Simple Markov-chain algorithms for generating bipartite graphs and tournaments

(Extended Abstract)

Ravi Kannan* Prasad Tetali† Santosh Vempala‡

Abstract

We consider two problems: randomly generating labeled bipartite graphs with a given degree sequence and randomly generating labeled tournaments with a given score sequence. We analyze simple Markov chains for both problems. For the first problem, we cannot prove that our chain is rapidly mixing in general, but in the (near-) regular case, i.e. when all the degrees are (almost) equal, we give a proof of rapid mixing. Our methods also apply to the corresponding problem for general (nonbipartite) regular graphs which was studied earlier by several researchers. One significant difference in our approach is that our chain has one state for every graph (or bipartite graph) with the given degree sequence; in particular, there are no auxiliary states as in the chain used by Jerrum and Sinclair.

For the problem of generating tournaments, we are able to prove that our Markov chain on tournaments is rapidly mixing, if the score sequence is near-regular. The proof techniques we use for the two problems are similar.

1 Introduction

The problems of counting and random generation of (0,1)-matrices with prescribed row and column sums are of interest to combinatorialists, statisticians, and computer scientists. Theorems of Gale, Ryser and several others (see [7], [19]) on properties of such matrices and subsequent applications to tournaments, Hadamard matrices and latin squares made the study of these matrices particularly attractive. The same problem when we allow non-negative integers instead of just 0 or 1, is of interest in statistics (Diaconis et al.). We study these problems here and present some partial results.

Suppose we are given the degrees of the vertices of a labeled bipartite graph. Then how many different graphs can have this 'degree sequence' ? How do we select one of these graphs uniformly at random? Stated another way, given a sequence of $m+n$ positive integers $r_1, r_2, \ldots, r_m, c_1, c_2, \ldots, c_n$

*School of Computer Science, Carnegie Mellon University, Pittsburgh PA 15213. Email: kannan+@cs.cmu.edu.

†School of Mathematics, Georgia Institute of Technology, Atlanta GA 30332. Research supported by the NSF Grant No. CCR-9503952 Email: tetali@math.gatech.edu.

‡School of Computer Science, Carnegie Mellon University, Pittsburgh PA 15213.Supported in part by NSF National Young Investigator grant CCR-9357793. Email: svempala@cs.cmu.edu.

such that $\sum r_i = \sum c_j$, we wish to count the number of m by n, (0,1)-matrices A which satisfy the following property: the sum of the entries of the i^{th} row of A is r_i and the sum of j^{th} column is c_j for $i = 1, \ldots, m$ and $j = 1, \ldots, n$. Gale [7] and Ryser [19] were the first to obtain a simple easy-to-check criterion to determine if there exists at least one (0,1)-matrix with the given row and column sums. While the constructive proof of the above fact produces one such matrix, Ryser remarks that "the exact number of them is undoubtedly an extremely intricate function of the row and column sums."

Such matrices represent labeled bipartite graphs and the integer sequence can be called a *bipartite degree sequence*, as the numbers correspond to the degrees of the vertices in a bipartite graph. In this work we consider the problem of almost uniform generation of such matrices. We are able to give the first fully polynomial time algorithm for this problem in the regular case, namely, when all the r_i and c_j are equal.

The closely related problem of generating/counting labeled graphs (not bipartite) with a given degree sequence has received considerable attention in recent years. The works of Bender and Canfield [1], Bollobás [2], and Wormald [22] give us an algorithm for generating regular graphs that is efficient for degrees upto $O((\log n)^{1/2})$. After a series of improvements, [5, 20, 15], Jerrum and Sinclair [9], were able to solve the problem for regular graphs in polynomial time by considering a random walk through an associated Markov chain. They go on to show that their approach is efficient only for near-regular sequences. Our algorithm is also based on a random walk; however, unlike the chain studied in [9], our chain has no "auxiliary" states. So although we are unable to prove it, it is possible that our algorithm works for *all* degree sequences. Our analysis for regular sequences can be extended to prove rapid-mixing for slightly more general degree sequences (and bipartite degree sequences). A sufficient condition is that the minimum and maximum degrees are close to the average degree. A precise statement of this and its proof are left for the full version of the paper.

The corresponding problem for bipartite multigraphs (where multiple edges are allowed) is the problem of uniform random generation of matrices with nonnegative integer entries with given row and column sums. This problem is of importance in statistics where such matrices arise as "contingency tables". A Markov chain similar to the one we analyze has been proposed by Diaconis for the contingency table problem and it is a long-standing open problem as to whether this mixes rapidly. We hope our techniques shed some light on this problem. We also hope our techniques help in generating a latin square of order n uniformly at random, in view of the fact that the set of latin squares of order n is precisely the set of (0,1)-*cubes* of size $n \times n \times n$, with all the line sums equal to 1.

The second problem we are able to solve with this approach is that of almost uniform generation of a labeled tournament, given a near-regular score sequence on n players (or vertices). A tournament is a complete oriented graph. The *score sequence* (also called the *score vector*) of a tournament is the sequence of outdegrees of the vertices of the tournament. Let $\mathbf{s} = \mathbf{s_n} = (s_1, \ldots s_n)$ denote a score sequence on n vertices. With no loss of generality, we assume $s_1 \leq s_2 \leq \cdots \leq s_n$. Let $\mathcal{T}(\mathbf{s})$ denote the set of labeled tournaments (on n vertices) with the score sequence $\mathbf{s}$. We provide a fully polynomial algorithm for generating a member of $\mathcal{T}(\mathbf{s})$ *almost* uniformly at random, for a given $\mathbf{s}$. While the Markov chain we describe is irreducible for all score sequences, at the moment we can prove rapid mixing only for near-regular score sequences – $(s_1, \ldots, s_n)$ is *near-regular* if $\forall i$, $|s_i - (n-1)/2| = O(n^{3/4+\epsilon})$, for a sufficiently small $\epsilon > 0$.

