

A local switch Markov chain on given degree graphs with application in connectivity of Peer-to-Peer networks

Tomás Feder Adam Guetz Milena Mihail Amin Saberi

Abstract

We study a switch Markov chain on regular graphs, where switches are allowed only between links that are at distance 3; we call this the Flip. The motivation for studying the Flip Markov chain arises in the context of unstructured peer-to-peer networks, which constantly perform such flips in an effort to randomize. We show that the Flip Markov chain on regular graphs is rapidly mixing, thus justifying this widely used peer-to-peer networking practice. Our mixing argument uses the Markov chain comparison technique, however, we extend this technique to embedding arguments where the compared Markov chains are defined on different state spaces. We give several conditions which generalize our results beyond regular graphs.

1 Introduction

In this paper, we study natural Markov chains on the set of simple graphs with a fixed degree sequence. All the Markov chains that we study here are variations of the “switch”. The moves in this Markov chain remove a pair of edges (i, j) and (k, l) from the graph and replace it with (i, k) and (j, l) , if the resulting graph remains simple. This Markov chain is known to connect all the graphs with the same degree sequence[1] (for any pair of graphs G and G' with the same degree sequence, it is possible to transform G into G' by a sequence of switches), it is therefore a natural candidate scheme for generating a random such graph. Extending and simplifying pioneering work of Jerrum and Sinclair[8, 10], Kannan, Tetali, and Vempala [14] analyzed a restriction of this chain on bipartite graphs and showed that it is rapidly mixing (while Bezakova, Bhatnagar and Vigoda[2] gave a related simulated annealing based scheme for sampling bipartite graphs with a wide range of degree distributions). Cooper, Dyer, and Greenhill [3] used a better canonical path argument and extended the results of [14] to non-bipartite graphs under certain conditions for the degree sequence; their conditions include the fundamental case of regular graphs.

However, beyond generating random graphs with a given degree sequence, the above type of Markov chains have recently arose in the context of maintaining well connected topologies in unstructured peer-to-peer networks[3, 15, 5, 17]. In this context, the graph G represents the topology of a peer-to-peer network at a particular instance, and a switch of (i, j) and (k, l) with (i, k) and (j, l) represents dropping and adding a few links from the network, thus slightly changing its topology. We have implemented a Gnutella

client and have measured, in practice, between 5 and 30 requests per second to each client for such additions and deletions; approximately 1% of these requests are satisfied. Given that current peer-to-peer networks, like Gnutella, KaZaa or eMule, have in the order of 1-5 million clients, this indicates a dramatic amount of constant change in the network topology. Peer-to-peer networks are apparently constantly trying to change their topology in an effort to randomize, and thus maintain the excellent connectivity properties of random graphs (such as low diameter and expansion). In this sense, the work of [3] was a first indication that the heuristics used by current peer-to-peer networks have some theoretical foundations.

However, in the context of peer-to-peer networks, modelling the changes in the network by a switch in the underlying graph $G(V, E)$, $|V| = n$, has a serious drawback. In particular, the assumption that two links can be picked uniformly at random violates the strong local nature of peer-to-peer networks. In such networks, each node has memory $O(\log n)$ (enough to remember its address), and can be assumed to have computational resources $O(\text{poly } \log n)$. Thus each node “knows” a neighborhood of size at most $O(\text{poly } \log n)$ around itself. How can, under these circumstances, the network pick two links uniformly at random? In addition, a switch operation can potentially make the graph disconnected. Without a central authority, it seems almost impossible to re-connect the graph.

In this paper we overcome the above problem by looking at a variation of the switch chain which can be implemented with local information about the network. In a restricted version of the switch which we call a “Flip”, the edges (i, j) and (k, l) will be replaced by (i, k) and (j, l) only if i and l are adjacent to each other. Another way to see this is the endpoints of the edge (i, l) exchange a random neighbor with each other. The exchange happens only if the resulting graph is still simple. This operation is fairly common in different variations of peer-to-peer networks and is part of standard java codes for peer-to-peer operations [17]. Another important property of this switch is that, if applied to a connected graph, it keeps the graph connected. Thus, when starting from a connected graph, the Flip operation defines a Markov chain on connected graphs. Moreover, in [15] it was shown that this Markov chain connects the state space of all connected graphs, i.e. it is possible to transform a connected graph G to a connected graph G' with the same degree sequence, using only Flip operations.

The main contribution of this paper is to show that the Flip Markov chain is rapidly mixing for regular graphs. In order to do that we first consider the switch chain induced on the set of connected graphs with a given degree sequence. Let us call this chain the restricted switch. We bound the mixing time of the Flip chain in terms of the mixing time of the restricted switch by a type of Markov chain comparison argument. It is easy to see that the nodes of the underlying graph of the Flip chain are the same as the restricted switch, but the graph underlying the chain of the restricted switch may have many more edges than the graph underlying the Flip. The technical difficulty of our argument is to show that it is possible to simulate and map every edge of the restricted switch chain to a path in the Flip chain between the same pair of vertices such that the the “congestion” of every edge, i.e. the number of times it is picked in all paths, is bounded.

Then we compare the restricted switch with the switch Markov chain. We show that there is a way to embed the transition graph of the switch chain in the transition graph of the restricted switch chain such that, at most $2n$ nodes of the switch chain map to a node of the restricted switch chain, and if two nodes are adjacent in the switch chain then their distance will be bounded by a constant in the restricted switch. It is also easy to see that degree of a node in both chains is polynomial in n . Therefore it is possible to simulate every transition in the switch chain with a series of transitions in the restricted switch such that the congestion of every edge is bounded by a polynomial. Using the above two arguments and the results of [3], we conclude that the Flip Markov chain is rapidly mixing for regular graphs. To the best of our knowledge, the embedding argument that we use to bound the mixing time of restricted switch Markov chain in terms of the mixing time of the switch Markov chain is the first Markov chain comparison argument where the compared Markov chains are not on the same state space.

