

Rapid Mixing of the Switch Markov Chain for Strongly Stable Degree Sequences and 2-Class Joint Degree Matrices*

Georgios Amanatidis

Centrum Wiskunde & Informatica (CWI)
Amsterdam, The Netherlands

georgios.amanatidis@cwi.nl

Pieter Kleer

Centrum Wiskunde & Informatica (CWI)
Amsterdam, The Netherlands

kleer@cwi.nl

October 29, 2018

Abstract

The *switch Markov chain* has been extensively studied as the most natural Markov Chain Monte Carlo approach for sampling graphs with prescribed degree sequences. We use comparison arguments with other—less natural but simpler to analyze—Markov chains, to show that the switch chain mixes rapidly in two different settings. We first study the classic problem of uniformly sampling simple undirected, as well as bipartite, graphs with a given degree sequence. We apply an embedding argument, involving a Markov chain defined by Jerrum and Sinclair (TCS, 1990) for sampling graphs that *almost* have a given degree sequence, to show rapid mixing for degree sequences satisfying *strong stability*, a notion closely related to *P*-stability. This results in a much shorter proof that unifies the currently known rapid mixing results of the switch chain and extends them up to sharp characterizations of *P*-stability. In particular, our work resolves an open problem posed by Greenhill (SODA, 2015).

Secondly, in order to illustrate the power of our approach, we study the problem of uniformly sampling graphs for which—in addition to the degree sequence—a *joint degree distribution* is given. Although the problem was formalized over a decade ago, and despite its practical significance in generating synthetic network topologies, small progress has been made on the random sampling of such graphs. The case of a single degree class reduces to sampling of regular graphs, but beyond this almost nothing is known. We fully resolve the case of two degree classes, by showing that the switch Markov chain is *always* rapidly mixing. Again, we first analyze an auxiliary chain for strongly stable instances on an augmented state space and then use an embedding argument.

*This work is supported by NWO Gravitation Project NETWORKS, Grant Number 024.002.003.

Contents

1	Introduction	1
2	Preliminaries	3
2.1	Graphical Degree Sequences	4
2.2	Joint Degree Matrix Model	5
3	Sampling Undirected Graphs	6
3.1	Flow for the Jerrum-Sinclair Chain	7
3.2	Flow Transformation	7
4	Sampling Graphs with a Given JDM	10
5	Discussion	11
A	Missing Material from Section 3	12
A.1	On the Proofs of Corollaries 5 and 6	12
A.2	Proof of Theorem 7	13
B	Sampling Bipartite Graphs	19
C	The JDM Model: An Example	22
D	The PAM Model and the Hinge Flip Markov Chain	23
D.1	Proof of Theorem 23	25
D.1.1	Canonical Paths	25
D.1.2	Encoding	31
D.1.3	Bounding the Congestion	34
E	Strongly Stable Families for the PAM Model	35
F	Switch Markov Chain for Two Regular Degree Classes	40
	References	41

1 Introduction

The problem of (approximately) uniform sampling of simple graphs with a given degree sequence has received considerable attention, and has applications in domains as diverse as hypothesis testing in network structures [35] and ecology, see, e.g., [37] and references therein. We refer the interested reader to [5] for more pointers to possible applications.

Several variants of the problem have been studied, including sampling undirected, bipartite, connected, or directed graphs. In particular, there is an extensive line of work on Markov Chain Monte Carlo (MCMC) approaches, see, e.g., [29, 30, 20, 7, 25, 34, 19, 18, 9, 6]. In such an approach one studies a random walk on the space of all graphical realizations.¹ This random walk is induced by making small local changes to a given realization using a probabilistic procedure that defines the Markov chain. The idea, roughly, is that after a sufficient number of steps, the so-called *mixing time*, the resulting graph corresponds to a sample from an almost uniform distribution over all graphical realizations of the given degree sequence. The goal is to show that the chain mixes *rapidly*, meaning that one only needs to perform a polynomial (in the size of the graph) number of transitions of the chain in order to obtain an approximately uniform sample. One of the most well-known probabilistic procedures for making these small changes uses local operations called *switches* (also known as *swaps* or *transpositions*); see, e.g., [37] and Figure 1 for an example. The resulting *switch Markov chain* has been shown to be rapidly mixing for various degree sequences [7, 24, 25, 34, 19, 18], but it is still open whether it is rapidly mixing for all degree sequences.




Figure 1: Example of a switch in which edges $\{v, w\}, \{x, y\}$ are replaced by $\{v, y\}, \{x, w\}$. Note that the degree sequence is preserved when applying a switch operation.

In this work, besides the problem of sampling undirected, as well as bipartite, simple graphs with a given degree sequence, we also focus on the problem of sampling undirected simple graphs with a given *joint degree distribution*. That is, in addition to the degrees, the number of edges between nodes of degree i and degree j is also specified for every pair (i, j) .² The motivation for using such a metric is that this extra information restricts the space of possible realizations to graphs with more desirable structure. This was first observed by Mahadevan et al. [32] who argued that the joint degree distribution is a much more reliable metric for a synthetic graph to resemble a real network topology, compared to just using the degree sequence. The joint degree matrix model of Amanatidis, Green, and Mihail [1] formalizes this approach. Although there are polynomial-time algorithms that produce a graphical realization of a given joint degree distribution [1, 2, 41, 11, 22], it is not known how to uniformly sample such a realization efficiently. In particular, bounding the mixing time of the natural restriction of the switch Markov chain for this setting has been an open problem since the introduction of the model [1, 41, 19].

Our Contribution. The proofs of the results in [34, 24, 19, 18] for the analysis of the switch Markov chain in undirected (or bipartite) graphs are all using conceptually similar ideas to the ones introduced by Cooper, Dyer and Greenhill [7] for the analysis of the switch chain for regular undirected graphs, and are based on the multicommodity flow method of Sinclair [39]. The individual parts of this template can become quite technical and require long proofs. In this work we take a different approach for proving that the switch chain

¹Sometimes the state space is augmented with a set of auxiliary states, as in [29].

²We refer the reader to Subsection 2.2 for an exact definition and also Appendix C for an example.

is rapidly mixing. First we analyze some easier auxiliary Markov chain; such a chain can be used to sample graphical realizations that *almost* have a given fixed degree sequence or joint degree distribution. We show that there exists an efficient multicommodity flow for the auxiliary chain when the given instance is *strongly stable*,³ and then show how it can be transformed into an efficient multicommodity flow for the switch chain. Note that this last step compares two Markov chains with different state spaces, as the auxiliary chain samples from a strictly larger set of graphs than the switch chain. Thus, for the flow transformation we use embedding arguments similar to those in Feder et al. [20].⁴

Using the aforementioned approach we obtain the following two main results:

- 1) We show rapid mixing of the switch chain for *strongly stable* families of degree sequences (Theorem 4), thus providing a rather short proof that unifies and extends the results in [7, 24] (Corollaries 5 and 6).

We introduce strong stability as a stricter version of the notion of P -stability [29]. The strong stability condition is satisfied by the degree sequences in the works [30, 7, 24, 34, 19, 18] and by characterizations of P -stability [27]. These characterizations are employed to get Corollaries 5 and 6. In particular, investigating the mixing time for the degree sequences of Corollary 6 was posed as an open question by Greenhill [24]; here we resolve it in the positive.

- 2) We show that the switch chain restricted on the space of the graphic realizations of a given joint degree distribution with two degree classes is *always* rapidly mixing (Theorem 10).

Despite being for the case of two classes, this is the very first rapid mixing result for the problem. Establishing the rapid mixing of the auxiliary chain in this case presents significant challenges. To attack this problem, we rely on ideas introduced by Bhatnagar, Randall, Vazirani and Vigoda [4]. At the core of this approach lies the *mountain-climbing problem* [26, 43]. We should note that the auxiliary chain used here is analyzed in a much more general model than the joint degree matrix model.

Besides the above results we also unify the results in [30, 34, 19, 18] for bipartite degree sequences and extend them for the special case of an equally sized bipartition; see Corollaries 18 and 19 in Appendix B. We should note that the unification of the existing results mentioned so far is qualitative rather than quantitative, in the sense that our simpler, indirect approach provides weaker polynomial bounds for the mixing time. For examples of explicit mixing time bounds we refer the reader to [7, 8, 25].

Finally, as a byproduct of our analysis for the auxiliary Markov chain used to prove Theorem 10, we obtain the first *fully polynomial almost uniform generator* [40] for sampling graphical realizations of certain sparse *partition adjacency matrix* instances with two partition classes [10, 15] (this is a generalization of the joint degree distribution problem; see Appendix D for definitions). See Corollary 34 in Appendix E.

Related Work. Here the focus is on MCMC approaches. As such approaches have impractical mixing times in general, we should note that there is a line of work on non-MCMC sampling algorithms which, albeit having weaker properties, do have practical running times. See, e.g., [3, 21] and references therein.

