $\{b_{ij} : i \in [m], j \in [n]\}$. Then either it is the case that all $b_{ij} < n^{2m\beta}$, or there is a “jump” $\mathcal{J} = [n^q, n^r]$, with $q \geq \beta$, $q + \beta \leq r \leq 2m\beta$, such that $b_{ij} \notin \mathcal{J}$ for all $(i, j) \in B$.*

Proof: Let (i_0, j_0) be such that $b_{i_0j_0}$ is a largest cell bound. Assume $b_{i_0j_0} > n^{2m\beta}$, from Lemma 6 there exists $(i_0, j_1) \in B$ with $b_{i_0j_1} \geq n^{2m\beta-2}$. There are at most $2(m-1)$ cells of B . We may assume that some one cell has bound at most $n^{2\beta}$, otherwise we can set $\mathcal{J} = [n^\beta, n^{2\beta}]$. So assume there exists $(i_2, j_2) \in B$ such that $b_{i_2j_2} \leq n^{2\beta}$. There are at most $2m-3$ other cells in B . If there is no jump \mathcal{J} , then $\max_{(i,j) \in B} b_{ij} \leq n^{2\beta} n^{\beta(2m-3)} = n^{\beta(2m-1)} < n^{2m\beta-2} < b_{i_0j_1}$, a contradiction. \square

If $b_{ij} < n^{2m\beta}$ for every $i \in [m], j \in [n]$, then we determine $|\Sigma_{r,c,b}|$ exactly by dynamic programming (see §3.1). We now assume the jump \mathcal{J} exists. For any row $i \in [m]$, let $b_{ij'} = \max_{(i,j) \in B} b_{ij}$. By Lemma 6, if $b_{ij'} < n^q$, then $r_i < mn^{q+1}$. Alternatively, if $b_{ij'} > n^r$, then $r_i > n^r$. Thus all row totals r_i satisfy either $r_i < mn^{q+1}$ or $r_i > n^r$. Therefore the jump in the cells in B guaranteed by Lemma 7 induces a corresponding, though smaller, jump in the row totals.

Definition 4 *For any cell with $b_{ij} < n^q$, we call (i, j) a small cell. The set of all small cells is denoted by S . The remaining cells $A \setminus S$ are called large cells. We call any row i a small row if every cell $(i, j) \in B$ is small. We assume without loss that $[\sigma]$ is the set of small rows, for some $\sigma \in [m]$. Thus $r_i < mn^{q+1}$ if $i \in [\sigma]$. Rows $[m] \setminus [\sigma]$, called large rows, satisfy $r_i > n^r$.*

In §3.2 we show that the flow in small cells and rows can be set arbitrarily, without greatly influencing the number of flows in the residual table, provided the small row totals are satisfied and the column totals are sufficiently large. We show in §3.1 how the number of solutions for

the small rows and small cells can be determined by dynamic programming. However, to ensure that we deal with tables in which all column totals are sufficiently large, we first partition the columns of the table into small and large columns.

Definition 5 *If $b_{ij} < n^{2m\beta}$ for all $i \in [m], j \in [n]$, we define the set of small columns to be $[n]$. Otherwise q is defined and $p = q + \beta/2$. Then the set of small columns is the set of all columns j satisfying $c_j < n^p$. We assume without loss that these are columns $[\nu]$ for some $\nu \in [n]$.*

Let $\mathbf{R} = \{(i, j) \in \mathbf{A} : \sigma < i \leq m, \nu < j \leq n\}$. We refer \mathbf{R} as the *residual table*. In §3.1 we will use dynamic programming to decompose the problem into smaller subproblems on \mathbf{R} .

First we eliminate the small columns of the table in polynomial time, to express the value of $|\Sigma_{r,c,b}|$ as the weighted sum of a polynomial number of cell-bounded contingency problems, each of size m by $n - \nu$, on the large columns. Let S_ν be the set of all feasible partial row sums for the small columns. Thus S_ν is the set of all ordered partitions $t = (t_1, \dots, t_m)$ of $\sum_{j=1}^\nu c_j$ into m parts such that $\Sigma_{t,c(\nu),b(\nu)} \neq \emptyset$. For every $t \in S_\nu$, we will determine

$$w_t =_{\text{def}} |\Sigma_{t,c(\nu),b}|.$$

Next we consider the subproblems on the large columns. Let s denote $(r - t)$ throughout. For each $t \in S_\nu$, we have row sums s_i ($i \in [m]$) for the table on the large columns. We perform dynamic programming on the small rows ($r_i \leq mn^{q+1}$) to count exactly the total number of assignments to all small rows, given that the partial row sum over the small columns is t . Let S' be the set of small cells in \mathbf{R} , i.e. $S' = \{(i, j) \in \mathbf{R} : b_{ij} \leq n^q\}$. For the remainder of this section we are working with the large columns, and we abbreviate $\Sigma_{s,c(\bar{\nu}),b}$ to $\Sigma_{s,c,b}$.

Definition 6 *Let $\Sigma_{s(\sigma),*,b}$ be the set of tables on the small rows which have row sums s_i for $i \in [\sigma]$, arbitrary column sums, and satisfy the cell bounds for all $i \in [\sigma], j \in [n] \setminus [\nu]$. Let S' denote the set of small cells in \mathbf{R} , and let $\Sigma_{S'}$ be the set of assignments to all cells $(i, j) \in S'$ which satisfy the cell bounds.*

The second step in §3.1 will be to calculate, for every $t \in S_\nu$, the term

$$W_t =_{\text{def}} |\Sigma_{s(\sigma),*,b}| \times |\Sigma_{S'}|.$$

Definition 7 *A partial assignment \mathbf{x} to all small cells and all cells in small columns and rows will be called good if it satisfies the cell bounds, each small column $j \in [\nu]$ has sum c_j , and each small row $i \in [\sigma]$ has sum r_i . Let \mathcal{G} be the set of all good assignments.*

For any $\mathbf{x} \in \mathcal{G}$, let $\mathcal{N}^{\mathbf{x}}$ denote the residual network obtained by fixing the values of the small columns, small rows and small cells to their values in \mathbf{x} . Let $\mathcal{P}^{\mathbf{x}}$ be the flow polytope for $\mathcal{N}^{\mathbf{x}}$, and let $\mathcal{I}(\mathcal{P}^{\mathbf{x}})$ be the set of lattice points in $\mathcal{P}^{\mathbf{x}}$. Clearly

$$|\Sigma_{r,c,b}| = \sum_{\mathbf{x} \in \mathcal{G}} |\mathcal{I}(\mathcal{P}^{\mathbf{x}})|. \quad (4)$$

We approximate $|\Sigma_{r,c,b}|$ in the following way. We choose any fixed $\mathbf{y} \in \mathcal{G}$. The residual network $\mathcal{N}^{\mathbf{y}}$ is not necessarily connected. However, we make two useful observations:

- (a) For every good assignment $\mathbf{x} \in \mathcal{G}$, $\mathcal{N}^{\mathbf{x}}$ has the same component structure.
- (b) Every component contains at least two rows, so there are at most $m/2$ components.