Asymptotic enumeration of labeled tournaments has been of interest to combinatorialists and has been studied in some detail (see [21], [12], [13], [14]). The focus of the study in these papers is in providing an asymptotic estimate (of the type $(1+o(1))$ or $(1+O(1))$ of some explicit formula) of the exact number of labeled tournaments with a prescribed score sequence. For example, in [12], McKay proves that for any $\epsilon > 0$,

$$RT(n) = \left(\frac{2^{n+1}}{\pi n}\right)^{(n-1)/2} n^{1/2}e^{-1/2}(1+O(n^{-1/2+\epsilon})),$$

where $RT(n)$, defined for odd n, is the number of labeled regular tournaments which have the score sequence, $((n-1)/2, \cdots, (n-1)/2)$. The techniques used in these papers include the use of generating functions and integration in n complex dimensions, and thus are quite different from the Markov chain Monte-Carlo approach. The graph underlying our Markov chain for the tournaments has also been studied by Brualdi and Li [3]; in particular, bounds on the diameter of this graph for some special score sequences can be found there.

We use the Jerrum-Sinclair technique of canonical path counting to analyze the Markov chains involved in our schemes.

2 The Markov chain on bipartite graphs

Let $\mathbf{d} = \mathbf{d}_{mn} = (r_1, \ldots, r_m, c_1, \ldots c_n)$ be a bipartite degree sequence. We consider a Markov chain $M(\mathbf{d})$ whose states correspond to distinct graphs with degree sequence $\mathbf{d}$. A state $u \in M(\mathbf{d})$ has an associated matrix M_u and a bipartite graph G_u. Two states u, v are connected iff $M_v - M_u$ has exactly 4 non-zero entries. In other words if G_u and G_v are the bipartite graphs represented by the states u and v, they are adjacent in the chain iff the symmetric difference of their edges is a cycle of length 4 with edges alternately from G_u and G_v. (Note that this is the smallest possible difference between any two states of $M(\mathbf{d})$.) We sometimes refer to an edge of the Markov chain as a *switch* as it switches the edges on a cycle of length 4. For brevity of notation, we refer to $M(\mathbf{d})$ by M.

The random walk can now be described as follows. At a state u we choose two entries $(i,j), (k,l), 1 \leq i \leq k \leq m, 1 \leq j \leq l \leq n$ of M_u uniformly at random. If the entries $(i,j), (k,l)$ are both 1 and the entries $(i,l), (k,j)$ are both 0 then we go to the state where $(i,j), (k,l)$ are 0 and $(i,l), (k,j)$ are 1, i.e. we switch the cycle i, j, k, l, i. In any other case we stay at u.
In the next section we will set up paths between every pair of states in the Markov chain hence showing that it is irreducible. (The fact that the moves as described above can be used to transform a given (0,1)-matrix to any other with the same row and column sums was also another elegant observation of Ryser [18].)

2.1 Canonical Paths.

Let $\mathcal{C}$ be the set of all cycles in the complete bipartite graph on $m \times n$ vertices. Order these cycles, first by length and arbitrarily among cycles of the same length. Fix an order on the vertices of the graph. Each cycle is then a sequence of vertices starting with the vertex in the cycle that first appears in the vertex-ordering, and following the vertices along the cycle in the order which gives the earlier lexicographic sequence. The first vertex in the sequence of a cycle is also called the "start" vertex of the cycle.

Consider two states u and v of M. Let H be the symmetric difference of G_u and G_v. For illustration let the edges in H from G_u be red edges and those from G_v be blue edges. An *alternating cycle* in H is a cycle which has edges alternately from G_u and G_v, i.e., it is a sequence of vertices $w_1, \ldots, w_{2k}, w_1$ such that the edges $(w_1, w_2), (w_3, w_4), \ldots, (w_{2k-1}, w_{2k})$ are all red or all blue) and the edges $(w_2, w_3), (w_4, w_5), \ldots, (w_{2k}, w_1)$ are all of the other color. Such an alternating cycle also corresponds to a sequence of edges that are alternately present and absent in G_u (and G_v). The following simple observation plays a key role in defining canonical paths.

Lemma 1 *H is the union of a set of edge-disjoint alternating cycles.*

For each graph H we fix a unique decomposition into cycles. Let the cycle decomposition of H into edge-disjoint alternating cycles be $C_1, C_2, \ldots, C_k$. The same decomposition is fixed for all graphs H obtained by swapping the red and blue edges on some subset of cycles. This mapping from symmetric difference graphs to alternating cycle decompositions will be extended to graphs where a small number of cycles are allowed to overlap.

Let $H_1 = G_u \Delta C_1, H_i = H_{i-1} \Delta C_i, 1 < i < k$, then $G_v = H_{k-1} \Delta C_k$. Our canonical path from G_u to G_v will visit $H_1, \ldots, H_{k-1}$ in that order. To go from G_u to H_1 we look at the cycle C_1 in more detail. If C_1 is of length 4 then we can switch it to go from G_u to H_1. Otherwise, let $C_1 = (w_1, \ldots, w_{2l}, w_1)$ where the edges $(w_1, w_2), (w_3, w_4), \ldots, (w_{2l-1}, w_{2l})$ are present in G_u and the edges $(w_2, w_3), (w_4, w_5), \ldots, (w_{2l}, w_1)$ are present in H_1 (and also G_v). Here w_1 is the designated start vertex.