Jerrum and Sinclair [29] give a fully polynomial almost uniform generator for generating graphical realizations of degree sequences coming from any P -stable family of sequences (see preliminaries). The auxiliary chain we use to prove Theorem 4, henceforth referred to as Jerrum-Sinclair (JS) chain, is presented in [29] as a more straightforward implementation of this generator. One drawback is that the algorithms in [29] work with auxiliary states. Kannan, Tetali and Vempala [30] introduce the switch chain as a simpler and more direct generator that does not have to work with auxiliary states. They addressed the mixing time

³In the case of sampling graphs with a given degree sequence, the existence of such a flow for the auxiliary chain we use was already claimed in [29], at least for the degree sequences satisfying a more restrictive analog of inequality (2). For completeness, we give a detailed proof in Appendix A.2.

⁴We also refer the reader to the survey of Dyer et al. [13] for more on Markov chain comparison techniques.

of such a switch-based Markov chain for the regular bipartite case. Cooper et al. [7] then gave a rapid mixing proof for regular undirected graphs, and later Greenhill [24] extended this result to certain ranges of irregular degree sequences (see also Greenhill and Sfragara [25]). Miklós, Erdős and Soukup [34] proved rapid mixing for the *half-regular* bipartite case, and Erdős, Miklós and Toroczkai [19] for the *almost half-regular* case. Very recently, Erdős et al. [18] presented a range of bipartite degree sequences unifying and generalizing the results in [34, 19].

Switch-based Markov chain Monte Carlo approaches have also been studied for other graph sampling problems. Feder et al. [20] as well as Cooper et al. [9] study the mixing time of a Markov chain using a switch-like probabilistic procedure (called a *flip*) for sampling connected graphs. For sampling perfect matchings, switch-based Markov chains have also been studied, see, e.g., the recent work of Dyer, Jerrum and Müller [14] and references therein. It is interesting to note that Dyer et al. [14] also use a lemma on the mountain-climbing problem that is very similar to Lemma 26.

The joint degree matrix model was first studied by Patrinos and Hakimi [36], albeit with a different formulation and name, and was reintroduced in Amanatidis et al. [1]. While it has been shown that the switch chain restricted on the space of the graphic realizations of any given joint degree distribution is irreducible [1, 11], almost no progress has been made towards bounding its mixing time. Stanton and Pinar [41] performed experiments based on the autocorrelation of each edge, suggesting that the switch chain mixes quickly. The only relevant result is that of Erdős et al. [19] showing fast mixing for a related Markov chain over the severely restricted subset of the so-called *balanced* joint degree matrix realizations; this special case, however, lacks several of the technical challenges that arise in the original problem.

Outline. In Section 2 we give all the necessary preliminaries and we formally describe the JS chain, the switch chain, and the restricted switch chain. Our first main result is Theorem 4 in Section 3, where we show that the switch chain is rapidly mixing for families of strongly stable degree sequences. Given the rapid mixing of the JS chain (Appendix A.2), the proof of Theorem 4 in Section 3 is self-contained. The corresponding result for the bipartite case is completely analogous and is deferred to Appendix B. In Section 4 we state our second main result, Theorem 10. For the sake of presentation, we defer the proof of Theorem 10 to Appendices D, E and F, and we only include a short discussion of our approach in Section 4.

2 Preliminaries

We begin with the preliminaries regarding Markov chains and the multicommodity flow method of Sinclair [39]. For Markov chain definitions not given here, see for example [31].

Let $\mathcal{M} = (\Omega, P)$ be an ergodic, time-reversible Markov chain over state space Ω with transition matrix P and stationary distribution π . We write $P^t(x, \cdot)$ for the distribution over Ω at time step t given that the initial state is $x \in \Omega$. The *total variation distance* at time t with initial state x is

$$\Delta_x(t) = \max_{S \subseteq \Omega} |P^t(x, S) - \pi(S)| = \frac{1}{2} \sum_{y \in \Omega} |P^t(x, y) - \pi(y)|,$$

and the *mixing time* $\tau(\epsilon)$ is defined as $\tau(\epsilon) = \max_{x \in \Omega} \{\min\{t : \Delta_x(t') \leq \epsilon \text{ for all } t' \geq t\}\}$. Informally, $\tau(\epsilon)$ is the number of steps until the Markov chain is ϵ -close to its stationary distribution. A Markov chain is said to be *rapidly mixing* if the mixing time can be upper bounded by a function polynomial in $\ln(|\Omega|/\epsilon)$.

It is well-known that, since the Markov chain is time-reversible, the matrix P only has real eigenvalues $1 = \lambda_0 > \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{|\Omega|-1} > -1$. We may replace the transition matrix P of the Markov chain by $(P + I)/2$, to make the chain *lazy*, and hence guarantee that all its eigenvalues are non-negative. It then follows that the second-largest eigenvalue of P is λ_1 . In this work we always consider the lazy versions of the Markov chains involved. It follows directly from Proposition 1 in [39] that $\tau(\epsilon) \leq \frac{1}{1-\lambda_1} (\ln(1/\pi_*) + \ln(1/\epsilon))$,

where $\pi_* = \min_{x \in \Omega} \pi(x)$. For the special case where $\pi(\cdot)$ is the uniform distribution, as is the case here, the above bound becomes $\tau(\epsilon) \leq \ln(|\Omega|/\epsilon)/(1 - \lambda_1)$. The quantity $(1 - \lambda_1)^{-1}$ can be upper bounded using the *multicommodity flow method* of Sinclair [39].

We define the state space graph of the chain \mathcal{M} as the directed graph \mathbb{G} with node set Ω that contains exactly the arcs $(x, y) \in \Omega \times \Omega$ for which $P(x, y) > 0$ and $x \neq y$. Let $\mathcal{P} = \cup_{x \neq y} \mathcal{P}_{xy}$, where \mathcal{P}_{xy} is the set of simple paths between x and y in the state space graph \mathbb{G} . A *flow* f in Ω is a function $\mathcal{P} \rightarrow [0, \infty)$ satisfying $\sum_{p \in \mathcal{P}_{xy}} f(p) = \pi(x)\pi(y)$ for all $x, y \in \Omega, x \neq y$. The flow f can be extended to a function on oriented edges of \mathbb{G} by setting $f(e) = \sum_{p \in \mathcal{P}: e \in p} f(p)$, so that $f(e)$ is the total flow routed through $e \in E(\mathbb{G})$. Let $\ell(f) = \max_{p \in \mathcal{P}: f(p) > 0} |p|$ be the length of a longest flow carrying path, and let $\rho(e) = f(e)/Q(e)$ be the *load* of the edge e , where $Q(e) = \pi(x)P(x, y)$ for $e = (x, y)$. The maximum load of the flow is $\rho(f) = \max_{e \in E(\mathbb{G})} \rho(e)$. Sinclair ([39], Corollary 6') shows that $(1 - \lambda_1)^{-1} \leq \rho(f)\ell(f)$.

We use the following standard technique for bounding the maximum load of a flow in case the chain \mathcal{M} has uniform stationary distribution π . Suppose θ is the smallest positive transition probability of the Markov chain between two distinct states. If b is such that $f(e) \leq b/|\Omega|$ for all $e \in E(\mathbb{G})$, then it follows that $\rho(f) \leq b/\theta$. Thus, we have

$$\tau(\epsilon) \leq \frac{\ell(f) \cdot b}{\theta} \ln(|\Omega|/\epsilon).$$

Now, if $\ell(f), b$ and $1/\theta$ can be bounded by a function polynomial in $\log(|\Omega|)$, it follows that the Markov chain \mathcal{M} is rapidly mixing. In this case, we say that f is an *efficient* flow. Note that in this approach the transition probabilities do not play a role as long as $1/\theta$ is polynomially bounded.

2.1 Graphical Degree Sequences

A sequence of non-negative integers $d = (d_1, \dots, d_n)$ is called a *graphical degree sequence* if there exists a simple, undirected, labeled graph on n nodes having degrees d_1, \dots, d_n ; such a graph is called a *graphical realization* of d . For a given degree sequence d , $\mathcal{G}(d)$ denotes the set of all graphical realizations of d . Throughout this work we only consider sequences d with positive components, and for which $\mathcal{G}(d) \neq \emptyset$. Let $\mathcal{G}'(d) = \cup_{d'} \mathcal{G}(d')$ with d' ranging over the set $\{d' : d'_j \leq d_j \text{ for all } j, \text{ and } \sum_{i=1}^n |d_i - d'_i| \leq 2\}$. That is, we have (i) $d' = d$, or (ii) there exist distinct κ, λ such that $d'_i = d_i - 1$ if $i \in \{\kappa, \lambda\}$ and $d'_i = d_i$ otherwise, or (iii) there exists a κ so that $d'_i = d_i - 2$ if $i = \kappa$ and $d'_i = d_i$ otherwise. In the case (ii) we say that d' has two nodes with degree deficit one, and in the case (iii) we say that d' has one node with degree deficit two. A family \mathcal{D} of graphical degree sequences is called *P-stable* [29], if there exists a polynomial $q(n)$ such that for all $d \in \mathcal{D}$ we have $|\mathcal{G}'(d)|/|\mathcal{G}(d)| \leq q(n)$, where n is the number of components of d .