Facts (a) and (b) follow indirectly from Theorem 8 in §3.1. Fact (a) follows since Theorem 8 implies that \mathcal{N}^x has the component structure of $\mathbb{R} \setminus \mathbb{S}$ for all $x \in \mathcal{G}$. Fact (b) follows, since Theorem 8 implies \mathcal{P}^x is full-dimensional for every component, but a one-row table has a unique rational flow. For each component C of \mathcal{N}^y , we consider the flow polytope \mathcal{P}_C^y of \mathcal{N}_C^y . We use volume estimation [14, 22] to estimate $\text{vol}(\mathcal{P}_C^y)$ within relative error $\epsilon/2m$ in time polynomial in n , ϵ^{-1} and $\log(\max_{ij} b_{ij})$. Denote this estimate by $\widehat{\text{vol}}(\mathcal{P}_C^y)$. Next we define

$$\widehat{\text{vol}}(\mathcal{P}^y) =_{\text{def}} \prod_C \widehat{\text{vol}}(\mathcal{P}_C^y).$$

Finally we estimate $|\mathcal{I}(\mathcal{P}^x)|$ by $\widehat{\text{vol}}(\mathcal{P}^y)$, for every assignment $x \in \mathcal{G}$. Note $\sum_{x \in \mathcal{G}} 1 = \sum_{t \in S_\nu} w_t W_t$. Therefore our estimate $|\Sigma'|$ of $|\Sigma_{r,c,b}|$ is

$$\begin{aligned} |\Sigma'| &=_{\text{def}} \sum_{x \in \mathcal{G}} \widehat{\text{vol}}(\mathcal{P}^y) = \widehat{\text{vol}}(\mathcal{P}^y) \sum_{x \in \mathcal{G}} 1 \\ &= \widehat{\text{vol}}(\mathcal{P}^y) \sum_{t \in S_\nu} w_t W_t. \end{aligned} \quad (5)$$

The w_t will be computed (in polynomial time) in the first dynamic programming phase and the W_t in the second dynamic programming phase (see §3.1). The product $\widehat{\text{vol}}(\mathcal{P}^y)$ is computed by at most $m/2$ calls to a volume estimation algorithm.

To prove that our algorithm is an *fpras*, we must show that, with high probability,

$$(1 - \epsilon)|\Sigma_{r,c,b}| \leq |\Sigma'| \leq (1 + \epsilon)|\Sigma_{r,c,b}|.$$

In Theorem 9 we prove the following, where $x, y \in \mathcal{G}$ and C is any component of \mathcal{N}^y ,

- (i) $(1 - \epsilon/2m)\text{vol}(\mathcal{P}_C^x) \leq |\mathcal{I}(\mathcal{P}_C^x)| \leq (1 + \epsilon/2m)\text{vol}(\mathcal{P}_C^x)$.
- (ii) $(1 - \epsilon/2m)\text{vol}(\mathcal{P}_C^y) \leq \text{vol}(\mathcal{P}_C^x) \leq (1 + \epsilon/2m)\text{vol}(\mathcal{P}_C^y)$.

The proof of Theorem 9 depends strongly on the jump \mathcal{J} of Lemma 7, and a careful choice of values for p , q and r . Combining (i) and (ii), we find that for any $x \in \mathcal{G}$ and component C ,

$$(1 - \epsilon/2m)^2 \text{vol}(\mathcal{P}_C^y) \leq |\mathcal{I}(\mathcal{P}_C^x)| \leq (1 + \epsilon/2m)^2 \text{vol}(\mathcal{P}_C^y).$$

By construction we know that, with high probability, the estimate $\widehat{\text{vol}}(\mathcal{P}_C^y)$ lies within $(1 \pm \epsilon/2m)\text{vol}(\mathcal{P}_C^y)$ for every C . Therefore, with high probability, for every $x \in \mathcal{G}$,

$$(1 - \epsilon/2m)^3 \widehat{\text{vol}}(\mathcal{P}_C^y) \leq |\mathcal{I}(\mathcal{P}_C^x)| \leq (1 + \epsilon/2m)^3 \widehat{\text{vol}}(\mathcal{P}_C^y).$$

For all $x \in \mathcal{G}$ we have

$$|\mathcal{I}(\mathcal{P}^x)| = \prod_C |\mathcal{I}(\mathcal{P}_C^x)|.$$

There are at most $m/2$ components so, with high probability,

$$(1 - \epsilon/2m)^{3m/2} \prod_C \widehat{\text{vol}}(\mathcal{P}_C^y) \leq |\mathcal{I}(\mathcal{P}^x)| \leq (1 + \epsilon/2m)^{3m/2} \prod_C \widehat{\text{vol}}(\mathcal{P}_C^y).$$

We have $\epsilon \in (0, \frac{1}{2})$, and $m \geq 2$, $(1 - \epsilon/2m)^{3m/2} \geq 1 - \epsilon$ and $(1 + \epsilon/2m)^{3m/2} \leq 1 + \epsilon$. Hence for all $x \in \mathcal{G}$,

$$(1 - \epsilon)\widehat{\text{vol}}(\mathcal{P}^y) \leq |\mathcal{I}(\mathcal{P}^x)| \leq (1 + \epsilon)\widehat{\text{vol}}(\mathcal{P}^y).$$

So by (4) and (5), the value $|\Sigma'|$ lies within $(1 \pm \epsilon)|\Sigma_{r,c,b}|$, and our algorithm is indeed an *fpras*.

3.1 Dynamic Programming

Phase 1: Compute $w_t = |\Sigma_{t,c(\nu),b}|$ for every $t \in S_\nu$.

Recall $s = r - t$. There are two cases where we apply dynamic programming:

- (1) If there is some $i \in [m], j \in [n]$ such that $b_{ij} \geq n^{2m\beta}$, we apply dynamic programming to the small columns $[\nu]$ to calculate $|\Sigma_{t,c(\nu),b}|$ for every partition $t \in S_\nu$.
- (2) If $b_{ij} < n^{2m\beta}$ for all $i \in [m], j \in [n]$, we apply dynamic programming to calculate $|\Sigma_{r,c,b}|$ exactly. We will refer to this case as the $\nu = n$ case.