Lemma 2 *C_1 contains a vertex w_i such that the cycle $w_i, w_{i+1}, w_{2l-i}, w_{2l-i+1}, w_i$ can be switched in G_u.*

Proof. By induction on the size of C_1. If $|C_1| = 4$, then it has two red edges and two blue edges and the whole cycle can be switched in one step in G_u. If $|C_1| > 4$, consider the cycle $w_1, w_2, w_{2l-1}, w_{2l}, w_1$. We know that the edges (w_1, w_2) and (w_{2l-1}, w_{2l}) are red edges (present in G_u and not in G_v) and the edge (w_{2l}, w_1) is blue (present in G_v and not G_u). So if (w_2, w_{2l-1}) is not present in G_u, we can switch the 4-length cycle in G_u. However if (w_2, w_{2l-1}) is present in G_u, then $w_2, w_3, \ldots, w_{2l-1}, w_2$ is a shorter alternating cycle in G_u. By the induction hypothesis, we find a switchable 4-length cycle as required. □

Now we have the tools to define the canonical path between G_u and G_v. We find the first vertex w_i such that the cycle $w_1, w_2, \ldots, w_{i+1}, w_{2l-i}, \ldots, w_{2l}, w_1$ is alternating in G_u. So $w_i, w_{i+1}, \ldots, w_{2l-i+1}, w_i$ can be switched, and we switch it to get a graph G_z such that $G_z \Delta G_v$ consists of one (if w_i is w_1) or in general two smaller cycles $w_1, \ldots, w_i, w_{2l-i+1}, \ldots, w_{2l}, w_1$ and $w_{i+1}, \ldots, w_{2l-i}, w_{i+1}$. The first cycle is "erased" from the symmetric difference by a series of switches $w_{j-1}, w_j, w_{2l-j+1}, w_{2l-j+2}, w_{j-1}$ in the order $j = i, i-1, \ldots, 2$. This is because we switched at the first point in C_1 where a switch was possible and so all the edges (w_j, w_{2l-j+1}) must be alternately absent and present in G_u for $j = 1, 2, \ldots, i$. We call this part of the path *phase I.*

At the end of phase I we are left with a graph whose symmetric difference with H_1 is an alternating cycle smaller than C_1. The rest of the canonical path can be seen by a suitable induction hypothesis on the number of edges in the symmetric difference H.

3 A Condition for Rapid-Mixing

Let u, v be any two states of M. Let z be a state on the canonical path from u to v in M. We first prove some useful observations.

As before let $H = G_u \Delta G_v$. We represent an alternating cycle in H by an m by n matrix with 1's for the edges in G_v and -1's for the edges in G_u. Then $M_v = M_u + C_1 + C_2 + \cdots + C_k$. Each matrix C_i has every row and column sum equal to zero. Let $M_i = M_u + C_1 + C_2 + \cdots + C_i$. M_i corresponds to the graph H_i as defined in the previous section.

Lemma 3 *The matrix $\hat{M} = M_u + M_v - M_z$ has at most 3 entries that are not 0 or 1.*

Proof. If z corresponds to some M_i then $\hat{M} = M_u + M_v - M_z = M_u + C_{i+1} + \cdots + C_k$. Since these cycles are edge-disjoint by lemma 1, and they are alternating for G_u, all the entries of $\hat{M}$ are 0 or 1.

Assume z is a state between the two states corresponding to M_{i-1} and M_i. This means that at the point when the canonical path from u to v passes through z, it is working on the cycle C_i. Let $C_i = (w_1, w_2, \ldots, w_{2l}, w_1)$.

The state z could be of three types:

1. As described in the previous section, when the canonical path reaches C_i it finds the first place along C_i where a switch can be performed, i.e. the first vertex w_j with such that $w_j, w_{j+1}, w_{2l-j}, w_{2l-j+1}, w_j$ can be switched. We call this 4-cycle S_j. Such an S_j exists by lemma 2

   So $M_z = M_u + C_1 + \cdots + C_{i-1} + S_j$ and

   $\hat{M} = M_u + M_v - M_z = M_u + C_{i+1} + \cdots + C_k + C_i - S_j$. $\hat{M} + S_j$ is a (0,1)-matrix. So according to the nature of the switch $\hat{M}$ has the entries corresponding to (w_{j+1}, w_{2l-j}) and (w_{2l-j+1}, w_j) both equal to 2 or both equal to -1.

2. Now we have two shorter alternating cycles in G_u, call them $D_1 = w_1, w_2, \ldots w_j, w_{2l-j+1}, \ldots, w_{2l}, w_1$ and D_2. Since S_j was the first possible switch, $D_1 = S_{j-1} + S_{j-2} + \cdots + S_1$ where the switches are legal only in the order $j-1, j-2, \ldots, 1$. Each step sets one odd entry in $\hat{M}$ (starting with the entry corresponding to the edge (w_{2l-j+1}, w_j)) to 0 or 1, depending on whether it was -1 or 2, and creates a new odd entry keeping the total number of odd entries the same. At the end of this process, i.e. when we have switched S_1, $\hat{M}$ has one odd entry (corresponding to the edge (w_{j+1}, w_{2l-j})).