Jerrum and Sinclair [29] define the following Markov chain on $\mathcal{G}'(d)$, which will henceforth be referred to as the *JS chain*.⁵ Let $G \in \mathcal{G}'(d)$ be the current state of the chain:

- With probability $1/2$, do nothing.
- Otherwise, select an ordered pair i, j of nodes uniformly at random and
 - if $G \in \mathcal{G}(d)$ and (i, j) is an edge of G , then delete (i, j) from G ,
 - if $G \notin \mathcal{G}(d)$, the degree of i in G is less than d_i , and (i, j) is not an edge of G , then add (i, j) to G .
 If the new degree of j exceeds d_j , then select an edge (j, k) uniformly at random and delete it.

The graphs $G, G' \in \mathcal{G}'(d)$ are *JS adjacent* if G can be obtained from G' with positive probability in one transition of the JS chain and vice versa. The following properties of the JS chain are easy to check.

⁵A slightly different definition of stability is given by Jerrum, McKay and Sinclair [27]. Based on this variant, one could define the corresponding variant of the JS chain. Nevertheless, the definitions of stability in [27] and [29] (and their corresponding definitions of strong stability) are equivalent. To avoid confusion, here we only use the definitions in [29].

Theorem 1 (Follows by [29]). *The JS chain is irreducible, aperiodic and symmetric, and, hence, has uniform stationary distribution over $\mathcal{G}'(d)$. Moreover, $P(G, G')^{-1} \leq 2n^3$ for all JS adjacent $G, G' \in \mathcal{G}'(d)$, and also the maximum in- and out-degrees of the state space graph of the JS chain are bounded by n^3 .*

We say that two graphs G, G' are within distance r in the JS chain if there exists a path of at most length r from G to G' in the state space graph of the JS chain. By $\text{dist}(G, d)$ we denote the minimum distance of G from an element in $\mathcal{G}(d)$. The following parameter will play a central role in this work. Let

$$k_{JS}(d) = \max_{G \in \mathcal{G}'(d)} \text{dist}(G, d). \quad (1)$$

Based on the parameter $k_{JS}(d)$, we define the notion of *strong stability*. The simple observation in Proposition 3 justifies the terminology. For the different settings studied in this work, i.e., for sampling bipartite graphs or joint degree matrix realizations, the definition of k_{JS} is adjusted accordingly (see Appendices B and D).

Definition 2 (Strong stability). A family of graphical degree sequences \mathcal{D} is called *strongly stable* if there exists a constant ℓ such that $k_{JS}(d) \leq \ell$ for all $d \in \mathcal{D}$.

Proposition 3. *If \mathcal{D} is strongly stable, then it is P-stable.*

Proof. Suppose \mathcal{D} is strongly stable with respect to the constant ℓ . Let $d \in \mathcal{D}$ be a degree sequence with n components. For every $G \in \mathcal{G}'(d) \setminus \mathcal{G}(d)$ choose some $\varphi(G) \in \mathcal{G}(d)$ within distance $k = k_{JS}(d)$ of G . As the in-degree of any node in the state space graph of the JS chain is bounded by n^3 , the number of paths with length at most k that end up at any particular graph in $\mathcal{G}(d)$ is upper bounded by $(n^3)^k$. Therefore, $|\mathcal{G}'(d)|/|\mathcal{G}(d)| \leq n^{3k} \leq n^{3\ell}$, meaning that \mathcal{D} is stable, since ℓ is constant. \square

Finally, the lazy version of the *switch chain on $\mathcal{G}(d)$* is defined as follows (see, e.g., [30, 7]). Let $G \in \mathcal{G}(d)$ be the current state of the chain:

- With probability $1/2$, do nothing.
- Otherwise, select two edges $\{a, b\}$ and $\{x, y\}$ uniformly at random, and select a perfect matching M on nodes $\{x, y, a, b\}$ uniformly at random (there are three possible options). If $M \cap E(G) = \emptyset$, then delete $\{a, b\}, \{x, y\}$ from $E(G)$ and add the edges of M . This local operation is called a *switch*.

The graphs $G, G' \in \mathcal{G}(d)$ are *switch adjacent* if G can be obtained from G' with positive probability in one transition of this chain and vice versa. It is well-known that the switch chain is irreducible, aperiodic and symmetric (e.g., [25] and references therein), and, thus, has uniform stationary distribution over $\mathcal{G}(d)$. Furthermore, is it a matter of simple counting that $P(G, G')^{-1} \leq 6n^4$ for all switch adjacent $G, G' \in \mathcal{G}(d)$, and the maximum in- and out-degrees of the state space graph of the switch chain are bounded by n^4 .

2.2 Joint Degree Matrix Model

Here in addition to the degrees, we would also like to specify the number of edges between nodes of degree i and nodes of degree j for every pair (i, j) . Let $V = \{1, \dots, n\}$ be a set of nodes. An instance of the joint degree matrix (JDM) model is given by a partition $V_1 \cup V_2 \cup \dots \cup V_q$ of V into pairwise disjoint (degree) classes, a symmetric *joint degree matrix* $c = (c_{ij})_{i, j \in [q]}$ of non-negative integers, and a sequence $d = (d_1, \dots, d_q)$ of non-negative integers.⁶ We say that the tuple $((V_i)_{i \in [q]}, c, d)$ (or just (c, d) when it is clear what the partition is) is graphical, if there exists a simple, undirected, labeled graph $G = (V, E)$ on the nodes in V such that all nodes in V_i have degree d_i and there are precisely c_{ij} edges between nodes in V_i and V_j . Such a G is called

⁶This is shorthand notation for the degree sequence. Alternatively, we could write $\hat{d} = (d_1^1, \dots, d_1^{|V_1|}, \dots, d_q^1, \dots, d_q^{|V_q|})$ corresponding to the definition of a graphical degree sequence. In such a case, $d_i^j = d_i$ for $i \in V$ and $j \in \{1, \dots, |V_i|\}$.

a graphical realization of the tuple. We let $\mathcal{G}((V_i)_{i \in q}, c, d)$ (or just $\mathcal{G}(c, d)$) denote the set of all graphical realizations of $((V_i)_{i \in q}, c, d)$. We focus on the case of $q = 2$, i.e., when two degree classes are given.

While switches maintain the degree sequence, this is no longer true for the joint degree constraints. However, some switches do respect these constraints as well, e.g., if w, y in Figure 1 are in the same degree class. Thus, we are interested in the following (lazy) *restricted switch Markov chain* for sampling graphical realizations of $\mathcal{G}(c, d)$. Let $G \in \mathcal{G}(c, d)$ be the current state of the chain:

- With probability $1/2$, do nothing.
- Otherwise, try to perform a *switch move*: select two edges $\{a, b\}$ and $\{x, y\}$ uniformly at random, and select a perfect matching M on nodes $\{x, y, a, b\}$ uniformly at random. If $M \cap E(G) = \emptyset$ and $G + M - (\{a, b\} \cup \{x, y\}) \in \mathcal{G}(c, d)$, then delete $\{a, b\}, \{x, y\}$ from $E(G)$ and add the edges of M .

This chain is irreducible, aperiodic and symmetric [1, 11]. Like the switch chain defined above, $P(G, G')^{-1} \leq n^4$ for all adjacent $G, G' \in \mathcal{G}(c, d)$, and also the maximum in- and out-degrees of the state space graph are less than n^4 . Since bounding the mixing time of this chain on $\mathcal{G}(c, d)$ has been elusive [1, 41, 19], we follow a similar approach as in the case of the switch chain for undirected and bipartite graphs. In Appendix D the simpler *hinge flip Markov chain* is defined on a strictly larger state space (and even for a more general model than JDM). This also gives rise to the corresponding strong stability definition. As the whole analysis of this auxiliary chain takes place in the appendix, we defer any further definitions there.

3 Sampling Undirected Graphs

The result in Theorem 4 below is our main result regarding the mixing time of the switch chain for strongly stable degree sequences. Its proof is divided in two parts. First, in Section 3.1, by giving an efficient multicommodity flow, we show that for any d in a family of strongly stable degree sequences the JS chain is rapidly mixing on $\mathcal{G}'(d)$. Then, in Section 3.2, we show that such an efficient flow for the JS chain on $\mathcal{G}'(d)$ can be transformed into an efficient flow for the switch chain on $\mathcal{G}(d)$. This yields the following theorem.