We should note that for non-regular graphs, the Flip chain is not always irreducible (does not connect the state space of all connected graphs with a given degree sequence). We find a general sufficient condition for degree sequences for which the Flip chain is irreducible and ergodic. Using this result and a generalization of the canonical path argument given in [3], we can prove the mixing of the switch Markov chain and the Flip Markov chain for a wide range of degree distributions. We note, however, that, in the context of peer-to-peer network applications which mainly motivate the study of the Flip Markov chain, the most interesting case is that of regular graphs (unstructured peer-to-peer networks as are known to typically have degrees between 5 and 30).

The rest of this paper is organized as follows. In Section 2 we give definitions and outline of results. In Section 3 we show that the Flip Markov chain is rapidly mixing for regular graphs. In Section 3.2 we bound the mixing time of the Flip markov chain in terms of the mixing time of the switch Markov chain restricted to connected graphs. This is a Markov chain comparison argument. In Section 3.1 we bound the mixing time of the restricted switch Markov chain in terms of the mixing time of the general unrestricted switch Markov chain. This is a generalized Markov chain comparison argument, in the sense that the state spaces of the two Markov chains are not the same. In Section 4 we characterize degree sequences for which the Flip Markov chain mixes rapidly, and we extend the allowed moves of the Flip Markov chain (always preserving locality of information as motivated by the peer-to-peer networking application) so that the state space indeed covers all connected realizations of a given degree sequence.

2 Definitions and discussion of related results

In this paper we will be studying properties of three different related Markov chains on graphs.

Switch Markov chain: For a given graphical degree sequence (d_1, \dots, d_n) , let Ω_S be the set of graphs satisfying this degree sequence. We define a Markov chain \mathcal{M}_S on Ω_S as follows. From $G \in \Omega$, with probability $\frac{1}{2}$ do nothing. Otherwise choose two distinct edges $(i, j), (k, l)$ uniformly at random. Then, choose a perfect matching M of $\{i, j, k, l\}$

uniformly at random, and if $M \cap E(G) = \emptyset$ then delete the edges $(i, j), (k, l)$ from G and add the edges of M to G . Otherwise do nothing. This operation is called a switch. We write \mathcal{G}_S for the underlying graph of \mathcal{M}_S , so $\mathcal{G}_S = (\Omega_S, T)$ where each edge $e \in T$ corresponds to a transition of \mathcal{M}_S . It is known that \mathcal{M}_S is irreducible [1] and hence ergodic, converging to the uniform stationary distribution over all connected graphs. [3] show that \mathcal{M}_S is rapidly mixing for degree sequences corresponding to regular graphs (and for several more general cases of degree sequences).

Switch Markov chain restricted to connected graphs: The Markov chain \mathcal{M}_{SC} is defined identically to \mathcal{M}_S , except that a switch is not accepted if it disconnects the graph. Its underlying graph \mathcal{G}_{SC} is an induced subgraph of \mathcal{G}_S restricted to connected graphs. It has been shown that \mathcal{G}_{SC} is connected [22], and therefore \mathcal{M}_{SC} is irreducible and ergodic. \mathcal{M}_{SC} is also time-reversible with uniform stationary distribution. \mathcal{M}_{SC} has been previously studied experimentally in [6, 19] in the context of generating connected realizations of degree sequences representing internet topologies. However, it was not known if \mathcal{M}_{SC} is rapidly mixing.

Flip Markov chain: The flip Markov chain \mathcal{M}_F is a chain on the set of connected graphs in which a switch between edges $(i, j), (k, l)$ is allowed only if i and l are adjacent. Let \mathcal{G}_F be the underlying graph of \mathcal{M}_F . It is noted in [15] that \mathcal{M}_F for regular graphs is ergodic and time-reversible with uniform stationary distribution.

The main contribution of the paper is to prove that the Flip Markov chain is rapidly mixing for regular graphs using Markov chain comparison techniques (For definitions and an excellent survey of these techniques, see [9, 4]). We will use the \mathcal{M}_{SC} chain as an intermediate step to compare \mathcal{M}_S and \mathcal{M}_F . We will use the known rapid mixing of \mathcal{M}_S [3] to show that \mathcal{M}_{SC} is rapidly mixing, and we will then use the rapid mixing of \mathcal{M}_{SC} to show that \mathcal{M}_F is rapidly mixing.

3 The Flip Markov Chain

In order to bound the mixing time of \mathcal{M}_F in terms of \mathcal{M}_{SC} we do the following: for every edge (u, v) in \mathcal{G}_{SC} we choose a path between u and v in \mathcal{G}_F such that the ‘‘congestion’’ of every edge, i.e. the number of times it is picked in all paths, is bounded. We will use that to prove that the mixing times of the \mathcal{M}_F and \mathcal{M}_{SC} are polynomially related.