3. Then we find the first switch on the remaining (shorter) cycle $D_2 = w_{j+1}, w_{j+2}, \ldots, w_{2l-j}, w_{j+1}$. When we find one and do it, we may create two new odd entries in $\hat{M}$ taking the total number of odd entries to 3 and also create two new shorter cycles. Working backwards to erase the first one keeps the number of odd entries equal to 3 and then it drops to 1.

When the cycle C_i is completely erased from the symmetric difference $\hat{M}$ is once again a purely (0,1)-matrix. ∎

Let $M' = M'(\mathbf{d})$ be the set of matrices that have the same row and column sums as matrices in M and further each matrix of M' has up to 3 entries that are not 0 or 1. These 'odd' entries have to be -1 or 2.

To derive a condition for rapid mixing, we need to look at a more complicated set-up. Given a bipartite degree sequence $\mathbf{d}$, we construct a graph G_d such that the perfect matchings in G_d map to (0,1)-matrices for $\mathbf{d}$.

Given $\mathbf{d} = (r_1, \ldots, r_m, c_1, \ldots, c_n)$, start with the complete bipartite graph with m and n vertices in the two color classes. Label the vertices on one side $a_1, \ldots, a_m$ and on the other $b_1, \ldots, b_m$. Subdivide each edge (a_i, b_j) with two new vertices x_{ji} and y_{ij} so that the edge from a_i to b_j is now a path a_i, x_{ij}, y_{ij}, b_j. Finally, replace each vertex a_i with $n - r_i$ vertices each with edges to each x_{ij} for $1 \leq j \leq n$. Similarly replace each vertex b_j with $m - c_j$ vertices each with edges to each y_{ij} for $1 \leq i \leq m$. Let the graph obtained be G_d.

Lemma 4 *A perfect matching in the graph G_d maps to a (0,1)-matrix with the sequence* $\mathbf{d}$. *Moreover, there are exactly $N = \Pi(n - r_i)!\Pi(m - c_i)!$ perfect matchings mapping to a single matrix.*

We shall use this construction to set up N^2 paths between every pair of states in the Markov chain, i.e. one path between every pair of perfect matchings. These paths will have the property that no more than $N^2|M'|$ go through any single edge of the Markov chain.

Consider the set of all perfect matchings of G_d. We will describe a path between every pair of perfect matchings. In general, these paths will use two kinds of moves. One set correspond to switches of the Markov chain. The second set are moves that connect perfect matchings corresponding to the same state of the chain. We connect two perfect matchings if their symmetric difference consists of exactly 4 edges, and none of them is in the induced graph G_{xy}. These moves are called *short* moves and the ones corresponding to switches are called *long* moves.

Here is a sketch of the scheme: Let F_u and F_v be perfect matchings of G_d such that F_u maps to u and F_v maps to v. Consider the set of edges $F_u \Delta F_v$. It is a set of vertex-disjoint even cycles $C_1, \ldots, C_k$ in G_d. We can erase these cycles from the symmetric difference following a fixed global order on all cycles of G_d and the two kinds of moves available. Let F_z be an intermediate perfect matching on the path from F_u to F_v. Given the graph $F' = F_u \Delta F_v \Delta F_z$, the matching F_z and the current move, we can uniquely determine the matchings F_u and F_v. First compute $F_u \Delta F_v = F' \Delta F_z$. From the current move and the global order on cycles of G_d we know that if the current move is affecting the cycle C_i in the sequence $C_1, \ldots, C_i, \ldots, C_k$ then $C_1, \ldots C_{i-1}$ have already been switched while $C_{i+1}, \ldots, C_k$ have not been switched. The next theorem makes this precise.

Theorem 1 *We can set up paths between pairs of perfect matchings of G_d such that if F_u, F_v are perfect matchings and $(F_z, F_{z'})$ is a move on the path from F_u to F_v then,*

1. *$F' = F_u \Delta F_v \Delta F_z$ is almost a perfect matching of G_d, i.e it maps to a matrix with less than 3 entries that are not 0 and 1, and these "odd" entries are each 2 or -1.*

2. *Given F' and $(F_z, F_{z'})$ we can uniquely determine F_u and F_v.*

Proof. Fix an order on all cycles in G_d, with cycles that don't have any edges in G_{xy} appearing before the remaining cycles. Also let each cycle have a start vertex.

For two perfect matchings F_u, F_v, $F_u \Delta F_v$ consists of vertex-disjoint cycles $C_1, \ldots, C_k$ in the sequence fixed by the global order. The paths we define are based on those in the proof of lemma 3. The subset of edges of each C_i in the graph G_{xy} maps to an alternating cycle in $G_u \Delta G_v$.

First we erase all cycles that can be erased with only short moves. These cycles have no edge in G_{xy}. Let C_i be the first of the remaining cycles. It maps to a cycle in $G_u \Delta G_v$ and the proof of lemma 3 gives us a sequence of switches that erase it. These switches correspond to long moves. We extend this sequence of moves to erase C_i from $F_u \Delta F_v$ in the following manner: When we make a long move we may create upto 4 cycles that can be erased using only short moves. We perform these short moves immediately following the long move. In this manner we get a series of short and long moves that have the net result of erasing C_i.

To prove the first part of the theorem we observe that only the long moves affect the edges of any $C_i \cap G_{xy}$ and since we are using the same sequence of long moves, it follows from lemma 3 that F' maps to an element of M'.