We consider each column h ($1 \leq h \leq \nu$) in increasing order, and compute $|\Sigma_{t,c(h),b}|$ for each ordered partition $t \in S_h$. This is similar to the dynamic programming phase of the *fpras* of [5] for standard contingency tables with a constant number of rows.

By definition of a small column, we know that every cell bound is at most $c_j < n^p = n^{q+\beta} \leq n^{(2m-1/2)\beta}$ in the $\nu < n$ case, or less than $n^{2m\beta}$ in the $\nu = n$ case. In either case we have $b_{ij} < n^{2m\beta}$ for every cell in a small column.

If $h = 1$, then $|\Sigma_{t,c(1),b}| = 1$ for every partition t of c_1 into m parts which satisfies $t_i \leq b_{i1}$ for all $i \in [m]$. Therefore the cardinality of S_1 can be bounded by

$$|S_1| \leq \prod_{i=1}^{m-1} (b_{ij} + 1) \leq \prod_{i=1}^{m-1} n^{2m\beta} = n^{2m(m-1)\beta},$$

which is polynomial in n and ϵ^{-1} for m constant. Thus we can list S_1 in polynomial time.

If $2 \leq h \leq \nu$, we use the results from the computation on column $(h-1)$ to compute $|\Sigma_{t,c(h),b}|$. Let \hat{t} denote an element of S_{h-1} . Then the dynamic programming recurrence is

$$|\Sigma_{t,c(h),b}| = \sum_{\hat{t} \in S_{h-1}: (t-b_h) \leq \hat{t} \leq t} |\Sigma_{\hat{t},c(h-1),b}|, \quad (6)$$

since there is a unique extension to column h with values $x_{ih} = t_i - \hat{t}_i$ provided these satisfy $0 \leq x_{ih} \leq b_{ih}$. Therefore we can use the $|\Sigma_{\hat{t},c(h-1),b}|$ values constructed for column $(h-1)$ to compute $|\Sigma_{t,c(h),b}|$ for column h .

We now bound the running time of the dynamic programming algorithm. First we bound the number of possible $t \in S_h$. We have $b_{ij} < n^{2m\beta}$ for all $i \in [m], j \in [\nu]$. For any $i \in [m-1]$, $\min\{r_i, \sum_{j=1}^h b_{ij}\}$ is an upper bound on t_i , for any $t \in S_h$. Therefore any $t \in S_h$ must have $t_i < n^{2m\beta+1}$, since $h \in [n]$.

Therefore the cardinality of S_h can be bounded by

$$|S_h| \leq \prod_{i=1}^{m-1} \left(\sum_{j=1}^h b_{ij} + 1 \right) \leq \prod_{i=1}^{m-1} n^{2m\beta+1} = n^{(m-1)(2m\beta+1)}.$$

Thus S_h can be listed in $O(n^{(m-1)(2m\beta+1)})$ time, which is polynomial in n and ϵ^{-1} . For each $t \in S_h$, we can calculate $|\Sigma_{t,c(h),b}|$ using equation (6). There are at most $2n^{(m-1)(2m\beta+1)}$ elements of S_h , and for each such element we sum over at most $2n^{(m-1)(2m\beta+1)}$ elements of S_{h-1} . Therefore the h^{th} phase of the dynamic programming algorithm takes $O(n^{2(m-1)(2m\beta+1)})$ time. There are at most n phases of dynamic programming. Therefore the running time of the entire algorithm is bounded above by $O(n^{2(m-1)(2m\beta+1)+1})$, which is polynomial in n and ϵ^{-1} .

If $\nu = n$, then we are done. If $\nu < n - 1$, we obtain a polynomial-sized set of weights $\{w_t : t \in S_\nu\}$ with $w_t = |\Sigma_{t,c(\nu),b}|$, such that

$$|\Sigma_{r,c,b}| = \sum_{t \in S_\nu} w_t |\Sigma_{s,c(\bar{\nu}),b}|. \quad (7)$$

Phase 2: Compute $W_t = |\Sigma_{s(\sigma),*,b}| \cdot |\Sigma_{S'}|$ for every $t \in S_\nu$.

We assume that there is at least one cell with $b_{ij} > n^{2m\beta}$. Then, since the original cell bounds are tight, there is some $j' \neq j$ such that $b_{ij'} \geq n^{2m\beta-1}$. By definition of p , both j and j' are large columns. So there are at least two large columns, i.e. $\nu \leq n - 2$.

For the remainder of this section, we omit the $\bar{\nu}$ from $c(\bar{\nu})$, $b(\bar{\nu})$, since we are always working with the large columns $[n] \setminus [\nu]$. For each $t \in S_\nu$, we compute the total number of assignments to the small rows and small cells of the table on the large columns with row sums $s = r - t$. We will use dynamic programming to count the following quantity exactly:

$$|\Sigma_{s(\sigma),*,b}| = \prod_{i=1}^{\sigma} |\Sigma_{s_i,*,b}^{(i)}|, \quad (8)$$

where $\Sigma_{s_i,*,b}^{(i)} =_{\text{def}} \{(z_{\nu+1}, \dots, z_n) : \sum_{j=\nu+1}^n z_j = s_i, 0 \leq z_j \leq b_{ij}\}$.

For every $i \in [\sigma]$, we compute $|\Sigma_{s_i,*,b}^{(i)}|$ by dynamic programming. For $\nu + 1 \leq h \leq n$, and for every $s'_i \leq s_i$, let $\Sigma_{s'_i,*,b}^{(i,h)} =_{\text{def}} \{(z_{\nu+1}, \dots, z_h) : \sum_{j=\nu+1}^h z_j = s'_i, 0 \leq z_j \leq b_{ij}\}$. Define $|\Sigma_{s'_i,*,b}^{(i,\nu)}| = 1$. Then we use the following recurrence to compute $|\Sigma_{s_i,*,b}^{(i)}| = |\Sigma_{s_i,*,b}^{(i,n)}|$:

$$|\Sigma_{s'_i,*,b}^{(i,h+1)}| = \sum_{x_{h+1}=0}^{\min\{s'_i, b_{i,h+1}\}} |\Sigma_{(s'_i-x_{h+1}),*,b}^{(i,h)}|. \quad (9)$$

Every row $i \in [\sigma]$ satisfies $r_i < mn^{q+1}$, so $s_i < mn^{q+1}$. Therefore we consider at most mn^{q+1} values for s'_i . For each s'_i , we consider at most $b_{i,h+1} + 1 \leq mn^{q+1}$ values for x_{h+1} . Therefore computing $|\Sigma_{s'_i,*,b}^{(i,h)}$ for a single value of s'_i requires $O(n^{q+1})$ time. For given h , we compute $|\Sigma_{s'_i,*,b}^{(i,h)}$ for all $s'_i \leq s_i$ in $O(n^{2(q+1)})$ time. There are $O(n)$ values of h , so we compute $|\Sigma_{s_i,*,b}^{(i)}|$ in $O(n^{2q+3})$ time. We can compute $|\Sigma_{s_i,*,b}^{(i)}$ for all $i \in [\sigma]$ in $O(n^{2q+3})$ time, since $\sigma = O(1)$. Thus we can compute $|\Sigma_{s(\sigma),*,b}|$ in $O(n^{2q+3})$ time, using (8).