We then compare the restricted switch markov chain \mathcal{M}_{SC} with the switch Markov chain \mathcal{M}_S . We show that there is a function h that embeds \mathcal{G}_{SC} into \mathcal{G}_S such that $h(\mathcal{G}_{SC}) = \mathcal{G}_{SC}$, each graph in \mathcal{G}_{SC} is the image of at most $2n$ graphs in \mathcal{G}_S , and that if two graphs are adjacent in \mathcal{G}_S the distance of their images remains bounded by a constant factor in \mathcal{G}_{SC} . It is also easy to see that degree of a node in both graphs is polynomial. Therefore it is possible to simulate every transition in the switch Markov chain with a series of transitions in the restricted switch Markov chain such that the congestion of every edge is bounded by a polynomial. Using the above two arguments and the work of Cooper *et al.*, we conclude that the Flip Markov chain is rapidly mixing for regular graphs.

In the first part of the section, we bound the mixing time of the switch restricted to the




Figure 1: Mapping $h : \mathcal{G}_S \mapsto \mathcal{G}_{SC}$

set of connected graphs. We then use this result to prove that the Flip Markov chain is mixing.

3.1 Rapid mixing of \mathcal{M}_{SC}

We now give an explicit construction that maps every edge in \mathcal{G}_S to a path of bounded length in \mathcal{G}_{SC} , while no more than $2n$ graphs in \mathcal{G}_S map to a single graph in \mathcal{G}_{SC} .

Lemma 1 *For a graphical degree sequence d , if $d_i \geq 2$ for every i , there exists a mapping $h : \Omega_S \mapsto \Omega_{SC}$ such that $h(G) = G$ and $|h^{-1}(G) \leq 2n|$ for $G \in \Omega_{SC}$, and if G, \bar{G} are adjacent in \mathcal{G}_S , then $h(G), h(\bar{G})$ are at distance at most 17 in \mathcal{G}_{SC} .*

Proof. Let k be the number of connected components of G . Impose an arbitrary numbering on the set of vertices and sort the connected components H_i of $G \in \Omega_S$ in ascending order of their highest numbered vertex. For each component H_i , let $e(H_i) = (u_i, v_i)$ where e is an edge on a cycle in H_i such that the pair (u_i, v_i) is lexicographically highest. Since all degrees are greater than 2, each component must have such an $e(H_i)$, which we call a *bridge edge*.

Now for mapping a graph G to $h(G)$, we remove each $e(H_i)$ and connect the components together by adding (v_i, u_{i+1}) for $1 \leq i \leq k - 1$ plus (v_k, u_1) .

Suppose G and \bar{G} are adjacent in \mathcal{G}_S . Let $G' = h(G)$ and $\bar{G}' = h(\bar{G})$. We will show that the distance between G' and \bar{G}' in \mathcal{G}_{SC} is at most 17. First, we need to define two useful transformations.

1. An *edge rotation* a series of switches within a component to change its bridge edge. We do this as follows: choose an edge $e \in \bar{G}$ that is incident to neither the vertices of $e(H_i)$ nor the vertices of $e(\bar{H}_i)$. If \bar{G} has more than one component such an edge must exist, otherwise \bar{G} and G are adjacent in \mathcal{G}_{SC} anyways.

If $e(\bar{H}_i)$ shares a vertex with $e(H_i)$, we can complete the rotation of the edge with just one switch. Now, it might be necessary to change the direction of H_i by switching the edges (v_{i-1}, v_i) and (u_i, u_{i+1}) to (v_{i-1}, u_i) and (v_i, u_{i+1}) .




Figure 2: edge rotation

If $e(\bar{H}_i)$ does not share a vertex with $e(H_i)$, we can perform a switch to remove $e(H_i)$ and e and another two switches to remove the edges from the vertices of $e(\bar{H}_i)$ linking to H_i' 's neighboring components. Then we perform one more switch to restore $e(\bar{H}_i)$ and e , resulting in the desired graph. An edge rotation therefore requires at most 4 switches.

2. A *component reordering* changes the order of a component H_i in the chain and moves it to a position j . First we switch the edge connecting H_i to H_{i+1} and the one connecting H_j to H_{j+1} . The second switch switches the edge connecting H_{i+1} and H_{j+1} and the edge connecting H_{i-1} and H_i . Now, It may be necessary to change the direction of H_i itself with one switch described in the edge rotation, for a total of at most 3 switches.

Now, let us name the edges in $E(G) - E(\bar{G})$ as e, f and $E(\bar{G}) - E(G)$ as e' and f' . Consider the following cases for the switch from G to \bar{G}

Case 1: *switch is within the component:* e, f, e', f' are chosen within a single component H_i , and the switch does not disconnect H_i .

The component orderings in G and \bar{G} are the same, as their corresponding vertex sets do not change with this switch. To create a path from $h(\bar{G})$ to $h(G)$, we may need to first perform an edge rotation to an edge $e'' \neq \{e', f'\}$ if either e' or f' is the bridge edge. Then we can safely switch e', f' with e, f without disconnecting the graph, and then rotate again to restore $e(H_i)$ as the link. This takes at most $4 + 1 + 4 = 9$ switches.

Case 2: *Switch disconnects component:* e, f, e', f' are chosen within a single component H_i , and the switch disconnects H_i into new components R_1, R_2 .

Note that a single switch can split a component into at most two components. Assume that R_1 has the higher numbered vertex. Then R_1 is in the correct position in \bar{G}' , R_2 may not be. If either component is using e' or f' as a bridge, we now perform an edge rotation to a different edge within the component.

We now reorder R_2 to be adjacent to R_1 . This only takes two switches, since we're not concerned about direction at this point. We may need to change direction however, of one of R_1 or R_2 to ensure that the link between them is neither e' nor f'' . Now, switch by removing the edge linking R_1 and R_2 and the edge linking R_1 and H_{i-1}' , adding an

edge linking R_2 and H'_{i-1} and restoring R_1 's missing bridge edge. Next, we remove e' and f' and add e and f . All that's left is to do a final edge rotation, and the path in \mathcal{G}_{SC} is complete. This takes at most $4 + 4 + 3 + 1 + 1 + 4 = 17$ switches.