Now we prove the second part. Given F' and $(F_z, F_{z'})$, we can calculate $F_u \Delta F_v = F' \Delta F_z$. This gives us the sequence of cycles. Then the move $(F_z, F_{z'})$ also tells us which cycle we are currently working on. Let this be C_i. To determine exactly how far we have proceeded in this cycle, we apply the erasure procedure to the cycle till we reach the given move. It is a simple observation that any single move is used no more than once in erasing any particular cycle. To get F_u we undo the cycles $C_1, \ldots, C_{i-1}$ and the moves upto the current move of C_i in F_z, and to get F_v we apply the remaining moves of C_i and switch the cycles $C_{i+1}, \ldots, C_k$ to F_z. □

Corollary 2 *The conductance of the Markov chain is at least* $\frac{|M|}{2|M'|}$.

Proof. The total number of paths going through a single move $(F_z, F_{z'})$ is at most the number of graphs of type F'. $|F'| = N|M'|$, and using that N edges of F map to a single edge of M, the total number of paths through an edge of the M is at most $N^2|M'|$. Now we send a flow of $\frac{1}{N^2}$ through each path between every pair of states of M. Hence we send a flow of 1 unit between every pair of states. The flow through a single edge, i.e. the congestion is less than $\frac{1}{N^2}.N^2|M'| = |M'|$.

Let S be a subset of states such that $|S| \leq |M|/2$. Let $(S, M-S)$ denote the set of edges from S to $M-S$. The conductance of M is the minimum $\frac{|(S,M-S)|}{|S|}$ over all such subsets S. For any subset S, there are $|S||M-S|$ pairs of vertices that send a unit of flow across the cut $(S, M-S)$. Since the maximum flow through any edge is $|M'|$, the number of edges in the cut is at least $\frac{|S||M-S|}{|M'|} \geq \frac{|S||M|}{2|M'|}$. This implies that the conductance is at least as $\frac{|M|}{2|M'|}$, as required. □

Hence the chain is rapid-mixing if the ratio $\frac{|M'|}{|M|}$ is bounded by a polynomial in m, n.

3.1 Generating regular bipartite graphs.

A degree sequence is called *regular* if all the integers in the sequence are equal. In this section we show that the Markov chain M is rapid-mixing for regular sequences.

Let M_i be the set of matrices with exactly i entries equal to 2 or -1, and no more than one such non-0,1 entry in any row or column.

Theorem 3 *For a regular bipartite degree sequence* $(k, k, \ldots, k)$ *of length* $2n$, $|M_i| \leq (nk)^{2i}|M|$.

Proof. By induction on i, the number of odd entries. Base: $|M_0| = |M|$. Assuming the theorem is true for M_{i-1}, we shall prove it for M_i.

Assume w.l.o.g. that the entry (1,1) is odd, specifically let it be 2. All but $k-2$ other entries in the first row are zero. Similarly all but $k-2$ entries in the first column are zero. Let j be a column with a zero in the first row, i.e. $(1,j)$ is zero. Since this column sum is k, and it can have at most one odd entry, it has 1's in atleast $k-1$ rows and therefore there must be a row i such that $(i,1)$ is zero and (i,j) is 1. Now switching the cycle corresponding to the 4 four entries $(1,1),(1,j),(i,j),(i,1)$ gives a matrix in the class M_{i-1}.

Similarly suppose that the $(1,1)$ entry is -1. So $k+1$ entries in the first row are 1 and the remaining $n-k-2$ entries are zero. Let j be a column with a 1 in the first row. Since the column sum is k and it can have at most one odd entry, there are atleast $n-k-1$ zeros in this column. One of these zeros must occur in a position (i,j) such that $(i,1)$ is 1. The cycle corresponding to the entries $(1,1),(1,j),(i,j),(i,1)$ can be switched to get a matrix in the class M_{i-1}.

So each matrix in the class M_i is one switch away from some matrix in the class M_{i-1}. Since the total number of possible switches from a matrix in M_{i-1} is less than$(nk)^2$, we have $|M_i| = (nk)^2|M_i|$ and the theorem follows. □

Corollary 4 *For a regular bipartite degree sequence $(k,k,\ldots,k)$ of length $2n$, the conductance of the Markov chain is $\Omega(1/poly(n))$.*

Proof. From lemma 3 we have that for a bipartite sequence $|M'| \le |M_3|$. For a regular sequence $(k,k,\ldots,k)$ of length $2n$ we have that $|M_3| = O((nk)^6)|M|$ from the previous theorem. Using these two facts and corollary 2, the conductance of the chain is $\Omega(1/poly(n))$. □

4 The Markov chain on tournaments

In view of good existing surveys on results on tournaments, we shall keep the discussion on general facts about tournaments brief, and recommend that the reader refer to [16] or [17]. Let T_n denote a labeled tournament on n vertices. Let $S(T_n)$ denote the score sequence of T_n. Given an arbitrary sequence of n nonnegative integers, it is easy to determine if the sequence is in fact a score sequence using the following theorem, first proved by Landau [11].

Theorem 5 (Landau '53) *Let $0 \le s_1 \le \ldots \le s_n$ be a sequence of integers. Then $\mathcal{T}(s_1,\ldots,s_n)$ is nonempty if and only if for $i = 1,2,\ldots,n$,*

$$s_1 + s_2 + \cdots + s_i \ge \binom{i}{2},$$

with equality for $i = n$.