Finally, consider the set of small cells $S' \subseteq \mathbf{R}$. For every $(i, j) \in S'$, we have $b_{ij} < n^q$. Thus there are at most $b_{ij} + 1 \leq n^q$ feasible values for each cell $(i, j) \in S'$. To determine $|\Sigma_{S'}|$, we treat each $(i, j) \in S'$ separately. It can be computed in $O(mn)$ time as:

$$|\Sigma_{S'}| = \prod_{(i,j) \in S'} (b_{ij} + 1).$$

Hence we can compute $W_t = |\Sigma_{s(\sigma),*,b}| \cdot |\Sigma_{S'}|$ for all $t \in S_\nu$.

3.2 Volume Estimation

We state here some facts which were used earlier to show that our algorithm is an *fpras*. Let \mathcal{N} be the network whose flows correspond to elements of $\Sigma_{r,c,b}$. Consider any $\mathbf{x} \in \mathcal{G}$, and let $\mathcal{N}^{\mathbf{x}}$

be its residual network. Any small cell $(i, j) \in S'$ in R will have been assigned a value by x . We refer to these as *blocked* cells, the remainder being *unblocked*. By the definitions of small rows, small cells and the small columns, it seems likely that the values assigned by x should not greatly influence the number of flows in \mathcal{N}^x . We will prove that this is in fact the case. We first prove the existence of a “central” rational flow in \mathcal{N}^x . For any component C of \mathcal{N}^x , let $T_C = B_C \uplus D_C$ be a maximum spanning tree on C . Let $T' = \bigcup_C T_C$, $B' = \bigcup_C B_C$ and $D' = \bigcup_C D_C$.

Theorem 8 *Let f be the rational flow on \mathcal{N} defined in §2. Suppose $x \in \mathcal{G}$ yields residual network \mathcal{N}^x . Then there is a rational flow g^x in \mathcal{N}^x such that*

$$\begin{aligned} g_{ij}^x &\in f_{ij} \pm (m+1)n^{p+1}, & (i, j) \in B', \\ g_{ij}^x &\in f_{ij} \pm m^2n^{q+1}, & (i, j) \in D', \\ g_{ij}^x &= f_{ij}, & (i, j) \in R \setminus T'. \end{aligned}$$

Also, the tight bounds for any cell (i, j) in \mathcal{N}^x satisfy $l_{ij} \leq g_{ij}^x - b_{ij}/4mn$ and $b'_{ij} \geq g_{ij}^x + b_{ij}/4mn$.

Proof: Let $T = B \uplus D$ be the original spanning tree for \mathcal{N} . Suppose C is a component of \mathcal{N}^x . We first show that there is a maximum spanning tree $T_C = B_C \uplus D_C$ in C of unblocked cells, constructed using the original b_{ij} , such that $b_{ij} > n^p/m$ for every $(i, j) \in D_C$ and $b_{ij} > n^r$ for every $(i, j) \in B_C$.

Let I_C be the set of rows and J_C the set of columns spanned by C , so B_C is a subtree of T_C spanning I_C . All cells in D_C satisfy $b_{ij} > n^p/m$, since $c_j > n^p$ for all $j \in J_C$ and the b_{ij} were tight for \mathcal{N}^x . Only cells with $b_{ij} < n^p$ have been removed from \mathcal{N} , so all cells in the subtree $B' = C \cap B$ of T spanning I_C remain, since $r > p + 3$. Therefore every cell $(i, j) \in B_C$ also satisfies $b_{ij} > n^r$. Otherwise we can obtain a spanning tree of greater weight by removing any $(i, j) \in B_C$ with $b_{ij} < n^r$ from T_C , and reconnecting the resulting components by the appropriate edge of B' .

Let f be the rational flow defined in §2, having slack $b_{ij}/2mn$ on every cell $(i, j) \in A$. We modify f to obtain a rational flow g^x in \mathcal{N} so that $g_{ij}^x = x_{ij}$ for all $(i, j) \notin R \setminus S$ and every cell $(i, j) \in R \setminus S$ has slack at least $b_{ij}/4mn$.

Let C be any component of \mathcal{N}^x . The residual row and column sums of the table on $I_C \times J_C$ depend on x . For every $i \in I_C$, the residual row sums $\hat{r}_i(x)$ satisfy

$$\begin{aligned} r_i \geq \hat{r}_i(x) &= r_i - \sum_{j=1}^{\nu} x_{ij} - \sum_{j>\nu, (i,j) \in S} x_{ij} \\ &\geq r_i - \sum_{j=1}^{\nu} b_{ij} - \sum_{j>\nu, (i,j) \in S} b_{ij} \\ &> r_i - \nu n^p - (n - \nu)n^q \\ &\geq r_i - n^{p+1}. \end{aligned}$$

The residual sums \hat{c}_j for columns $j \in J_C$ satisfy

$$\begin{aligned}
c_j \geq \hat{c}_j(\mathbf{x}) &= c_j - \sum_{i=1}^{\sigma} x_{ij} - \sum_{i>\sigma, (i,j) \in S} x_{ij} \\
&\geq c_j - \sum_{i=1}^{\sigma} b_{ij} - \sum_{i>\sigma, (i,j) \in S} b_{ij} \\
&> c_j - \sigma n^q - (m - \sigma)mn^{q+1} \\
&\geq c_j - m^2 n^{q+1}.
\end{aligned}$$

Clearly $|\hat{r}_i(\mathbf{x}) - \hat{r}_i(\mathbf{f})| \leq n^{p+1}$ for $i \in [m] \setminus [\sigma]$ and $|\hat{c}_j(\mathbf{x}) - \hat{c}_j(\mathbf{f})| \leq m^2 n^{q+1}$ for $j \in [n] \setminus [\nu]$.