Theorem 4. *Let \mathcal{D} be a strongly stable family of degree sequences with respect to some constant k . Then there exists a polynomial $q(n)$ such that, for any $0 < \epsilon < 1$, the mixing time τ_{sw} of the switch chain for a graphical sequence $d = (d_1, \dots, d_n) \in \mathcal{D}$ satisfies*

$$\tau_{\text{sw}}(\epsilon) \leq q(n)^k \ln(1/\epsilon).$$

We discuss two direct corollaries of Theorem 4. Both corollaries are consequences of the corresponding results in [27], where it is shown that the families of sequences satisfying (2) and (3), respectively, are (strongly) stable. We work with a slightly different definition of stability here than the one used in [27]. The reason why the results from [27] carry over to the definition used here (which was introduced in [29]) is explained in Appendix A.1. The equivalence of these definitions of P -stability is also claimed in [27].

Corollary 5 below extends the rapid mixing results in [7, 25]. In particular, the condition of [25] for rapid mixing is $\delta \geq 1$ and $3 \leq \Delta \leq \frac{1}{3} \sqrt{2m}$, which is a special case of the condition (2) below.⁷

Corollary 5. *Let $\mathcal{D} = \mathcal{D}(\delta, \Delta, m)$ be the set of all graphical degree sequences $d = (d_1, \dots, d_n)$ satisfying*

$$(2m - n\delta)(n\Delta - 2m) \leq (\Delta - \delta)[(2m - n\delta)(n - \Delta - 1) + (n\Delta - 2m)\delta] \quad (2)$$

where δ and Δ are the minimum and maximum component of d , respectively, and $m = \frac{1}{2} \sum_{i=1}^n d_i$. For any $d \in \mathcal{D}$, we have $k_{\text{JS}}(d) \leq 6$. Hence, the switch chain is rapidly mixing for sequences in \mathcal{D} .

⁷The condition in Theorem 1.1 in [25] is a special case of the condition of Theorem 3.1 in [29] which in turn is a special case of the condition of Corollary 5. See also the remark after Theorem 8 in [27].

The next corollary is a special case of the Corollary 5 and answers an open question posed in [25]. It is a result of similar flavor, but the corresponding condition is stated only in terms of δ and Δ .

Corollary 6. *Let $\mathcal{D} = \mathcal{D}(\delta, \Delta)$ be the set of all graphical degree sequences $d = (d_1, \dots, d_n)$ satisfying*

$$(\Delta - \delta + 1)^2 \leq 4\delta(n - \Delta - 1) \quad (3)$$

where δ and Δ are the minimum and maximum component of d , respectively. For any $d \in \mathcal{D}$, we have $k_{JS}(d) \leq 6$. Hence, the switch chain is rapidly mixing for sequences in \mathcal{D} .

Explicit families satisfying these conditions are given in [27]. For instance, all sequences $d = (d_1, \dots, d_n)$ with (i) $\delta(d) \geq 1$ and $\Delta(d) \leq 2\sqrt{n} - 2$, or (ii) $\delta(d) \geq \frac{1}{4}n$ and $\Delta(d) \leq \frac{3}{4}n - 1$ satisfy (3) but not necessarily the conditions in [25]. We refer the reader to [27, 38] for more examples. The bounds in Corollaries 5 and 6 are in a sense best possible with respect to the graph parameters involved. Namely, there exist non-stable degree sequence families the members of which only slightly violate (3); see the discussion in [27] for details.

3.1 Flow for the Jerrum-Sinclair Chain

Jerrum and Sinclair [29] claim, without proof, that the JS chain is rapidly mixing for (some) families of stable degree sequences. For completeness, we prove in Theorem 7 that the chain is rapidly mixing for any family of strongly stable degree sequences. For the proof of the theorem see Appendix A.2.

Theorem 7 ([29]). *Let \mathcal{D} be a strongly stable family of degree sequences with respect to some constant k . Then there exist polynomials $p(n)$ and $r(n)$ such that for any $d = (d_1, \dots, d_n) \in \mathcal{D}$ there exists an efficient multicommodity flow f for the JS chain on $\mathcal{G}'(d)$ satisfying $\max_e f(e) \leq p(n)/|\mathcal{G}'(d)|$ and $\ell(f) \leq r(n)$.*

Our proof of Theorem 7 uses conceptually similar arguments to the ones used in [7] for the analysis of the switch chain on regular undirected graphs. However, the analysis done here for the JS chain is, in our opinion, easier and cleaner than the corresponding analysis for the switch chain. In particular, the so-called *circuit processing procedure* is simpler in our setting, as it only involves altering edges in the symmetric difference of two graphical realizations in a straightforward fashion. In the switch chain analyses [7, 25, 34, 19, 18] one also has to temporarily alter edges that are not in the symmetric difference and this significantly complicates things. Moreover, for the analysis of the JS chain, we can rely on arguments used (in a somewhat different context) by Jerrum and Sinclair [28] for the analysis of a Markov chain for sampling (near) perfect matchings of a given graph. This usage of arguments in [28] was suggested by Jerrum and Sinclair [29] for showing that the JS chain is rapidly mixing for stable degree sequences.

3.2 Flow Transformation

Next we show that, when d comes from a family of strongly stable degree sequences, an efficient multicommodity flow for the JS chain on $\mathcal{G}'(d)$ can be transformed into an efficient multicommodity flow for the switch chain on $\mathcal{G}(d)$. In combination with Theorem 7 this implies that if \mathcal{D} is strongly stable, then for any sequence in \mathcal{D} there exists an efficient flow for the switch chain. For the sake of simplicity, we did not attempt to optimize the bounds in the proof of Theorem 8.

Theorem 8. *Let \mathcal{D} be a strongly stable family of degree sequences with respect to some constant k , and let $p(n)$ and $r(n)$ be polynomials such that for any $d = (d_1, \dots, d_n) \in \mathcal{D}$ there exists an efficient multicommodity flow f_d for the JS chain on $\mathcal{G}'(d)$ with the property that $\max_e f(e) \leq p(n)/|\mathcal{G}'(d)|$ and $\ell(f) \leq r(n)$.*

Then there exists a polynomial $t(n)$ such that for all $d = (d_1, \dots, d_n) \in \mathcal{D}$ there is a feasible multicommodity flow g_d for the switch chain on $\mathcal{G}(d)$ with (i) $\ell(g_d) \leq 2k \cdot \ell(f_d)$, and (ii) for every edge e of the state space graph of the switch chain, we have

$$g_d(e) \leq t(n)^k \cdot \frac{p(n)}{|\mathcal{G}(d)|}. \quad (4)$$

Proof. Let $d \in \mathcal{D}$. For simplicity we will write f and g instead of f_d and g_d respectively. Since there are two Markov chains involved in the proof, each with a different state space graph, we should clarify that \mathcal{P}_{xy} refers to the set of simple paths between x and y in the state space graph of the JS chain. We first introduce some additional notation.

For every pair $(x, y) \in \mathcal{G}'(d) \times \mathcal{G}'(d)$ with $x \neq y$, and for any $p \in \mathcal{P}_{xy}$, we write $\alpha(p) = f(p)|\mathcal{G}'(d)|^2$. Recall that since the stationary distribution of the JS chain is uniform on $\mathcal{G}'(d)$ we have $\sum_{p \in \mathcal{P}_{xy}} f(p) = |\mathcal{G}'(d)|^{-2}$. Thus, $\sum_{p \in \mathcal{P}_{xy}} \alpha(p) = 1$. Moreover, we define $\alpha(e) = \sum_{p \in \mathcal{P}_{xy}: e \in p} \alpha(p) = f(e)|\mathcal{G}'(d)|^2$.

Now, for every $G \in \mathcal{G}'(d) \setminus \mathcal{G}(d)$ choose some $\varphi(G) \in \mathcal{G}(d)$ that is within distance k of G in the JS chain, and take $\varphi(G) = G$ for $G \in \mathcal{G}(d)$. Based on the arguments in the proof of Proposition 3, it follows that for any $H \in \mathcal{G}(d)$,

$$|\varphi^{-1}(H)| \leq n^{3k}, \quad (5)$$

using that the maximum in-degree of any element in the state space graph of the JS chain is upper bounded by n^3 . In particular, this implies that

$$\frac{|\mathcal{G}'(d)|}{|\mathcal{G}(d)|} \leq n^{3k}. \quad (6)$$

Let the flow h be defined as follows for any given pair (x, y) . If $(x, y) \in \mathcal{G}(d) \times \mathcal{G}(d)$, take $h(p) = \alpha(p)/|\mathcal{G}(d)|^2$ for all $p \in \mathcal{P}_{xy}$. If either x or y is not contained in $\mathcal{G}(d)$, take $h(p) = 0$ for every $p \in \mathcal{P}_{xy}$. Note that h is a multicommodity flow that routes $1/|\mathcal{G}(d)|^2$ units of flow between any pair $(x, y) \in \mathcal{G}(d) \times \mathcal{G}(d)$, and zero units of flow between any other pair of states in $\mathcal{G}'(d)$.