There is a constructive version of this theorem (see e.g. [16] or [17]), which gives a linear time algorithm to produce a tournament with the given score sequence. It follows from an elegant theorem of Ryser [18] on (0,1)-matrices that any tournament can be transformed into any other tournament with the same score sequence by successively reversing the arcs in 3-cycles (see also [6], and for a direct proof, see [3]).

Given all this, the irreducible Markov chain $M(\mathbf{s})$, whose state space $\Omega = \mathcal{T}(\mathbf{s})$ is the set of all labeled tournaments on n vertices with the score sequence $\mathbf{s} = (s_1,\ldots,s_n)$, suggests itself:

1. Start with an arbitrary T_n such that $S(T_n) = \mathbf{s}$.
2. (Δ-reversal). At each step, choose a 3-cycle uniformly at random (u.a.r.), and reverse the orientation of the arcs in the cycle with probability $1/2$; with the remaining probability do nothing.

An interesting fact (and an easy exercise in undergraduate combinatorics) is that the number of 3-cycles in a tournament is purely a function of its score sequence. More precisely, the number of 3-cycles in T_n with score sequence $(s_1, \ldots, s_n)$ equals $\binom{n}{3} - \sum_i \binom{s_i}{2}$. This shows that the stationary distribution of the Markov chain as defined above is uniform.

We remark that the steps of the above Markov chain are easy to implement. For, one can start with a list of all the 3-cycles, and at each step of the Markov chain, one needs to update the list which is easy to do. An alternative approach would be to use a slightly different chain as follows. At each step, pick a triple of vertices u.a.r and if the triple forms a 3-cycle, then apply a Δ-reversal (i.e. reverse the orientation of the 3-cycle), else do nothing. It is easy to check that this also gives a Markov chain with uniform stationary distribution. Moreover, the following analysis of rapid mixing remains essentially the same for this new chain.

4.1 The canonical paths.

Let T and T' be two arbitrary tournaments from $\mathcal{T}(\mathbf{s})$. Let $T \setminus T'$ denote the tournament with the same set of vertices, and whose arcs are those of T which are not in T'. It is easy to see that $T \setminus T'$ is an arc disjoint collection of directed cycles. Since Section 2.1 describes in detail how to deal with a *collection* of edge disjoint cycles in general, here we are only going to describe how to "take care of" a particular cycle C of length $3 \leq k \leq n$ in $T \setminus T'$. For this reason, in fact, let us assume that $T \setminus T'$ is a simple directed cycle, C. We are going to apply a particular sequence of Δ-reversals so that either T is transformed into T', or $T \setminus T'$ is reduced by at least one arc. In the latter case, we apply the procedure repeatedly. (Note that $T' \setminus T$ is simply the set of arcs obtained by reversing the arcs in $T \setminus T'$.)

Let $C = (w_1, \ldots, w_k)$, where the arcs $(w_1, w_2), \ldots, (w_k, w_1)$ are in T and $(w_1, w_k), \ldots, (w_2, w_1)$ are in T'. Further let w_1 be the start vertex of C (see Section 2.1). Set j to be $\min\{i : (w_i, w_1) \in T\}$. Then we can reverse the orientation of $(w_1, w_2), \ldots, (w_{j-1}, w_j)$ by applying a sequence of Δ-reversals on triples $(w_j, w_1, w_{j-1}), \ldots, (w_3, w_1, w_2)$, in that order. If $j = k$, then at the end of this sequence T is transformed into T', and if $3 \leq j < k$ then we are left with a smaller cycle, $C' = (w_1, w_j, w_{j+1}, \ldots, w_k)$. This finishes the description of canonical paths between every pair of states in $M(\mathbf{s})$.

4.2 Near-regular tournaments and proof of rapid mixing.

The proof of rapid mixing is similar to the proof (in Section 3) for the previous problem; so here we only outline the key differences. To get a bound on the *congestion* of the canonical paths defined above, it is convenient to consider the (0,1)-matrix representation of a tournament. Let the vertices be labeled, $1, 2, \ldots, n$. Given T_n, let $[T_n]$ represent the $n \times n$ matrix, where $[T_n]_{ij} = 1$ if arc $(i, j) \in T_n$, and 0 otherwise. Further, $[T_n]_{ii} = 0$, for all i. Thus $[T_n]$ is skew-symmetric (combinatorially speaking), where $[T_n]_{ij} + [T_n]_{ji} = 1$, whenever $i \neq j$. Let the score vector of T_n

be $(s_1, \ldots, s_n)$. Then the vector of row sums of $[T_n]$ is $(s_1, \ldots, s_n)$, and the vector of column sums is $(n-1-s_1, \ldots, n-1-s_n)$. Now, given a score sequence $\mathbf{s} = (s_1, \ldots, s_n)$, let $\mathbf{d}(\mathbf{s})$ denote the sequence $(s_1, s_2, \ldots, s_n, n-1-s_1, n-1-s_2, \ldots, n-1-s_n)$. Treating $[T_n]$ as a bipartite graph with the degree sequence $\mathbf{d}(\mathbf{s})$, we can construct the graph $G_{\mathbf{d}(\mathbf{s})}$, as in Section 3. Then the following lemma is an immediate corollary of Lemma 4.