Suppose that we modify \mathbf{f} to obtain a new function \mathbf{f}^x , as follows:

$$f_{ij}^x = \begin{cases} x_{ij} & \text{for } (i, j) \notin R \setminus S \\ f_{ij} & \text{for } (i, j) \in R \setminus S \end{cases}$$

The function \mathbf{f}^x is no longer even a rational flow in \mathcal{N} . However, we will modify \mathbf{f}^x to produce a rational flow \mathbf{g}^x , by changing only the cells in the trees T_C . First we ‘‘correct’’ the column sums for \mathbf{f}^x in J_C . Let $j \in J_C$ be such that $\sum_{i=1}^m f_{ij}^x \neq c_j$. Let $i^* \in [m]$ be such that $b_{i^*j} = \max_{i \in I_C} b_{ij}$. We may assume that $(i^*, j) \in \mathsf{T}_C$. Then $b_{i^*j} \geq n^p/m$, so

$$\frac{b_{i^*j}}{2mn} \geq \frac{n^{p-1}}{2m^2} = \frac{n^{q+\beta/2-1}}{2m^2} \geq \frac{32m^6 n^{q+2}}{2m^2} = 16m^4 n^{q+2}.$$

By definition of \mathbf{f} , $\frac{1}{2mn} b_{i^*j} \leq f_{i^*j}^x = f_{i^*j} \leq (1 - \frac{1}{2mn}) b_{i^*j}$. Thus we can add or subtract $16m^4 n^{q+2}$ to $f_{i^*j}^x$ and the resulting value satisfies the cell bounds. Since $|\hat{c}_j(\mathbf{x}) - \hat{c}_j(\mathbf{f})| \leq m^2 n^{q+1}$, we can add $\hat{c}_j(\mathbf{x}) - \hat{c}_j(\mathbf{f})$ to $f_{i^*j}^x$ and maintain the cell bounds for (i^*, j) . Since $m^2 n^{q+1} \leq b_{i^*j}/32m^3 n^2$, the resulting value in cell (i, j) will still have slack of at least $b_{i^*j}/3mn$. Define $h_{i^*j}^x = f_{i^*j} + \hat{c}_j(\mathbf{x}) - \hat{c}_j(\mathbf{f})$, $h_{ij}^x = f_{ij}$ ($i \neq i^*$). Then

$$\begin{aligned}
\sum_{i=1}^m h_{ij}^x &= \sum_{i=1}^{\sigma} x_{ij} + \sum_{i=\sigma+1}^m f_{ij} + \hat{c}_j(\mathbf{x}) - \hat{c}_j(\mathbf{f}) \\
&= c_j - \hat{c}_j(\mathbf{x}) + \sum_{i=\sigma+1}^m f_{ij} + \hat{c}_j(\mathbf{x}) - \hat{c}_j(\mathbf{f}) \\
&= c_j + \sum_{i=\sigma+1}^m f_{ij} - (c_j - \sum_{i=1}^{\sigma} f_{ij}) \\
&= c_j,
\end{aligned}$$

and column total j is satisfied. Note that we corrected column sum j by changing one cell in column j . Therefore, we can correct each of the column sums for all C and $j \in J_C$ independently. Letting $\mathsf{T}' = \bigcup_C \mathsf{T}_C$, we obtain a new function \mathbf{h}^x satisfying

$$\begin{aligned}
h_{ij}^x &= x_{ij}, & (i, j) &\notin R \setminus S, \\
h_{ij}^x &= f_{ij}, & (i, j) &\in R \setminus (S \cup \mathsf{T}'), \\
h_{ij}^x &\in f_{ij} \pm m^2 n^{q+1}, & (i, j) &\in \mathsf{T}'
\end{aligned}$$

with column sums c_j ($j \in [n]$) and row sums r_i ($i \in [\sigma]$). The row sums for h^x lie in $r_i \pm n^{p+1}$ ($i \in [m] \setminus [\sigma]$), since

$$\begin{aligned} r_i - \nu n^p - (n - \nu)m^2 n^{q+1} &\leq \sum_{j=1}^n h_{ij}^x \\ &\leq r_i + \nu n^p + (n - \nu)m^2 n^{q+1}. \end{aligned}$$

Now we correct the row sums for all C and $i \in I_C$. Define $J_C(i) = \{j \in J_C : (i, j) \notin S'\}$ and $I_C(j) = \{i \in I_C : (i, j) \notin S'\}$. Let

$$\text{err}_i(h^x) = \sum_{j \in J_C(i)} h_{ij}^x - \hat{r}_i(x)$$

be the signed *error* in row $i \in J_C$. Note that

$$\begin{aligned} \sum_{i \in I_C} \text{err}_i(h^x) &= \sum_{i \in I_C} \sum_{j \in J_C(i)} h_{ij}^x - \sum_{i \in I_C} \hat{r}_i(x) \\ &= \sum_{j \in J_C} \hat{c}_j(x) - \sum_{i \in I_C} \hat{r}_i(x) = 0. \end{aligned}$$

We now correct the row totals by considering pairs of rows i, i' such that $\text{err}_i > 0$ and $\text{err}_{i'} < 0$. Let $\varphi = \min\{\text{err}_i, -\text{err}_{i'}\}$, so $0 < \varphi \leq n^{p+1}$. Let $P_{i,i'}$ be the unique path in T_C from i to i' . Observe that $P_{i,i'} \subseteq B_C$ for every $i, i' \in I_C$, and $b_{kj} \geq n^r$ for every $(k, j) \in B_C$.