Note that

$$h(e) \leq \frac{|\mathcal{G}'(d)|^2}{|\mathcal{G}(d)|^2} \cdot f(e) \leq \frac{|\mathcal{G}'(d)|^2}{|\mathcal{G}(d)|^2} \frac{p(n)}{|\mathcal{G}'(d)|} = \frac{p(n)}{|\mathcal{G}(d)|} \frac{|\mathcal{G}'(d)|}{|\mathcal{G}(d)|} \leq n^{3k} \cdot \frac{p(n)}{|\mathcal{G}(d)|}, \quad (7)$$

using the definition of h in the first inequality, the assumption on f in the second inequality, and the upper bound of (6) in the last one.

Next, we merge the ‘‘auxiliary states’’ in $\mathcal{G}'(d) \setminus \mathcal{G}(d)$, i.e., the states not reached by the switch chain, with the elements of $\mathcal{G}(d)$. Informally speaking, for every $H \in \mathcal{G}(d)$ we merge all the nodes in $\varphi^{-1}(H)$ into a *supernode*. Self-loops created in this process are removed, and parallel arcs between states are merged into one arc that gets all the flow of the parallel arcs. Formally, we consider the graph Γ where $V(\Gamma) = \mathcal{G}(d)$ and $e = (H, H') \in E(\Gamma)$ if and only if H and H' are switch adjacent or if there exist $G \in \varphi^{-1}(H)$ and $G' \in \varphi^{-1}(H')$ such that G and G' are JS adjacent. Moreover, for a given h -flow carrying path $(G_1, G_2, \dots, G_q) = p \in \mathcal{P}_{xy}$, let $p'_\Gamma = (\varphi(G_1), \varphi(G_2), \dots, \varphi(G_q))$ be the corresponding (possibly non-simple) directed path in Γ . Any self-loops and cycles can be removed from p'_Γ and let p_Γ be the resulting simple path in Γ . Over p_Γ we route $h_\Gamma(p_\Gamma) = h(p)$ units of flow. Note that h_Γ is a flow that routes $1/|\mathcal{G}(d)|^2$ units of flow between any pair of states $(x, y) \in \mathcal{G}(d) \times \mathcal{G}(d)$ in the graph Γ and that $\ell(h_\Gamma) \leq \ell(f)$. Furthermore, the flow h_Γ on an edge $(H, H') \in E(\Gamma)$ is then bounded by

$$h_\Gamma(H, H') \leq \sum_{\substack{(G, G') \in \varphi^{-1}(H) \times \varphi^{-1}(H') \\ G \text{ and } G' \text{ are JS adjacent}}} h(G, G'), \quad (8)$$

where the inequality (instead of an equality) follows from the fact that when we map a path $p \in \mathcal{P}_{xy}$ to the corresponding path p_Γ , some edges of the intermediate path p'_Γ may be deleted. Using (5), it follows that $|\varphi^{-1}(H) \times \varphi^{-1}(H')| \leq n^{3k} \cdot n^{3k} = n^{6k}$ and therefore, in combination with (7) and (8), we have that

$$h_\Gamma(e) \leq n^{3k} \cdot n^{6k} \cdot \frac{p(n)}{|\mathcal{G}(d)|}. \quad (9)$$

Now recall how $E(\Gamma)$ was defined. An edge (H, H') might have been added because: (i) H and H' are switch adjacent (we call these edges of Γ *legal*), or (ii) H and H' are *not* switch adjacent but there exist




Figure 2: The dashed edge on the left represents an illegal edge, and the bold path represents a “short” detour. The shortcut path on the right is the result of removing any loops and cycles.

$G \in \varphi^{-1}(H)$ and $G' \in \varphi^{-1}(H')$ that are JS adjacent (we call these edges of Γ *illegal*). The final step of the proof consists of showing that the flow on every illegal edge in $E(\Gamma)$ can be rerouted over a “short” path consisting only of legal edges. In particular, for every flow carrying path p using e , we are going to show that the flow $h_\Gamma(p)$ is rerouted over some legal detour, the length of which is bounded by a multiple of k . Doing this iteratively for every remaining illegal edge on p , we obtain a directed path p'' only using legal edges, i.e., edges of the state space graph of the switch chain. Of course, p'' might not be simple, so any self-loops and cycles can be removed, as before, to obtain the simple legal path p' . Figure 2 illustrates this procedure for a path with a single illegal edge. Note that deleting self-loops and cycles only decreases the amount of flow on an edge.

The crucial observation here is that if $(H, H') \in E(\Gamma)$, then $|E(H) \Delta E(H')| \leq 4k$. That is, even though H and H' might not be switch adjacent, they are not too far apart. To see this, first note that the symmetric difference of any two JS adjacent graphs has size at most 2. Moreover, if one of any two JS adjacent graphs is in $\mathcal{G}(d)$, then their symmetric difference has size 1. In particular, for any $G^* \in \mathcal{G}'(d)$, we have $|E(G^*) \Delta E(\varphi(G^*))| \leq 2k - 1$.

Clearly, if $(H, H') \in E(\Gamma)$ is legal, then $|E(H) \Delta E(H')| = 4 \leq 4k$. Assume $(H, H') \in E(\Gamma)$ is illegal. Then there exist JS adjacent $G \in \varphi^{-1}(H)$ and $G' \in \varphi^{-1}(H')$ and according to the above we have

$$\begin{aligned} |E(H) \Delta E(H')| &\leq |E(H) \Delta E(G)| + |E(G) \Delta E(G')| + |E(G') \Delta E(H')| \\ &\leq 2k - 1 + 2 + 2k - 1 \leq 4k. \end{aligned}$$

Moreover, this implies that we can go from H to H' in a “small” number of moves in the switch chain. This easily follows from most results showing that the state space of the switch chain is connected, e.g., from [42].⁸ Specifically, here we use the following result of Erdős, Király, and Miklós [16] which implies that we can go from H to H' in $2k$ switches.

Theorem 9 (follows from Theorem 3.6 in [16]). *Let $d = (d_1, \dots, d_n)$ be a degree sequence. For any two graphs $H, H' \in \mathcal{G}(d)$, H can be transformed into H' using at most $\frac{1}{2}|E(H) \Delta E(H')|$ switches.*

For every *illegal* edge $e \in E(\Gamma)$, we choose such a (simple) path from H to H' with at most $2k$ transitions and reroute the flow of e over this path. Note that for any legal edge $e \in E(\Gamma)$, the number of illegal edge detours that use e for this rerouting procedure, is at most $(n^4)^{2k} \cdot (n^4)^{2k} = n^{16k}$, using the fact that in the state space graph of the switch chain the maximum degree of an element is at most n^4 and any illegal edge using e in its rerouting procedure must lie within distance $2k$ of e . Combining this with (9), we see that the resulting flow, g , satisfies

$$g(e) \leq \frac{p(n) \cdot n^{9k} + p(n) \cdot n^{16k}}{|\mathcal{G}(d)|}.$$

⁸To be precise, we can focus on the subgraph induced by the nodes with positive degree in the symmetric difference. Taylor’s proof on the connectivity of the state space of the switch chain [42] implies that we can find $O(k^2)$ switches to get from H to H' , only using edges in this induced subgraph.

Note that the $\ell(g) \leq 2k\ell(h_\Gamma)$. This holds because every illegal edge on a flow-carrying path gives rise to at most $2k$ additional edges as a result of rerouting the flow over legal edges, and the removal of loops and cycles from any resulting non-simple path can only decrease its length. Combining this inequality with $\ell(h_\Gamma) \leq \ell(f)$ (as we noted above), we get $\ell(g) \leq 2k \cdot \ell(f)$. This completes the proof of (4), as we have now constructed a feasible multicommodity flow g in the state space graph of the switch chain with the desired properties. \square

4 Sampling Graphs with a Given JDM

We may use a similar high level approach to that in Section 3 to show that the (restricted) switch chain defined in Subsection 2.2 is *always* rapidly mixing for JDM instances with two degree classes.

Theorem 10. *Let \mathcal{D} be the family of instances of the joint degree matrix model with two degree classes. Then the switch chain is rapidly mixing for instances in \mathcal{D} .*

In analogy to the JS chain we first analyze a simpler Markov chain, called the *hinge flip chain*, that adds and removes (at most) one edge at a time. Very much like the JS chain, the hinge flip chain might slightly violate the degree constraints. Now, however, the joint degree constraints might be violated as well. The definition of *strong stability* is appropriately adjusted to account for both deviations from the original requirements. Finally, we use a similar embedding argument as in Theorem 8. The relevant definitions, as well as the analysis of this auxiliary chain are deferred to Appendix D due to space constraints. Here we present a high level outline of the proof of Theorem 10.