Lemma 5 *A perfect matching in the graph $G_{\mathbf{d}(\mathbf{s})}$ maps to a tournament in $\mathcal{T}(\mathbf{s})$; moreover, there are exactly $N = \prod(n-s_i)! \prod(s_i+1)!$ perfect matchings mapping to a single tournament.*

The analog of Theorem 3.1 of Section 3.1 can now be derived from the following result of [13] (see also [21]). The results in [13] are actually more precise than as stated below. Given a score sequence $\mathbf{s} = (s_1, \ldots, s_n)$, let us define an *excess sequence* $\delta(\mathbf{s}) = (\delta_1, \ldots, \delta_n)$, where $\delta_j = s_j - [(n-1) - s_j] = 2s_j - (n-1)$, for $j = 1, \ldots, n$. Let $NT(n; \delta)$ denote the number of labeled tournaments on n vertices with the excess sequence $\delta = (\delta_1, \ldots, \delta_n)$. Thus $NT(n; (0, \cdots, 0))$ corresponds to $RT(n)$, the number of labeled regular tournaments. To avoid trivialities, assume in the following theorem that $\delta_1, \ldots, \delta_n$ are integers, of opposite parity to n, that sum to zero.

Theorem 6 (McKay-Wang) . *For $\epsilon > 0$ sufficiently small, suppose* $\max_j\{|\delta_j|\} = O(n^{3/4+\epsilon})$. *Then*

$$NT(n; \delta) = RT(n) \exp\left((-\frac{1}{2} + o(1)) \frac{\sum_{j=1}^{n} \delta_j^2}{n} \right).$$

Spencer was the first to obtain a result of the type above with $\max_j\{|\delta_j|\} = O(n^{3/5})$ (see Theorems 1 and 6 in [21]). We call $(s_1, \ldots, s_n)$ a *near-regular* score sequence if for all i, we have $|s_i - (n-1)/2| = o(n^{3/4})$.

Theorem 7 *For a near-regular score sequence* $\mathbf{s}$ *on n vertices, the conductance of the Markov chain $M(\mathbf{s})$ is $\Omega(1/poly(n))$.*

Proof sketch. Let S denote the set of score sequences that can be obtained from $\mathbf{s}$ by choosing two (not necessarily distinct) entries of $\mathbf{s}$, and increasing one score and decreasing the other score. Clearly, $|S|$ is at most n^2. Let $\mathcal{T}'(\mathbf{s})$ denote the set of (labeled) tournaments whose score sequence is in S. Now proceeding as in Corollary 3.1 of Section 3.1, it is straight forward to show that the conductance of the Markov chain $M(\mathbf{s})$ is $\Omega(|\mathcal{T}(\mathbf{s})|/|\mathcal{T}'(\mathbf{s})|)$. Using Theorem 4.2 we can arrive at the requisite bound:

$$\frac{|\mathcal{T}'(\mathbf{s})|}{|\mathcal{T}(\mathbf{s})|} \leq n^2 \exp\left((-\frac{1}{2} + o(1)) \frac{2 + o(n^{3/4})}{n} \right).$$

$\square$

5 Conclusion

Using some of the ideas from the end of [13], we can also efficiently generate a tournament almost uniformly at random from among all possible labeled tournaments on n vertices. While this problem is easy for labelled tournaments (pick an orientation of each edge at random), the technique might apply for generating unlabelled tournaments. The basic idea is that given a positive integer n, we can first select a score sequence almost uniformly at random, by selecting each s_j with the binomial distribution $B(n-1, 1/2)$, independently subject only to have the sum (of s_j's) equal to $n(n-1)/2$. From Theorems 4.4 and 4.5 of [13], it follows that the chosen (random) score sequence is, in fact, a near-regular score sequence with high probability (i.e. the failure probability is inverse polynomial in n). So, if the selected s_j's are not a genuine score sequence, then simply repeat the selection process. Now run the Markov chain of Section 4 to generate a random tournament which has the random score sequence.

Brualdi and Qiao showed that the diameter of the interchange graph of labeled tournaments is at most $(n-1)(n-2)/2$, and can be as large as $(n-1)^2/4$ – the two labeled tournaments they use for the lower bound are actually isomorphic, as noted in [8]. Recently, Guiduli et al [8] proved several interesting results about the interchange graph of *unlabeled* tournaments, wherein the vertices are non-isomorphic tournaments with a given score sequence, and (once again) two tournaments are adjacent, iff one can be obtained from the other by a Δ-reversal. They showed that the diameter of this graph can be as large as $(1-o(1))n^2/32$, through an explicit example of the Paley tournament on n vertices (for n prime, and congruent to 3 mod 4) and what they call a balanced 2-partition-transitive tournament, with the same (regular) score sequence. (Note that the upper bound in the labeled case implies an upper bound for the unlabeled case, and the lower bound of the unlabeled case implies a lower bound for the labeled case.) They also showed that one can find in $O(n^2)$ time a sequence of Δ- reversals which would convert a given tournament T to another (given) tournament T' with the same score sequence. The above lower bound(s) on the diameter are obviously lower bound(s) on the rate of mixing of the chain we discussed in this paper; however, providing tighter bounds on the rate of mixing should be a challenging task.

The original theorem of Ryser [18] on tournaments shows that one can transform any tournament into any other with the same score sequence by reversing 4-cycles and 3-cycles. Thus one can define a new chain using $\Box$-reversals (reversal of the arcs in 4-cycles) as well as Δ-reversals. Since any $\Box$-reversal can be accomplished by two Δ-reverals, our results of rapid mixing extend to this new chain. It might be interesting to compare the exact mixing rate of this chain with the ones we analyzed in this work.

We conclude with open problems of determining whether the Markov chains proposed are rapidly mixing for general degree/score sequences. The other natural open questions are whether *exact* counting in the problems we mentioned is #P-hard. One piece of strong evidence we have in this direction is a recent result ([4]) which shows that the problem of determining the exact number of *contingency tables* with prescribed row and column sums is #P-complete, even in the $2 \times n$ case.