We modify h^x in the cells of $P_{i,i'}$ by routing flow φ from i to i' . Note that $P_{i,i'}$ has odd length, since \mathcal{N}^x is bipartite. List the cells of $P_{i,i'}$ in order of appearance from i to i' , subtract φ from all cells in odd positions in this list and add φ to cells in even positions. The only change to sums over rows or columns are in rows i and i' . Denote the updated solution by \tilde{h}^x . Then

$$\begin{aligned} \text{err}_i(\tilde{h}^x) &= \text{err}_i(h^x) - \varphi, \\ \text{err}_{i'}(\tilde{h}^x) &= \text{err}_{i'}(h^x) + \varphi, \\ \text{err}_k(\tilde{h}^x) &= \text{err}_k(h^x) \quad (k \neq i, i'). \end{aligned}$$

Clearly no error is increased in absolute value. Furthermore, either the new error for row i is zero or the new error for row i' is zero. Thus the routing exactly corrects either row i or row i' . We perform this procedure iteratively, at each step choosing a pair of rows i, i' such that the sum in row i has positive error and the sum in row i' is negative. Since we correct at least one row sum at every step, we do this at most m times to obtain a function g^x with row sums $\hat{r}_i(x)$ and column sums $\hat{c}_j(x)$. The cells which are altered (the B_C cells) during the row-correcting process still satisfy their cell bounds. We know that the total amount of flow routed from all i to i' pairs is at most mn^{p+1} . Therefore no cell changes by more than mn^{p+1} from h^x to g^x . During the correction of column totals, from f^x to h^x , an element of B_C could have been changed once by at most $m^2 n^{q+1}$. Thus the total modification to any B_C cell is at most $mn^{p+1} + m^2 n^{q+1}$, which by definition of β, q , and p is at most $(m+1)n^{p+1}$. But the cells in B_C have upper bound $b_{ij} > n^r$. Hence, using the definitions of p, r , the slack for g^x in cell $(i, j) \in B_C$ is at least as big as $b_{ij}/2mn - (m+1)n^{p+1}$, which satisfies

$$\begin{aligned} b_{ij}/2mn - (m+1)n^{p+1} &\geq b_{ij}/4mn + (n^r/4mn - (m+1)n^{p+1}) \\ &\geq b_{ij}/4mn. \end{aligned}$$

Thus all cells $(i, j) \in \mathbf{B}_C$ satisfy their cell bounds with slack at least $b_{ij}/4mn$ in g^x . The cells in \mathbf{D}_C are modified only once by at most m^2n^{q+1} , in going from f^x to h^x . But these cells have upper bound $b_{ij} > n^p/m$. We have already shown that the cells altered in going from f^x to h^x have slack $b_{ij}/3mn$, therefore these cells will certainly have slack $b_{ij}/4mn$. No other cell is changed. Thus, for every $(i, j) \in \mathbf{T}_C$, the tight bounds have slack at least $b_{ij}/4mn$ in g^x . Now we can repeat the argument of §2 to show that, for any cell (i, j) , there are rational flows g', g'' in \mathcal{N}^x with $g'_{ij} = g_{ij}^x + b_{ij}/4mn$, $g''_{ij} = g_{ij}^x - b_{ij}/4mn$. \square

Theorem 9 *Let $x, y \in \mathcal{G}$, let C be any component of \mathbf{R} , and let $\mathcal{P}_C^x, \mathcal{P}_C^y$ be the corresponding flow polytopes. Then*

- (i) $\mathcal{I}(\mathcal{P}_C^x) \in (1 \pm \epsilon/2m)\text{vol}(\mathcal{P}_C^x)$,
- (ii) $\text{vol}(\mathcal{P}_C^x) \in (1 \pm \epsilon/2m)\text{vol}(\mathcal{P}_C^y)$.

Proof: (i): From Theorem 8, the tight cell bounds satisfy $b'_{ij} - \ell_{ij} \geq b_{ij}/2mn$ for all $(i, j) \in \mathbf{R} \setminus \mathbf{S}$. Now, since

$$\begin{aligned} b'_{ij} - \ell_{ij} &\geq \frac{b_{ij}}{2mn} \geq \frac{n^q}{2mn} \\ &\geq \frac{n^{6+2\log_n(32m^6\epsilon^{-1})}}{2mn} \\ &> \frac{16n^3m^4}{\epsilon} = (mn)^3 \frac{16m}{\epsilon}, \end{aligned}$$

\mathcal{N}_C^x satisfies the conditions of Theorem 2 with $\epsilon = \epsilon/8m$. Therefore, since $1 - \theta \leq e^{-\theta}$ and $e^\theta < 1 + 2\theta$ ($0 \leq \theta \leq \frac{1}{2}$),

$$(1 - \epsilon/2m)\text{vol}(\mathcal{P}_C^x) \leq \mathcal{I}(\mathcal{P}_C^x) \leq (1 + \epsilon/2m)\text{vol}(\mathcal{P}_C^x).$$

(ii): Observe that \mathcal{P}_C^x is the set of points z satisfying

$$\begin{aligned} \sum_{j \in J_C(i)} z_{ij} &= \hat{r}_i(x), & i \in I_C, \\ \sum_{i \in I_C(j)} z_{ij} &= \hat{c}_j(x), & j \in J_C, \\ 0 &\leq z_{ij} \leq b_{ij}, & (i, j) \in C. \end{aligned} \tag{10}$$

In Theorem 8 we constructed a rational flow $g^x \in \mathcal{P}_C^x$ satisfying

$$\frac{1}{4mn}b_{ij} \leq g_{ij}^x \leq b_{ij}(1 - \frac{1}{4mn}) \quad \text{for all } (i, j) \in C,$$

where the b_{ij} are the bounds which were tight for \mathcal{N} . They are not necessarily tight for \mathcal{N}_C^x , but here we choose to work with the original bounds. Rewriting (10) relative to g^x , by setting $Z_{ij} = z_{ij} - g_{ij}^x$ for all $(i, j) \in C$, gives

$$\begin{aligned} \sum_{j \in J_C(i)} Z_{ij} &= 0, & i \in I_C, \\ \sum_{i \in I_C(j)} Z_{ij} &= 0, & j \in J_C, \\ -g_{ij}^x &\leq Z_{ij} \leq b_{ij} - g_{ij}^x, & (i, j) \in C. \end{aligned} \tag{11}$$

For $(i, j) \in \mathbf{D}_C$, let $I'_C(j) = I_C(j) \setminus \{(i, j)\}$. Now, by eliminating Z_{ij} , for $(i, j) \in \mathbf{T}_C$, we get the following full-dimensional representation for (11). (See Schrijver [27, §13].)

$$\begin{aligned} -g_{ij}^x &\leq \sum_{(k, \ell) \in P_{ij}} Z_{k\ell} - \sum_{(k, \ell) \in N_{ij}} Z_{k\ell} \leq b_{ij} - g_{ij}^x & (i, j) \in \mathbf{B}_C, \\ -g_{ij}^x &\leq -\sum_{k \in I'_C(j)} Z_{kj} \leq b_{ij} - g_{ij}^x, & (i, j) \in \mathbf{D}_C, \\ -g_{ij}^x &\leq Z_{ij} \leq b_{ij} - g_{ij}^x, & (i, j) \in C \setminus \mathbf{T}_C. \end{aligned} \tag{12}$$

where $(k, \ell) \in N_{ij}$ if the directed path in \mathbb{T}_C from row k to column ℓ traverses arc (i, j) in the direction i to j , and $(k, \ell) \in P_{ij}$ if the directed path in \mathbb{T}_C from k to ℓ traverses (i, j) in the direction j to i .