Rapid Mixing of the Hinge Flip Chain. The first step of the proof is to show that the hinge flip chain defined on a *strict superset* of the target state space mixes rapidly for strongly stable instances. Appendix D is dedicated to this step. The fact that we do not want to deviate by more than a constant from the joint degree constraints makes the analysis much more challenging than the one for the JS chain presented in Appendix A.2. In order to overcome the difficulties that arise due to this fact, we rely on ideas introduced by Bhatnagar et al. [4] for uniformly sampling bichromatic matchings. In particular, in the circuit processing part of the proof, we process a circuit at multiple places *simultaneously* in case there is only one circuit in the canonical decomposition of a pairing, or we process multiple circuits *simultaneously* in case the decomposition yields multiple circuits. At the core of this approach lies a variant of the *mountain-climbing problem* [26, 43]. In our case the analysis is more involved than that of [4], and we therefore use different arguments in various parts of the proof.

It is interesting to note that the analysis of the hinge flip chain is not carried out in the JDM model but in the more general Partition Adjacency Matrix (PAM) model [10, 15]. The difference from the JDM model is that in each class V_i the nodes need not have the same constant degree but rather follow a given degree sequence of size $|V_i|$. Given that small deviations from the prescribed degrees cannot be directly handled—by definition—by the JDM model, the PAM model is indeed a more natural choice for this step.

Strong Stability of JDM Instances. Next we show that for any JDM instance, any graph in the state space of the hinge flip chain (i.e., graphs that satisfy or *almost* satisfy the joint degree requirements) can be transformed to a graphical realization of the original instance within 6 hinge flips at most. That is, the set of JDM instances is a strongly stable family of instances of the PAM model and thus the hinge flip chain mixes rapidly for JDM instances. See Theorem 31 in Appendix E.

Flow Transformation. The final step is an embedding argument, along the lines of the argument of Subsection 3.2, for transforming the efficient flow for the hinge flip chain to an efficient flow for the switch chain. As an intermediate step we need an analog of Theorem 9, but this directly follows from the proof of irreducibility of the switch chain in [1]. See Appendix F.

5 Discussion

We believe that our ideas can be also used to simplify the switch chain analyses in settings where there is some given forbidden edge set, the elements of which cannot be used in any (bipartite) graphical realization [23, 25, 17, 18]. This is an interesting direction for future work, as it captures the case of sampling directed graphs.

Moreover, we suspect it can be shown that the condition in (17) is essentially best possible in terms of δ_r and Δ_r in a similar sense as described in [27] for the results in Corollaries 5 and 6.⁹ While this is an interesting question, our goal here is not to give a full bipartite analogue of [27]. Even so, a deeper understanding of when a family of bipartite degree sequences is strongly stable is missing. In particular, is it possible to unify the results of Corollaries 18 and 19 under a single condition similar to (2)?

Further, it is not clear whether there exist degree sequence families—bipartite or not—that are P -stable but not strongly stable. For instance, in a recent work by Gao and Wormald [21], who provide a very efficient non-MCMC approximate sampler for certain power-law degree sequences, it is argued that these power-law degree sequences are P -stable. Is it the case these sequences are strongly stable as well? Theorem 4 would then directly imply that the switch chain is rapidly mixing for this family.

A central open question is how to go beyond (strong) stability. We suspect that the proof template of [7] cannot be used for proving rapid mixing of the switch chain for general families of degree sequences. The intuition is that it relies on the fact that there is a set of auxiliary states that is not much larger than the set of actual graphical realizations for a given degree sequence; this property seems very closely related to P -stability, and also arises explicitly in the analysis of the bipartite case in [30]. This observation suggests the need for a novel approach for studying the mixing of the switch chain on non-stable degree families.

Finally, the problem of sampling graphic realizations of a given joint degree distribution with three or more degree classes is also open. Although our proof breaks down for more than two classes, we hope that our high level approach can facilitate progress on the problem.

Acknowledgements

We are grateful to Péter Erdős, Tamás Mezei and István Miklós for their useful comments.

⁹We suggest using the same ideas as the proof of Theorem 6 in [27] based on a non-stable family of bipartite degree sequences presented in [30].

A Missing Material from Section 3

A.1 On the Proofs of Corollaries 5 and 6

Since a slightly different notion of stability (used in [27]) is defined below, to avoid confusion, we should clarify that whenever *strong* stability is mentioned it is meant in the sense of Definition 2.

We first introduce some notation, using the same terminology as in [27]. Let $G = (V, E)$ be an undirected graph. For distinct $u, v \in V$ we say that u, v are co-adjacent if $\{u, v\} \notin E$, and $\{u, v\}$ is called a co-edge. An alternating path of length q in G is a sequence of nodes v_0, v_1, \dots, v_q such that (i) $\{v_i, v_{i+1}\}$ is an edge when i is even, and a co-edge if i is odd, or (ii) $\{v_i, v_{i+1}\}$ is a co-edge when i is even, and an edge if i is odd. The path is called a cycle if $v_0 = v_q$. We will always specify if a path is of type (i) or (ii).

The definition of P -stability in [27]—which we here call $\pm P$ -stability—is based on the following definition. Let $\mathcal{G}''(d) = \cup_{d'} \mathcal{G}(d')$ with d' ranging over the set

$$\{d\} \cup \left\{ d' : \sum_{i=1}^n d_i = \sum_{i=1}^n d'_i \text{ and } \sum_{i=1}^n |d_i - d'_i| = 2 \right\}.$$

That is, either $d = d'$, or d' has one node with deficit one and one node with surplus one. A family of sequences is $\pm P$ -stable if there exists a polynomial p such that $|\mathcal{G}''(d)|/|\mathcal{G}(d)| \leq p(n)$ for every $d = (d_1, \dots, d_n)$ in the family. In order to show in [27] that the families corresponding to the conditions in Corollaries 5 and 6 are $\pm P$ -stable, it is shown that for any $G \in \mathcal{G}''(d)$ there always exists a *short alternating path* of even length, starting with an edge and ending with a co-edge, connecting the node with surplus one to the node with deficit one. The connection with our definition of strong stability is as follows.

Lemma 11. *Let \mathcal{D} be a family of sequences, and assume that there exists a constant k_0 such that for all $d \in \mathcal{D}$ and any $G \in \mathcal{G}''(d) \setminus \mathcal{G}(d)$, there exists an alternating path of length at most k_0 , starting with an edge and ending with a co-edge, connecting the node with surplus one to the node with deficit one. Then the family \mathcal{D} is strongly stable with respect to $k_0/2 + 1$.*

Proof. Let $H \in \mathcal{G}'(d)$ be such that there are two nodes $u \neq v$ with deficit one (the case of one node with deficit two is very similar). Since u has deficit one, there is some x so that $\{u, x\}$ is a co-edge. Then the graph $H + \{u, x\} \in \mathcal{G}''(d)$ as now x has surplus one, and v still has deficit one (note that if $x = v$ we are immediately done).

The assumption now implies that there is an alternating path of length at most k_0 starting at x with an edge, and ending at v with a co-edge. But this implies that in H , there exists an alternating path starting at u and ending at v of length at most $k_0 + 1$, where both the first and last edge on this path are co-edges. This certainly implies that $k_{JS}(d) \leq k_0/2 + 1$ by using a sequence of (see beginning of Appendix A.2 below for definitions) Type 1 transitions, and finally a Type 2 transition. \square

The following two corollaries now explain why the families presented in Corollaries 5 and 6 are strongly stable. The value of $k_0 = 10$ used below follows directly from the proofs of Theorems 8 and 2 in [27].

Corollary 12. *Let $\mathcal{D} = \mathcal{D}(\delta, \Delta, m)$ be the set of all graphical degree sequences $d = (d_1, \dots, d_n)$ satisfying*

$$(2m - n\delta)(n\Delta - 2m) \leq (\Delta - \delta)[(2m - n\delta)(n - \Delta - 1) + (n\Delta - 2m)\delta] \quad (2)$$

where δ and Δ are the minimum and maximum component of d , respectively, and $m = \frac{1}{2} \sum_{i=1}^n d_i$. For all $d \in \mathcal{D}$ and any $G \in \mathcal{G}''(d) \setminus \mathcal{G}(d)$, there exists an alternating path of length at most 10, starting with an edge and ending with a co-edge, connecting the node with surplus one to the node with deficit one. Hence, \mathcal{D} is strongly stable.

Corollary 13. Let $\mathcal{D} = \mathcal{D}(\delta, \Delta)$ be the set of all graphical degree sequences $d = (d_1, \dots, d_n)$ satisfying

$$(\delta_{\max} - \delta_{\min} + 1)^2 \leq 4\delta_{\min}(n - \delta_{\max} - 1) \quad (3)$$

where δ and Δ are the minimum and maximum component of d , respectively. For all $d \in \mathcal{D}$ and any $G \in \mathcal{G}''(d) \setminus \mathcal{G}(d)$, there exists an alternating path of length at most 10, starting with an edge and ending with a co-edge, connecting the node with surplus one to the node with deficit one. Hence, \mathcal{D} is strongly stable.