References

[1] E.A. Bender and E.R. Canfield, The asymptotic number of labeled graphs with given degree sequence, *J. Combin. Theory Ser. A*, **24** (1978), 296-307.

[2] B. Bollobás, A probabilistic proof of an asymptotic formula for the number of labeled graphs with given degree sequence, *European J. Combin.*, **1** (1980), 311-316.

[3] R.A. Brualdi and Qiao Li, The interchange graph of tournaments with the same score vector, in "Progress in Graph Theory," Academic Press, Canada (1984), 128-151.

[4] M. Dyer, R. Kannan, and J. Mount, Sampling contigency tables, preprint (October, 1995).

[5] A. Frieze, On random regular graphs with non-constant degree, Research report 88-2, Department of Mathematics, Carnegie Mellon University, 1988.

[6] D.R. Fulkerson, Upsets in round robin tournaments, *Canad. J. Math.*, **17** (1965), 957-969.

[7] D. Gale, A theorem on flows in networks, *Pacific J. Math.*, **7** (1957), 1073-1082.

[8] B. Guildi, A. Gyárfás, S. Thomassé, and P. Weidl, 2-partition-transitive tournaments, preprint (April, 1996).

[9] M. Jerrum and A. Sinclair, Fast uniform generation of regular graphs, *Theoretical Computer Science*,**73** (1990), 91-100.

[10] M. Jerrum and A. Sinclair, Polynomial-time approximation algorithms for the Ising model, *University of Edinborough, Internal Report*, **CSR 1-90** (1990).

[11] H.G. Landau, On dominance relations and the structure of animal societies, III: the condition for a score structure, *Bull. Math. Biophys.* **15** (1953), 143-148.

[12] B.D. McKay, The asymptotic number of regular tournaments, eulerian digraphs and eulerian oriented graphs, *Combinatorica*, **10** (1990), 367-377.

[13] B.D. McKay and X. Wang, Asymptotic enumeration of tournaments with a given score sequence, *J. Combin. Theory, Series A*, to appear (1996).

[14] B.D. McKay and N.C. Wormald, Asymptotic enumeration of by degree sequence of graphs of high degree, *European J. Combinatorics*, **11** (1990), 565-580.

[15] B.D. McKay and N.C. Wormald, Uniform generation of random regular graphs of moderate degree, *J. Algorithms*, **11** (1990), 52-67.

[16] J.W. Moon, "Topics on tournaments," Holt, Rinehart, and Winston, New York (1968).

[17] K.B. Reid and L.W. Beineke, Tournaments, Ch. 7 of *Selected Topics in Graph Theory* (ed. by L.W. Beineke and R.J. Wilson), Academic Press, New York (1978), 169-204.

[18] H.J. Ryser, Matrices of zeros and ones in combinatorial mathematics, *Recent advances in matrix theory* (ed. by H. Schneider), Univ. of Wisconsin Press, Madison (1964), 103-124.

[19] H.J. Ryser, "Combinatorial Mathematics," Carus Math. Monograph, No. 14, New York, Wiley (1963).

[20] A.J. Sinclair and M.R. Jerrum, Approximate counting, uniform generation and rapidly mixing Markov chains, *Inform. and Comput.*, **82** (1989), 93-133.

[21] J.H. Spencer, Random regular tournaments, *Period. Math. Hungar.*, **5** (1974), 105-120.

[22] N.C. Wormald, Generating random regular graphs, *J. Algorithms*, **5** (1984), 247-280.

A Appendix: A bottleneck

In this section we show a degree sequence $\mathbf{d}$ of length $2n$ for which the ratio $\frac{|M_1|}{|M|}$ is exponential in n. Although this does not imply that the chain is not rapid-mixing (as in the case of more complicated chains studied by Jerrum and Sinclair) it indicates the need for a more sophisticated analysis.

Let $\mathbf{d}$ be the sequence $(1, 2, \ldots, n, 1, 2, \ldots, n)$. The last row and column of a solution to this sequence must be all 1's. So,

$$\begin{aligned} |M(\mathbf{d})| &\leq |M(1, 2, \ldots, n-2, 1, 2, \ldots, n-2)| \\ &\leq |M(1, 2, \ldots, n-4, 1, 2, \ldots, n-4)| \end{aligned}$$

Continuing this way we see that $|M(\mathbf{d})| = 1$.

To estimate $|M_1|$, let us count the number of matrices with degree sequence $\mathbf{d}$ and the entry (n, n) equal to 2.

The next few steps generate solutions in M_1 by filling up rows and colums in descending order.

$$\begin{aligned} |M_1(\mathbf{d})| &\geq |M(1, 1, 2, 3, \ldots, n-2, 1, 1, 2, 3, \ldots, n-2)| \\ &\geq 2^2 |M(1, 1, 2, \ldots, n-4, 1, 1, 2, \ldots, n-4)| \\ &\geq 2^4 |M(1, 1, 2, \ldots, n-6, 1, 1, 2, \ldots, n-6)| \\ &\geq 2^{n-2} \end{aligned}$$

In each step above we get 4 different solutions by choosing one of the rows with row sum 1 and one of the columns with column sum 1 to be set at 0.

Since $\frac{|M_1|}{|M|} \geq 2^{n-2}$, we can not hope to prove rapid-mixing in general with the canonical paths used in this paper.