Case 3: *Switch merges two components:* e, f are chosen from different components H_i, H_j , e', f' are such that $H_i \cup H_j \cup \{e', f'\} - \{e, f\}$ is connected.

Note that one of e, f may disconnect their respective components in this case, but not both. The merged component is the correct order for one of the components, say i . We apply a rotation so that an edge of H'_i is the bridge edge. We then apply a switch, removing an edge of H'_j and the edge connecting $(H'_{j-1} \text{ to } H_{j+1})'$ and adding edges so that H'_j is in the correct order. We can then apply another switch to remove e', f' and add e, f . Finally, we apply an edge rotation to each of H'_i and H'_j to complete the path. This yields at most $4 + 1 + 1 + 4 + 4 = 14$ switches.

Case 4: *Switch exchanges parts between two components:* e, f chosen such that e disconnects H_i into R_1, R_2 and f disconnects H_j into R_3, R_4 , with e' connecting R_1, R_3 and f' connecting R_2, R_4 .

Note that none of e, f, e', f' may be on a cycle in \bar{G} or G . This implies that one of the R 's in \bar{G}' is in the correct position and is using the correct bridge edge, say R_1 . We now apply an edge rotation to the R_2, R_4 component such that an edge of R_4 is being used. Now, we can apply a switch adding e, f and removing e', f' without disconnecting the graph. A edge rotation of R_3, R_4 to the proper bridge edge $e(H_j)$ and then a reordering of H'_j completes the path. This gives at most $4 + 1 + 4 + 3 = 12$ switches.

Given a graph $h(G) \in \Omega_{SC}$, it is possible to recover G by traversing the cycles in $h(G)$ starting from the highest-numbered vertex v of H_1 and traversing the components in one of the two directions, and recover all the H_i 's by recognizing when a vertex larger than the current highest numbered vertex of an H_i is found. For the H_i 's, the edges in $h(G)$ joining graphs H_i, H_{i+1} or H_r, H_1 are the sets of edges in a cycle in $h(G)$ that participate in a cut of size two and are incident to a vertex of degree at least 3. Since there are two directions and n choices for v , the pre-images $h(G)$ consist of at most $2n$ graphs. \square

Now, it is easy to see that \mathcal{G}_S has maximum degree $(dn)^2$, since there are at most dn edges and we pick two edges for a switch. Therefore, each edge in \mathcal{G}_S can be picked in at most $16(dn)^{16}$ paths. The lengths of paths are also bounded by n^2 . $\ln(\Omega_S) \leq (nd)^2$. Combining these facts with Lemma 1 implies the following.

Theorem 2 *The mixing time of the switch Markov chain M_{SC} induced on the set of connected graphs Ω_{SC} is*

$$\tau_{SC}(\epsilon) = O(d^{34} n^{36} \tau_S(\epsilon)).$$

So if M_S is rapidly mixing for a given degree sequence, M_{SC} is also rapidly mixing.

3.2 Rapid mixing of the Flip chain

In this section, we bound the mixing time of $\mathcal{M}_{\mathcal{F}}$ in terms of the mixing time of \mathcal{M}_{SC} for regular graphs. For regular graphs, $\mathcal{M}_{\mathcal{F}}$ is ergodic, time reversible, and has uniform stationary distribution [15].

Lemma 3 *For regular degree sequences with $d \geq 2$, every edge (G, G') in \mathcal{G}_{SC} can be simulated with a path P from G to G' in $\mathcal{G}_{\mathcal{F}}$ whose intermediate graphs G_i differ in at most 12 edges from G or from G' .*

Proof. If an edge (G, G') of \mathcal{G}_{SC} is given by e, f, e', f' , then $G - \{e, f\}$ must have either a path joining x, z or a path joining y, t . Let P be the shortest such path, of length p , say joining x, z . There are several cases. If (y, P, t) is an induced path, of length $p + 2$, then we proceed by reversing the order of the vertices in P by repeatedly exchanging consecutive vertices in $P = (0, \dots, p)$, obtaining at intermediate stages $(j, j + 1, \dots, j + \ell, j - 1, j + \ell + 1, \dots, p, j - 1, \dots, 0)$ which are 5 edges extra and 5 deficit compared to G , and with 4 edges extra and 4 deficit compared to G' .

If (y, P, t) is an induced path plus the edges $(y, 1)$ and $(p - 1, t)$ with $p \geq 3$, then proceed as in the previous case for the path $(y, 1, \dots, p - 1, t)$, with 4 edges extra and 4 deficit compared to G , and with 6 edges extra and 6 deficit compared to G' , then finish by switching edges $(y, 0), (t, 1)$, and switching edges $(y, p - 1), (t, p)$.

If (y, P, t) is an induced path plus edges joining y to some of the vertices $1, \dots, p - 1$, the last of these being i , proceed with the induced path $y, i, i + 1, \dots, p, t$ as in the first case, then proceed inductively with the path $y, 0, 1, \dots, i, t$ plus extra edges joining y to the other vertices, with 5 edges extra and 5 deficit compared to both G and G' .

Finally, if x, y, z, t have a common neighbor v , then if $|d_i - d_j| \leq 1$ the vertex x must have a neighbor u not adjacent to v , switch $(x, u), (v, z)$ with $(x, z), (v, u)$, then switch $(x, y), (z, t)$ with $(x, t), (z, y)$, then switch $(x, z), (v, u)$ with $(x, u), (v, z)$, giving 4 edges extra and 4 deficit compared to both G and G' .