A.2 Proof of Theorem 7

For the reader's convenience, we repeat the description of the JS chain [29] here (ignoring the lazy part). We also introduce some shorthand terminology for the type of moves defining the chain: in state $G \in \mathcal{G}'(d)$, select an ordered pair i, j of nodes uniformly at random and then

- (i) if $G \in \mathcal{G}(d)$ and (i, j) is an edge of G , delete (i, j) from G (Type 0 transition),
- (ii) if $G \notin \mathcal{G}(d)$ and the degree of i in G is less than d_i , and (i, j) is not an edge of G , add (i, j) to G ; if this causes the degree of j to exceed d_j , select an edge (j, k) uniformly at random and delete it (Type 1 transition).

In case the degree of j does not exceed d_j in (ii), we call this a Type 2 transition.

Theorem 7. Let \mathcal{D} be a strongly stable family of degree sequences with respect to some constant k . Then there exist polynomials $p(n)$ and $r(n)$ such that for any $d = (d_1, \dots, d_n) \in \mathcal{D}$ there exists an efficient multi-commodity flow f for the JS chain on $\mathcal{G}'(d)$ satisfying $\max_e f(e) \leq p(n)/|\mathcal{G}'(d)|$ and $\ell(f) \leq r(n)$.

We will use the following idea from [28]—used in a different setting—in order to restrict ourselves to establishing flow between states in $\mathcal{G}(d)$, rather than between all states in $\mathcal{G}'(d)$. Assume that d is a degree sequence with n components that is a member of a strongly stable family of degree sequences (with respect to some k).

Lemma 14. Let f' be a flow that routes $1/|\mathcal{G}'(d)|^2$ units of flow between any pair of states in $\mathcal{G}(d)$ in the JS chain, so that $f'(e) \leq b/|\mathcal{G}'(d)|$ for all e in the state space graph of the JS chain. Then f' can be extended to a flow f that routes $1/|\mathcal{G}'(d)|^2$ units of flow between any pair of states in $\mathcal{G}'(d)$ with the property that for all e in the state space graph of the JS chain

$$f(e) \leq q(n) \frac{b}{|\mathcal{G}'(d)|},$$

where $q(\cdot)$ is a polynomial whose degree only depends on $k_{JS}(d)$. Moreover, $\ell(f) \leq \ell(f') + 2k_{JS}(d)$.¹⁰

We now continue with the proof of Theorem 7. It consists of four parts following, in a conceptual sense, the proof template in [7] developed for proving rapid mixing of the switch chain for regular graphs. Certain parts use similar ideas as in [28] where a Markov chain for sampling (near)-perfect matchings is studied. Whenever we refer to [28], the reader is referred to Section 3 of [28].

¹⁰We omit the proof of Lemma 14 as the lemma is actually not needed for proving Theorem 4. Careful consideration of the proof of Theorem 8 shows that we can only focus on flow between states in $\mathcal{G}(d)$, since the flow h given in the proof of Theorem 8 only has positive flow between states corresponding to elements in $\mathcal{G}(d)$. That is, when defining the flow h , we essentially forget about all flow in f between any pair of states where at least one state is an auxiliary state, i.e., an element of $\mathcal{G}'(d) \setminus \mathcal{G}(d)$. Said differently, in Theorem 8 we could start with the assumption that f routes $1/|\mathcal{G}'(d)|^2$ units of flow between any pair of states in $\mathcal{G}(d)$ in the state space graph of the JS chain, and then the transformation still works. However, the formulations of Theorems 7 and 8 are more natural describing a comparison between the JS and switch chains.

Proof of Theorem 7. We only need to define a flow f' as in Lemma 14 so that $b \leq p_1(n)$ and $\ell(f') \leq p_2(n)$ for some polynomials $p_1(\cdot), p_2(\cdot)$ whose degrees may only depend on $k = k_{JS}(d)$. Actually, we are going to show that we may use $p_1(n) = p_2(n) = n^2$. Then the theorem follows from the lemma and the fact that $\ln(|\mathcal{G}'(d)|)$ is upper bounded by a polynomial in n . The latter follows from Equation (1) of McKay and Wormald [33] that implies that

$$|\mathcal{G}(d')| \leq n^{n^2}$$

for any degree sequence d' with n components (see also [25]). So, by the definition of $|\mathcal{G}'(d)|$ we have

$$|\mathcal{G}'(d)| \leq \left(\frac{n(n-1)}{2} + n + 1 \right) n^{n^2},$$

and thus $\ln(|\mathcal{G}'(d)|) \leq 3n^3$.

Before we define f' , we first introduce some basic terminology similar to that in [7]. Let V be a set of labeled vertices, let \prec_E be a fixed total order on the set $\{\{v, w\} : v, w \in V\}$ of edges, and let \prec_C be a total order on all *circuits* on the complete graph K_V , i.e., \prec_C is a total order on the closed walks in K_V that visit every edge at most once. We fix for every circuit one of its vertices where the walk begins and ends.

For given $G, G' \in \mathcal{G}(d)$, let $H = G \Delta G'$ be their symmetric difference. We refer to the edges in $G \setminus G'$ as *blue*, and the edges in $G' \setminus G$ as *red*. A *pairing* of red and blue edges in H is a bijective mapping that, for each node $v \in V$, maps every red edge adjacent to v , to a blue edge adjacent to v . The set of all pairings is denoted by $\Psi(G, G')$, and, with θ_v the number of red edges adjacent to v (which is the same as the number of blue edges adjacent to v), we have $|\Psi(G, G')| = \prod_{v \in V} \theta_v!$.

Canonical Path Description [7]. Similar to the approach in [7], the goal is to construct for each pairing $\psi \in \Psi(G, G')$ a canonical path from G to G' that carries a $|\Psi(G, G')|^{-1}$ fraction of the total flow from G to G' in f' . For notational convenience, for the remaining of the proof we write uw instead of $\{u, v\}$ to denote an edge. For a given pairing ψ and the total order \prec_E given above, we first decompose H into the edge-disjoint union of circuits in a canonical way. We start with the lexicographically smallest edge $w_0 w_1$ in E_H and follow the pairing ψ until we reach the edge $w_k w_0$ that was paired with $w_0 w_1$. This defines the circuit C_1 . If $C_1 = E_H$, we are done. Otherwise, we pick the lexicographically smallest edge in $H \setminus C_1$ and repeat this procedure. We continue generating circuits until $E_H = C_1 \cup \dots \cup C_s$. Note that all circuits have even length and alternate between red and blue edges, and that they are pairwise edge-disjoint. We form a path

$$G = Z_0, Z_1, \dots, Z_M = G'$$

from G to G' in the state space graph of the JS chain, by processing the circuits C_i in turn according to the total order \prec_C . The *processing of a circuit* C is the procedure during which all blue edges on C are deleted, and all red edges of C are added to the current graphical realization, using the three types of transitions in the JS chain mentioned at the beginning of this section. All other edges of the current graphical realization remain unchanged. In general, this can be done similarly to the circuit processing procedure in [28].¹¹

Circuit Processing [28]. Let $C = vx_1 x_2 \dots x_q v$ be a circuit with start node v . We may assume, without loss of generality, that vx_1 is the lexicographically smallest blue edge adjacent to the starting node v . We first perform a type 0 transition in which we remove the blue edge vx_1 . Then we perform a sequence of $\frac{q-1}{2}$ type 1 transitions in which we add the red edge $x_i x_{i+1}$ and remove the blue edge $x_{i-1} x_i$ for $i = 1, 3, \dots, q$. Finally we perform a type 2 transition in which we add the red edge vx_q . In particular, this means that the elements on the canonical path right before and after the processing of a circuit belong to $\mathcal{G}(d)$. It is easy to see that

¹¹This is the main difference between the switch chain analyses [7, 25, 34, 17, 19, 18] and our analysis. The processing of a circuit is much more complicated if performed directly in the switch chain.

all the intermediate elements that we visit during the processing of the circuit C belong to $\mathcal{G}'(d) \setminus \mathcal{G}(d)$, i.e., every element has either precisely two nodes with degree deficit one, or one node with degree deficit two. This is illustrated in Figures 4, 5 and 6 for the circuit in Figure 3.




Figure 3: The circuit $C = vx_1x_2x_3x_4x_5x_6x_7x_8x_9v$ with $v = x_3$ and $x_5 = x_8$. The blue edges are represented by the solid edges, and the red edges by the dashed edges.




Figure 4: The edge vx_1 is removed using a Type 0 transition (left). The edge x_1x_2 is added and $x_2x_3 = x_2v$ is removed using a Type 1 transition (right). We have also indicated the non-zero degree deficits.