Similarly, for any other $\mathbf{y} \in \mathcal{G}$, $\mathcal{P}_C^{\mathbf{y}}$ is the set of points Z' satisfying

$$\begin{aligned} -g_{ij}^{\mathbf{y}} &\leq \sum_{(k,\ell) \in P_{ij}} Z'_{k\ell} - \sum_{(k,\ell) \in N_{ij}} Z'_{k\ell} \leq b_{ij} - g_{ij}^{\mathbf{y}}, & (i, j) \in \mathbb{B}_C, \\ -g_{ij}^{\mathbf{y}} &\leq -\sum_{k \in I'_C(j)} Z'_{kj} \leq b_{ij} - g_{ij}^{\mathbf{y}}, & (i, j) \in \mathbb{D}_C, \\ -g_{ij}^{\mathbf{y}} &\leq Z'_{ij} \leq b_{ij} - g_{ij}^{\mathbf{y}}, & (i, j) \in C \setminus \mathbb{T}_C. \end{aligned} \quad (13)$$

But our construction of \mathbf{g}^x and \mathbf{g}^y ensures that $g_{ij}^x = f_{ij} = g_{ij}^y$ for every $(i, j) \in C \setminus \mathbb{T}_C$. $Z'_{ij} = Z_{ij}$ for all $(i, j) \in C \setminus \mathbb{T}_C$ and (13) can be written as

$$\begin{aligned} -g_{ij}^{\mathbf{y}} &\leq \sum_{(k,\ell) \in P_{ij}} Z_{k\ell} - \sum_{(k,\ell) \in N_{ij}} Z_{k\ell} \leq b_{ij} - g_{ij}^{\mathbf{y}}, & (i, j) \in \mathbb{B}_C, \\ -g_{ij}^{\mathbf{y}} &\leq -\sum_{k \in I'_C(j)} Z_{kj} \leq b_{ij} - g_{ij}^{\mathbf{y}}, & (i, j) \in \mathbb{D}_C, \\ -g_{ij}^x &\leq Z_{ij} \leq b_{ij} - g_{ij}^x, & (i, j) \in C \setminus \mathbb{T}_C. \end{aligned} \quad (14)$$

Let $\varepsilon = \epsilon/4m^2n$. We now show $\mathcal{P}_C^x \subseteq (1 + \varepsilon)\mathcal{P}_C^y$. It follows from (14) that $(1 + \varepsilon)\mathcal{P}_C^y$ is the set of points Z satisfying

$$\begin{aligned} -(1 + \varepsilon)g_{ij}^{\mathbf{y}} &\leq \sum_{(k,\ell) \in P_{ij}} Z_{k\ell} - \sum_{(k,\ell) \in N_{ij}} Z_{k\ell} \leq (1 + \varepsilon)(b_{ij} - g_{ij}^{\mathbf{y}}), & (i, j) \in \mathbb{B}_C, \\ -(1 + \varepsilon)g_{ij}^{\mathbf{y}} &\leq -\sum_{k \in I'_C(j)} Z_{kj} \leq (1 + \varepsilon)(b_{ij} - g_{ij}^{\mathbf{y}}), & (i, j) \in \mathbb{D}_C, \\ -(1 + \varepsilon)g_{ij}^x &\leq Z_{ij} \leq (1 + \varepsilon)(b_{ij} - g_{ij}^x), & (i, j) \in C \setminus \mathbb{T}_C. \end{aligned} \quad (15)$$

We must show that every $Z \in \mathcal{P}_C^x$ satisfies (15). Clearly Z satisfies the inequalities for $C \setminus \mathbb{T}_C$. Consider the inequalities for \mathbb{D}_C . For $(i, j) \in \mathbb{D}_C$, we have $g_{ij}^y, g_{ij}^x \in f_{ij} \pm m^2n^{q+1}$ by Theorem 8, and therefore $|g_{ij}^x - g_{ij}^y| \leq 2m^2n^{q+1}$. Therefore, we must show that $\varepsilon \min\{b_{ij} - g_{ij}^y, g_{ij}^y\} \geq 2m^2n^{q+1}$. We have $\min\{b_{ij} - g_{ij}^y, g_{ij}^y\} \geq b_{ij}/4mn$ from Theorem 8 and we know that $b_{ij} > n^p/m$ for all $(i, j) \in \mathbb{D}_C$. Thus

$$\begin{aligned} \varepsilon \min\{b_{ij} - g_{ij}^y, g_{ij}^y\} &\geq \frac{\varepsilon n^p}{4m^2n} = \frac{\varepsilon n^{p-2}}{16m^4} \\ &= \frac{\varepsilon n^{q+1+\log_n(32m^6\varepsilon^{-1})}}{16m^4} \\ &= 2m^2n^{q+1}, \end{aligned}$$

as required. Finally, consider the cells $(i, j) \in \mathbb{B}_C$. We have $g_{ij}^y, g_{ij}^x \in f_{ij} \pm (m+1)n^{p+1}$ by Theorem 8, so $|g_{ij}^y - g_{ij}^x| \leq 2(m+1)n^{p+1}$. Therefore, we must show that $\varepsilon \min\{b_{ij} - g_{ij}^y, g_{ij}^y\} \geq 2(m+1)n^{p+1}$. We have $\min\{b_{ij} - g_{ij}^y, g_{ij}^y\} \geq b_{ij}/4mn$ from Theorem 8 and we know that $b_{ij} > n^r$

for all $(i, j) \in \mathbb{B}_C$. Thus

$$\begin{aligned} \varepsilon \min\{b_{ij} - g_{ij}^{\mathbf{y}}, g_{ij}^{\mathbf{y}}\} &\geq \frac{\varepsilon n^r}{4mn} = \frac{\varepsilon n^{r-2}}{16m^3} \\ &\geq \frac{\varepsilon n^{p+1+\log_n(32m^6\varepsilon^{-1})}}{16m^3} \\ &= 2m^3 n^{p+1} > 2(m+1)n^{p+1}. \end{aligned}$$