□

In $\mathcal{G}_{\mathcal{F}}$, an edge E belongs to a path corresponding to the edge (G, G') in \mathcal{G}_{SC} only if G has at most 5 edges extra and 5 deficit compared to the endpoints of E . There are at most $(dn/2)^5 8 \cdot 6 \cdot 4 \cdot 3$ such G , with $(dn)^2$ choices possible for G' as a neighbor of G , giving a factor $24d^7 n^7$ for $\mathcal{G}_{\mathcal{F}}$. An additional factor of $O(n^2)$ is incurred as the length of a simulation of a step of \mathcal{G}_{SC} in $\mathcal{G}_{\mathcal{F}}$ is $O(n^2)$. Therefore the congestion is bounded by $O(d^7 n^9)$.

Theorem 4 *The mixing time for the Markov chain $\mathcal{M}_{\mathcal{F}}$*

$$\tau_F(\epsilon) = O(d^7 n^9 \tau_{SC}(\epsilon)).$$



Figure 3: Bow-tie switch

4 Extensions to more general degree sequences

It is desirable to extend our results to more general degree sequences, for example, including power-law degree sequences observed in some of the known networks[6, 19] But is important to note that the Flip Markov chain is not necessarily reducible when the graph is not regular. The *bow-tie* graph in Figure(4) is a counter-example.

However, we can show that this anomaly is limited only to certain graphs. We show that if all the graphs with a certain degree sequence have diameter bigger than 3, then the Flip chain is irreducible and therefore ergodic, time reversible and with a uniform stationary distribution.

Furthermore, we extend the proof of [3] for the rapid mixing of \mathcal{M}_S for general degree sequences that encompass those of power-law networks. Our proof is mainly built upon the proof of [3] and we only sketch it here.

4.1 The Switch Markov chain for non-regular graphs

We use the same multicommodity flow approach as in [3]: given $G, G' \in \Omega$, let $H = G\Delta G'$ be the symmetric difference of G and G' . The paths p defined in [3] have the property that if p goes from G to G' , and if G'' is on p , then every edge of $G\Delta G'$ joins two vertices that have non-zero degree in $H = G\Delta G'$. Since the graphs $G \in \Omega$ can be made d -regular by adding additional vertices to obtain \hat{G} , where the additional edges involve at most one vertex in G , the paths in [3] can be considered paths for our case of degree sequences (d_1, \dots, d_n) , since $\hat{G}, \hat{G}', \hat{G}''$ will agree in edges incident to vertices not in G, G', G'' .

The paths are defined by choosing a *pairing* of edges in $G \setminus G'$ with edges in $G' \setminus G$ around each vertex. Let $\Phi(G, G')$ be the set of such pairings. For each pairing in $\Phi(G, G')$, a canonical path from G to G' is constructed. Each of these paths will carry $1/\Phi(G, G')$ of the total flow from G to G' .

We now obtain a bound on the mixing time. Fix a pairing $\phi \in \Phi(G, G')$ and let Z be any graph on the corresponding canonical path from G to G' . Identify each graph with its symmetric $n \times n$ adjacency matrix. Define a symmetric $n \times n$ matrix L by $L + Z = G + G'$. Entries of L belong to $\{-1, 0, 1, 2\}$. An edge in L is called *bad* if its label is -1 or 2 . We call L an *encoding* for Z (with respect to G, G'). Note that an edge receives label -1 if it is absent in both G and G' but it is present in Z , while an edge receives label 2 if it is present in both G and G' but is not present in Z . Thus, edges in the symmetric difference $G\Delta G'$ never receive bad labels.

Let Z' be the next graph after Z in the canonical path from G to G' . It is shown in [3] that given (Z, Z') , L and ϕ , we can uniquely recover G and G' . It is also shown in [3] that there are at most four bad edges in any encoding L , one of them with 2, two with -1 , and the remaining one with either -1 or 2.

Suppose we fix an edge $e = (Z, Z')$. We may choose an encoding L and uniquely recover G, G' such that $L + Z = G + G'$ from each possible ϕ belonging to the set $\Phi'(L)$ of possible pairings. We write $\mathcal{L}(Z)$ for the set of possible choices of encodings L . The pair G, G' gives rise to $|\Phi(G, G')|$ pairings. It is shown in [3] that $|\Phi'(L)| \leq d^6 |\Phi(G, G')|$. Therefore if we write $(e \in \gamma_\phi(G, G'))$ for the indicator variable which is 1 if e is in the path from G to G' corresponding to ϕ and 0 otherwise, we have

$$\begin{aligned} |\Omega|^2 f(e) &= \sum_{(G, G')} \sum_{\phi \in \Phi(G, G')} (e \in \gamma_\phi(G, G')) |\Phi(G, G')|^{-1} \\ &\leq \sum_{L \in \mathcal{L}(Z)} \sum_{\phi \in \Phi'(L)} d^6 |\Phi'(L)|^{-1} \\ &\leq \sum_{L \in \mathcal{L}(Z)} d^6 \\ &\leq d^6 |\mathcal{L}(Z)|. \end{aligned}$$

Let λ be the maximum possible value for $\mathcal{L}(Z)$. For any transition $e = (Z, Z')$ we have

$$1/Q(e) = |\Omega|/P(Z, Z') \leq d^2 n^2 |\Omega|.$$

Therefore

$$\rho(f) \leq d^8 n^2 \lambda / |\Omega|.$$

Also $\ell(f) \leq dn/2$, since each transition along a canonical path replaces an edge of G by an edge of G' in [3]. Since π is uniform we have

$$\log 1/\pi^* \leq dn \log(dn).$$

Then we have

Lemma 5

$$\tau(\epsilon) \leq \frac{d^9 n^3 \lambda}{2|\Omega|} (dn \log(dn) + \log(\epsilon^{-1})). \quad (1)$$

It thus remains to bound $\lambda/|\Omega|$. An encoding L has the same degree sum at each vertex as a graph $G \in \Omega$. If we replace $-1, 2$ with $0, 1$ respectively, we obtain instead of the degree sequence \mathbf{d} a sequence \mathbf{d}' such that $dist(\mathbf{d}, \mathbf{d}') = \sum |d_i - d'_i| \leq 8$. There are $O(n^8)$ choices of such \mathbf{d}' .

Explicit asymptotic formulas for $|\Omega_{\mathbf{d}}|$ have been obtained in several cases:

1. all d_i in \mathbf{d} are the same and equal to d [23, 7];

2. if m is the sum of d_i , and all d_i are $o(m^{1/4})$ [11];
3. if m is the sum of d_i , and all d_i are $o(m^{1/3})$ [13];
4. if the average d of the d_i satisfies $\min(d, n - d + 1) > cn/\log n$ and $|d_i - d| = O(N^{-1/2+\epsilon})$ for sufficiently small $\epsilon > 0$ and any $c > 2/3$ [12].

In case (1), it is shown in [3] that $\lambda/|\Omega| \leq 2d^6n^5$, giving

$$\tau(\epsilon) \leq d^{15}n^8(dn \log(dn) + \log(\epsilon^{-1})).$$

In cases (2,3), If \mathbf{d}' is obtained from \mathbf{d} by increasing one d_i by one, then $O(1) \leq |\Omega_{\mathbf{d}'}|/|\Omega_{\mathbf{d}}| = O(\sqrt{n})$. In case (4), If \mathbf{d}' is obtained from \mathbf{d} by increasing one d_i by one, then $O(1) \leq |\Omega_{\mathbf{d}'}|/|\Omega_{\mathbf{d}}| = O(\sqrt{\log n})$. Since $|\mathcal{Q}|/|\Omega_{\mathbf{d}'}| = O(n^8)$, and G', G differ at most by six additions of 1 (corresponding to the three -1) and two subtraction of 1 (corresponding to the two 2) we have:

Corollary 6 *In cases (2,3) $\lambda/|\Omega| \leq O(n^{11})$ with*

$$\tau(\epsilon) \leq O(d^9n^{14}(dn \log(dn) + \log(\epsilon^{-1}))),$$

and we have in cases (4) $\lambda/|\Omega| \leq O(n^8 \log^3 n)$ with

$$\tau(\epsilon) \leq O(d^9n^{11} \log^3 n(dn \log(dn) + \log(\epsilon^{-1}))).$$

Suppose all d_i satisfy $((d+1)/(n-d))d \leq d_i \leq d$ for some $d < n/2$. We first show that no valid graph has an independent set I of size $n-d$. Otherwise the vertices in I have degree at least $((d+1)/(n-d))d$, so at least $d(d+1)$ edges join I to its complement J of size d which has degrees at most d , giving at most d^2 edges joining I to J , a contradiction. Thus at least $d/2$ edges must be in I .

Suppose $r = \text{dist}(\mathbf{d}, \mathbf{d}') = \sum |d_i - d'_i| = 2$. There are several subcases:

(a) $d'_i = d_i - 2$. Then v_i has at least $n - d - 1$ non-neighbors in a graph for \mathbf{d}' , and two of these must be joined by an edge (u, w) (else we have a large independent set), so we may remove the edge (u, w) and add the edges (v_i, u) and (v_i, w) , obtaining \mathbf{d} . Thus $|\Omega_{\mathbf{d}'}|/|\Omega_{\mathbf{d}}| \leq n^2$.

(b) $d'_i = d_i + 2$. If v_i has two non-adjacent neighbors u, w we proceed as in (a) by complementation. If the adjacent neighbors u, w have a vertex x adjacent to u and not to w , then we remove $(v_i, u), (v_i, w), (x, u)$, and add (x, w) , giving for $u = v_j$ the situation $d'_j = d_j - 2$ from (a). Thus $|\Omega_{\mathbf{d}'}|/|\Omega_{\mathbf{d}}| \leq n^5$. If u, w have the same neighbors other than v_i , then as in (a) there must be two non-neighbors of v_i forming an edge (x, y) , so we remove $(x, y), (v_i, u), (v_i, w)$ and add $(u, x), (w, y)$. Thus $|\Omega_{\mathbf{d}'}|/|\Omega_{\mathbf{d}}| \leq n^4$.

(c) $d'_i = d_i + 1, d'_j = d_j + 1$. If v_i, v_j are joined by an edge we remove this edge. If there is a vertex u adjacent to v_i and not to v_j , we remove edge (u, v_i) and add edge (u, v_j) , resulting in case (b). Thus $|\Omega_{\mathbf{d}'}|/|\Omega_{\mathbf{d}}| \leq n^6$. Otherwise some vertex w is adjacent to both v_i, v_j , and removing edges (w, v_i) and (w, v_j) results in case (a). Thus $|\Omega_{\mathbf{d}'}|/|\Omega_{\mathbf{d}}| \leq n^3$.

(d) $d'_i = d_i - 1, d'_j = d_j - 1$. This case is analogous to (c) by complementing and exchanging the roles of subcases (a) and (b) that arise.