We have now shown that $\mathcal{P}_C^{\mathbf{x}} \subseteq (1+\varepsilon)\mathcal{P}_C^{\mathbf{y}}$. Therefore

$$\begin{aligned} \text{vol}(\mathcal{P}_C^{\mathbf{x}}) &\leq \text{vol}((1+\varepsilon)\mathcal{P}_C^{\mathbf{y}}) \\ &\leq (1+\varepsilon)^{d_C} \text{vol}(\mathcal{P}_C^{\mathbf{y}}), \end{aligned}$$

where $d_C = (|I_C| - 1)(|J_C| - 1) - |S'| \leq mn$ is the dimension of $\mathcal{P}_C^{\mathbf{x}}$ (and $\mathcal{P}_C^{\mathbf{y}}$). Therefore

$$\begin{aligned} \text{vol}(\mathcal{P}_C^{\mathbf{x}}) &\leq (1+\varepsilon)^{mn} \text{vol}(\mathcal{P}_C^{\mathbf{y}}) \\ &= \left(1 + \frac{\varepsilon}{4m^2n}\right)^{mn} \text{vol}(\mathcal{P}_C^{\mathbf{y}}) \\ &\leq \left(1 + \frac{\varepsilon}{2m}\right) \text{vol}(\mathcal{P}_C^{\mathbf{y}}), \end{aligned}$$

using $(1 + \theta/\kappa)^\kappa \leq 1 + 2\theta$ for $0 \leq \theta \leq \frac{1}{2}$ and $\kappa > 0$.

Switching the roles of \mathbf{x} and \mathbf{y} , we also have $\text{vol}(\mathcal{P}_C^{\mathbf{y}}) \leq (1+\varepsilon)^{mn} \text{vol}(\mathcal{P}_C^{\mathbf{x}})$. Then it follows, using $(1 + \theta/\kappa)^{-\kappa} \geq 1 - 2\theta$ for $0 \leq \theta \leq \frac{1}{2}$ and $\kappa > 0$, that

$$\begin{aligned} \text{vol}(\mathcal{P}_C^{\mathbf{x}}) &\geq (1+\varepsilon)^{-mn} \text{vol}(\mathcal{P}_C^{\mathbf{y}}) \\ &\geq \left(1 - \frac{\varepsilon}{2m}\right) \text{vol}(\mathcal{P}_C^{\mathbf{y}}). \end{aligned}$$

□

References

- [1] S. Aoki, Exact methods and Markov chain Monte Carlo methods of conditional inference for contingency tables, PhD Thesis, University of Tokyo, 2004.
- [2] A. Barvinok, A polynomial-time algorithm for counting integral points in polyhedra when the dimension is fixed, *Mathematics of Operations Research* **19**, pp. 769-779, 1994.
- [3] W. Baldoni-Silva, J. De Loera and M. Vergne, Counting integer flows in networks, *Foundations of Computational Mathematics*, **4**(3), pp. 277-314, 2004.
- [4] I. Bezáková, N. Bhatnagar and E. Vigoda, Sampling binary contingency tables with a greedy start, in *Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2006.
- [5] 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**(2): pp. 291-310, 2003.
- [6] 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 Annual IEEE Symposium on Foundations of Computer Science*, pp. 711-720, 2002.

- [7] J. De Loera and B. Sturmfels, Algebraic unimodular counting, *Mathematical Programming, Series B*, **96**, pp. 183–203, 2003.
- [8] P. Diaconis and B. Efron, Testing for independence in a two-way table: new interpretations of the chi-square statistic (with discussion), *Annals of Statistics* **13**, pp. 845–913, 1995.
- [9] P. Diaconis and A. Gangolli, Rectangular arrays with fixed margins, in *Discrete probability and algorithms* (D. Aldous, P. Diaconis, J. Spencer and M. Steele, eds.), IMA Volumes on Mathematics and its Applications **72**, Springer-Verlag, New York, pp. 15–41, 1995.
- [10] P. Diaconis and L. Saloff-Coste, Random walk on contingency tables with fixed row and column sums, Department of Mathematics, Harvard University, 1995.
- [11] P. Diaconis and B. Sturmfels, Algebraic algorithms for sampling from conditional distributions, *Annals of Statistics* **26**, pp. 363–397, 1998.
- [12] M. Dyer, Approximate counting by dynamic programming, in *Proc. 35th Annual ACM Symposium on the Theory of Computing*, pp. 693–699, 2003.
- [13] M. Dyer and C. Greenhill, Polynomial-time counting and sampling of two-rowed contingency tables. *Theoretical Computer Science* **246**, pp. 265–278, 2000.
- [14] M. Dyer, A. Frieze and R. Kannan, A random polynomial time algorithm for approximating the volume of convex bodies, *Journal of the ACM* **38**, pp. 1–17, 1991.
- [15] M. Dyer, A. Frieze, R. Kannan, A. Kapoor, L. Perkovic and U. Vazirani, A mildly exponential time algorithm for approximating the number of solutions to a multidimensional knapsack problem, *Combinatorics, Probability and Computing* **2**, 271–284, 1993.
- [16] M. Dyer, R. Kannan and J. Mount, Sampling contingency tables. *Random Structures & Algorithms* **10**, pp. 487–506, 1997.
- [17] M. Grötschel, L. Lovász and A. Schrijver, *Geometric algorithms and combinatorial optimization*, Springer-Verlag, 1991.
- [18] R. Holmes and L. Jones, On uniform generation of two-way tables with fixed margins and the conditional volume test of Diaconis and Efron, *Annals of Statistics* **24**, pp. 64–68, 1996.
- [19] 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), pp. 671–697, 2004.
- [20] M. Jerrum, L. Valiant and V. Vazirani, Random generation of combinatorial structures from a uniform distribution, *Theoretical Computer Science* **43**, pp. 169–188, 1986.
- [21] R. Kannan and S. Vempala, Sampling lattice points, in *Proc. of 29th Annual ACM Symposium on Theory of Computing*, pp. 696–700, 1997.
- [22] L. Lovász and S. Vempala, Simulated annealing in convex bodies and a $O^*(n^4)$ volume algorithm, in *Proc. 44th Annual IEEE Symposium on Foundations of Computer Science*, pp. 650–659, 2003.
- [23] B. Morris, Improved bounds for sampling contingency tables, in *Random Structures & Algorithms*, **21**(2), pp. 135–146, 2002.
- [24] 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.)
- [25] J. Mount, Fast unimodular counting, *Combinatorics, Probability and Computing* **9**, pp. 277–285, 2000.

- [26] F. Rapollo, Markov bases and structural zeros, *Journal of Symbolic Computation*, **41**(2), pp. 164-172, 2006.
- [27] A. Schrijver, *Combinatorial optimization—polyhedra and efficiency*, Springer-Verlag, Berlin, 2003.
- [28] L. Valiant, The complexity of computing the permanent, *Theoretical Computer Science* **8**, 189–201, 1979.