(e) $d'_i = d_i + 1, d'_j = d_j - 1$. Removing or adding the edge (v_i, v_j) results in cases (a) or (b) respectively.

Summarizing, $|\Omega_{\mathbf{d}'}|/|\Omega_{\mathbf{d}}| \leq n^6$ for a change of two in degrees, so $|\Omega_{\mathbf{d}'}|/|\Omega_{\mathbf{d}}| \leq n^{24}$ for a change of 8 in degrees. Since $\lambda/|\Omega_{\mathbf{d}'}| = O(n^8)$, we have $\lambda/|\Omega| = O(n^{32})$, giving the following.

Theorem 7 *Suppose for every i we have*

$$\frac{d+1}{n-d} d \leq d_i \leq d$$

for some $d < n/2$ (for example, $1 \leq d_i \leq \sqrt{n} - 1$, or $n/6 \leq d_i \leq n/3 - 1$). Then the mixing rate of the switch Markov chain on the set of graphs with the degree sequence d_1, d_2, \dots, d_n is

$$\tau(\epsilon) \leq O(d^9 n^{35} (dn \log(dn) + \log(\epsilon^{-1}))).$$

4.2 The Flip chain on general graphs

The path argument given in Lemma 3 for the Flip Markov chain on regular graphs breaks down for non-regular degree sequences. However, it is possible to extend this lemma for general degree sequences when we add a bow-tie operation. as in Figure(3). The bow-tie switch is defined as follows:

if there is a vertex u adjacent to all four of x, y, z, t then switch (x, y) and (z, t) with (x, z) and (y, t) if the graph remains simple.

In the appendix we show that for degree sequences such that the graph diameter is at least 4, any bow-tie switch can be simulated by at path of most 5 switches in $\mathcal{G}_{\mathcal{F}}$. Using that and Theorem 7, we will have

Theorem 8 *For any graphical degree sequences $d_1 \geq d_2 \dots \geq d_n \geq 2$ that enforces a diameter bigger than 3, and satisfies the condition of Theorem 7, the Flip Markov chain is rapidly mixing.*

Acknowledgement: We would like to thank Persi Diaconis, Hector Garcia-Molina, Christian Schiedler, and Prasad Tetali for inspiring discussions.

References

- [1] Berge, C, Theory of Graphs.
- [2] I. Bezáková, N. Bhatnagar, E. Vigoda, “Sampling binary contingency tables with a greedy start”, submitted, 2005.

- [3] C. Cooper, M. Dyer and C. Greenhill, “Sampling regular graphs and a peer-to-peer network.”, SODA 2005, 980-988. (journal version available at www.maths.unsw.edu.au/csg/publications.html)
- [4] M. Dyer, L. Goldberg, M. Jerrum, R. Martin, “Markov chain comparison”, arxiv.org, 2004.
- [5] Gkantsidis, C., Mihail, M. and Saberi, A., On the Random Walk Method in Peer-to-Peer Networks, Proc. Infocom 2004.
- [6] Gkantsidis, C., Mihail, M., and Zegura, E. (2003). The Markov chain simulation method for generating connected power law random graphs.
- [7] S. Janson, T. Luczak and A. Rucinski, “Random graphs,” Wiley and Sons, New York, 2000.
- [8] Jerrum, M. and Sinclair, A. (1990). Fast uniform generation of regular graphs. Theoret. Comput. Sci. 73, 1, 91.100.
- [9] Jerrum, M. and Sinclair, A. in ”Approximation Algorithms for NP-hard Problems,” D.S.Hochbaum ed., PWS Publishing, Boston, 1996.
- [10] Jerrum, M., Sinclair, A., and McKay, B. (1992). When is a graphical sequence stable? In Random graphs, Vol. 2 (Poznan, 1989). Wiley-Intersci. Publ. Wiley, New York, 101.115.
- [11] B.D. McKay, “Asymptotics for symmetric 0-1 matrices with prescribed row sums,” *Ars Combinatoria*, 19A (1985) 15–25.
- [12] B.D. McKay and N.C. Wormald, “Asymptotic enumeration by degree sequence of graphs of high degree,” *European J. Combinatorics*, 11 (1990) 565–580.
- [13] B.D. McKay and N.C. Wormald, “Asymptotic enumeration by degree sequence of graphs with degrees $o(n^{1/2})$,” *Combinatorica* 11-4 (1991) 369–382.
- [14] R. Kannan, P. Tetali, S. Vempala, “Simple Markov-chain algorithms for generating bipartite graphs and tournaments”, *Random Struct. Algorithms* 14(4): 293-308 (1999).
- [15] P. Mahlmann, C. Schindelhauer, “Peer-to-peer networks based on random transformations of connected regular undirected graphs,” SPAA (Symposium on Parallelism in Algorithms and Architectures), Las Vegas, NV 2005.
- [16] McKay, B. D. and Wormald, N. C. (1990b). Uniform generation of random regular graphs of moderate degree. J. Algorithms 11, 1, 52.67.
- [17] Hector Garcia-Molina, private communication.
- [18] M. Mihail, N. Vishnoi, “On Generating Graphs with Prescribed Degree Sequences for Complex Network Modeling Applications”, Position Paper, ARACNE (Approx. and Randomized Algorithms for Communication Networks), Rome, IT, 2002.

- [19] Viger, F. and Latapy, M. (2005). Fast generation of random connected graphs with prescribed degrees. Preprint.
- [20] A. Sinclair, “Improved bounds for mixing rates of Markov chains and multicommodity flow,” *Combinatorics, Probability and Computing 1* (1992), 351–370.
- [21] A. Steger, N. Wormald, “Generating random regular graphs quickly,” *Combinatorics, Probability and Computing 8*, 1999, 377-396.
- [22] R. Taylor. “Constrained switchings in graphs,” *SIAM J. Alg. DISC. METH.*, 3(1):115.121, 1982.
- [23] N. Wormald, “Models of random regular graphs,” *Surveys in Peer-to-peer networks based on random tranCombinatorics*, 1999 (J. Lamb and D Preece eds.), London Math. Soc. Lecture Notes Series 267, CUP, 1999, 239–298.

5 Appendix

5.1 Flip on graphs with diameter bigger than 3

We now show that for degree sequences such that the graph diameter is at least 4, any bow-tie switch can be simulated by at path of most 5 switches in $\mathcal{G}_{\mathcal{F}}$.

In order to do this, it is useful to define the following transforms in the case where we wish to complete the switch $\{(x, y), (z, t)\} \rightarrow \{(x, z), (y, t)\}$. Note that if x and t are adjacent or z and y are adjacent, we can perform this switch in one step, so in all future examples we do not consider this case.

1. If x, y, t are adjacent to v but z is not, then perform the following series of switches:
 - $\{(x, v), (t, z)\} \rightarrow \{(x, t), (v, z)\}$
 - $\{(y, x), (t, v)\} \rightarrow \{(y, t), (x, v)\}$
 - $\{(t, x), (v, z)\} \rightarrow \{(t, v), (x, z)\}$

2. If some vertex q is adjacent to x but not adjacent to v , then perform the following series of switches:
 - $\{(q, x), (v, t)\} \rightarrow \{(q, v), (x, t)\}$
 - $\{(x, y), (z, t)\} \rightarrow \{(x, z), (y, t)\}$
 - $\{(q, v), (x, t)\} \rightarrow \{(q, x), (v, t)\}$

Now, we prove the following lemma by explicitly building a path for each subcase.

Lemma 9 *Consider $\mathcal{G}_{\mathcal{F}}$ for degree sequences such that the diameter of each $G \in \Omega_{\mathcal{F}}$ is at least 4. Then every bow-tie switch $(G, G') \in \mathcal{G}_{SC}$ can be simulated with a path of length at most 5 $p'(G, G')$ in $\mathcal{G}_{\mathcal{F}}$.*

Proof. Since G has diameter at least 4, then either there exists a vertex q such that $d(v, q) = 3$, or there must exist two vertices q and r such that $d(v, q) = d(v, r) = 2$ and $d(q, r) = 4$. Otherwise we could construct a path of length 3 or less between every two vertices in G .

Case 1: There exists a q such that $d(q, v) = 3$.

Either q is distance 2 from one of x, y, z, t , or q is distance at least 3 from one of x, y, z, t , or q is distance 4 from each of x, y, z, t .

If q is distance 2 from x , let r be the intermediate node. Then we can switch $\{(q, r), (x, v)\} \rightarrow \{(q, x), (r, v)\}$. Now x and v are no longer connected, so we can apply transform 1. Finally switch $\{(q, x), (r, v)\} \rightarrow \{(q, r), (x, v)\}$ to complete the series.

In the second subcase, let $p = (q, r, s, v)$ be a shortest path between q and v . Without loss of generality, assume that $d(x, q) = 3$. First switch $\{(q, r), (s, v)\} \rightarrow \{(q, s), (r, v)\}$. Now, s is no longer adjacent to v but is adjacent to x , so we can apply transform 2, then perform $\{(q, s), (r, v)\} \rightarrow \{(q, r), (s, v)\}$ to complete the series.

In the third subcase, let $p = (q, r, s, v)$ be a shortest path between q and v . s is not adjacent to x , so we can switch $\{(r, s), (v, x)\} \rightarrow \{(r, v), (s, x)\}$. Now apply transform 1, then switch back $\{(r, v), (s, x)\} \rightarrow \{(r, s), (v, x)\}$ to complete the series.

Case 2: There exists two vertices q and r such that $d(v, q) = d(v, r) = 2$ and $d(q, r) = 4$.

If either q or r are distance 1 from one of x, y, z, t , we can apply transform 1 directly to complete the switch.

If either q or r are distance 3 from one of x, y, z, t , we can apply the same switches as in the third subcase of the previous case to complete the switch.

Otherwise, q and r must be distance 2 from each of x, y, z, t . Let $p = (q, s, v, w, r)$ be a shortest path from q to r such that $d(w, x) = 1$. We can then switch $\{(q, s), (v, w)\} \rightarrow \{(q, v), (s, w)\}$. Now w is adjacent to x but not v , so we can apply transform 2, then switch $\{(q, v), (s, w)\} \rightarrow \{(q, s), (v, w)\}$ to complete the series. \square

Combining Lemma 9 with the arguments from Theorem 3 implies the following result.

Theorem 10 *For any degree sequences such that the diameter is at least 4 and $d \geq 2$, every edge (G, G') in \mathcal{G}_{SC} can be simulated with a path $p'(G, G')$ in $\mathcal{G}_{\mathcal{F}}$ whose intermediate G_i differ in at most 22 edges from G and from G' . Moreover, the mixing time for the Markov chain $\mathcal{M}_{\mathcal{F}}$ defined by $\mathcal{G}_{\mathcal{F}}$ is*

$$\tau_{\mathcal{F}}(\epsilon) = O(d^{12}n^{14}\tau_{SC}(\epsilon)).$$