Figure 5: The edge x_3x_4 is added and x_4x_5 is removed using a Type 1 transition (left). The edge x_5x_6 is added and x_6x_7 is removed using a Type 1 transition (right).




Figure 6: The edge $x_7x_8 = x_5x_8$ is added and $x_5x_9 = x_8x_9$ is removed using a Type 1 transition (left). The edge vx_9 is added using a Type 2 transition (right).

For the next part, we define the notion of an encoding that can be used to bound the congestion of an edge in the state space graph of the JS chain using an injective mapping argument.

Encoding [28]. Let $t = (Z, Z')$ be a given transition of the Markov chain. Suppose two graphs G and G' use the transition t over some canonical path for some pairing $\psi \in \Psi(G, G')$. Let $H = G \Delta G'$. We define the encoding

$$L_t(G, G') = \begin{cases} (H \Delta (Z \cup Z')) - e_{H,t} & \text{if } t \text{ is a Type 1 transition,} \\ H \Delta (Z \cup Z') & \text{otherwise,} \end{cases}$$

where $e_{H,t}$ is the first blue edge on the circuit that is currently being processed on the canonical path from G to G' (for the given pairing ψ). This encoding is of a similar nature as the encoding used in [28]. An example is given in Figures 7, 8 and 9. We also refer the reader to Figure 1 in [28] for a detailed example.¹² The following lemma is crucial for the analysis.

Lemma 15. *Given $t = (Z, Z')$, L , and ψ , we can uniquely recover G and G' . That is, if L is such that $L_t = L_t(G, G')$ for some pair (G, G') , then (G, G') is the unique pair for which this is the case, given t, L, ψ .*

Proof. We give the proof for when t is a Type 1 transition. The cases of the two other types are similar, and arguably somewhat easier. The proof uses the arguments in [28] interpreted in our setting. First note that $L \Delta (Z \cup Z')$ is a graph in which there are precisely two nodes with odd degree. In particular, the edge $e_{H,t}$ is the unique edge (having as endpoints these odd degree nodes) that has to be added to $L \Delta (Z \cup Z')$ to obtain $H = G \Delta G'$. That is, we have $(L \Delta (Z \cup Z')) + e_{H,t} = H$. The pairing ψ then yields a unique circuit decomposition of $E(H)$ as explained at the beginning of the proof. From the transition t it can be inferred which circuit is currently being processed, and, moreover, we can infer which edges of that circuit belong to G and which to G' . Furthermore, the global ordering $<_C$ on all circuits can then be used to determine for every other circuit whether it has been processed already or not. For every such circuit, we can then infer which edges on it belong to G and which to G' by comparing with Z (or Z'). Therefore, G and G' can be uniquely recovered from t, L and ψ . \square




Figure 7: Symmetric difference $H = G \Delta G'$ where the solid edges represent the edges G and the dashed edges the edges of G' . From left to right the circuit are numbered C_1, C_2 and C_3 , and assume that this is also the order in which they are processed.

¹²Although the perfect matching setting might seem different at first glance, it is actually closely related to our setting, with the only difference that the symmetric difference of two perfect matchings is the union of *node-disjoint cycles*, whereas in our setting the symmetric difference of two graphical realizations is the union of *edge-disjoint circuits*. This is roughly why the notion of pairings is needed, as they allow us to uniquely determine the circuits. That is, the edge-disjoint circuits determined by the pairing are the analogue of the node-disjoint cycles in the perfect matching setting in [28].




Figure 8: The transition $t = (Z, Z')$ that removes the edge x_6x_7 and adds the edge x_5x_6 as part of the processing of C_2 . Note that C_1 has already been processed. The edges in $(E(G) \cup E(G')) \setminus E(H)$ are left out.




Figure 9: The encoding $L = L_t(G, G')$, where again the edges in $(E(G) \cup E(G')) \setminus E(H)$ are left out. Note that in this case $e_{H,t} = vx_1$ and that L is itself an element of $\mathcal{G}'(d)$.

Injective Mapping Argument [7, 28]. We complete the proof by using an injective mapping argument to bound the congestion of the flow f' on the edges of the state space graph of the JS chain. The arguments used are a combination of ideas from [28] and the proof of Lemma 2.5 in [7] (see also Lemma 1 in [8]).¹³ We again focus on Type 1 transitions t as the proofs for the other two types are similar but simpler.

For a tuple (G, G', ψ) , let $p_\psi(G, G')$ denote the canonical path from G to G' for pairing ψ . Let

$$\mathcal{L}_t = \{L_t(G, G') \mid (G, G', \psi) \in \mathcal{F}_t\}$$

be the set of all (distinct) encodings L_t , where

$$\mathcal{F}_t = \{(G, G', \psi) : t \in p_\psi(G, G')\}$$

is the set of all tuples (G, G', ψ) such that the canonical path from G to G' under pairing ψ uses the transition t . A crucial observation is that every encoding $L_t(G, G')$ itself is an element of $\mathcal{G}'(d)$ (see Figure 9 for an

¹³Lemma 2.5 in [7] contained a flaw for which the corrigendum [8] was published.

example). This implies that

$$|\mathcal{L}_t| \leq |\mathcal{G}'(d)|. \quad (10)$$

Moreover, with $H = G \Delta G'$ and $L = L_t(G, G')$, the pairing ψ has the property that it pairs up the edges of $E(H) \setminus E(L)$ and $E(H) \cap E(L)$ in such a way that for every node v (with the exception of at most two nodes) each edge in $E(H) \setminus E(L)$ that is incident to v is paired up with an edge in $E(H) \cap E(L)$ that is incident to v . However, there are either two nodes for which the incident edges in $E(H) \setminus E(L)$ exceed by 2 the incident edges in $E(H) \cap E(L)$, or one node for which the incident edges in $E(H) \setminus E(L)$ exceed by 4 the incident edges in $E(H) \cap E(L)$. These are exactly the two nodes with degree deficit 1 or the one node with degree deficit 2 in L ; for the example in Figure 9 these are nodes x_1 and x_6 . There ψ pairs up each edge of $E(H) \cap E(L)$ to an edge of $E(H) \setminus E(L)$ but also two edges of $E(H) \setminus E(L)$ with each other; or in the case of one node with degree deficit 2 ψ pairs up each edge of $E(H) \cap E(L)$ to an edge of $E(H) \setminus E(L)$ but also makes two pairs out of the remaining 4 edges in $E(H) \setminus E(L)$. Let $\Psi'(L)$ be the set of all pairings with this property.¹⁴ Note that not every such pairing has to correspond to a tuple (G, G', ψ) for which $t \in p_\psi(G, G')$.

By simply counting, we can upper bound $|\Psi'(L)|$ in terms of $|\Psi(H)|$. We show the calculation for the case where L has two nodes with degree deficit 1. The case of one node with degree deficit 2 is very similar and the same upper bound works there as well. Suppose that u, w are the two nodes of L with degree deficit 1. Then

$$|\Psi'(L)| = (\prod_{v \in V \setminus \{u, w\}} \theta_v!) \cdot \frac{(\theta_u + 1)!}{2} \cdot \frac{(\theta_w + 1)!}{2} = |\Psi(H)| \cdot \frac{(\theta_u + 1)(\theta_w + 1)}{4} \leq n^2 \cdot |\Psi(H)|. \quad (11)$$

Putting everything together, we have

$$\begin{aligned} |\mathcal{G}'(d)|^2 f'(e) &= \sum_{(G, G')} \sum_{\psi \in \Psi(G, G')} \mathbf{1}(e \in p_\psi(H)) |\Psi(H)|^{-1} \\ &\leq \sum_{L \in \mathcal{L}_t} \sum_{\psi' \in \Psi'(L)} |\Psi(H)|^{-1} \quad (\text{using Lemma 15}) \\ &\leq n^2 \sum_{L \in \mathcal{L}_t} 1 \quad (\text{using (11)}) \\ &\leq n^2 \cdot |\mathcal{G}'(d)| \quad (\text{using (10)}) \end{aligned}$$

The usage of Lemma 15 for the first inequality works as follows. Every tuple $(G, G', \psi) \in \mathcal{F}_t$ with encoding $L_t(G, G')$ generates a unique tuple in $\{L_t(G, G')\} \times \Psi'(L_t(G, G'))$. But since, by Lemma 15, we can uniquely recover G and G' from L , t and ψ , we have that $\sum_{L \in \mathcal{L}_t} |\{L\} \times \Psi'(L)| = \sum_{L \in \mathcal{L}_t} \sum_{\psi' \in \Psi'(L)} 1$ is an upper bound on the number of canonical paths that use t .

By rearranging (11) we get the upper bound for f' required in Lemma 14. What is left to show is that $\ell(f')$ is not too large. This, however, is determined by the way we defined the canonical paths. It is easy to see that for any canonical path between any two graphs $G, G' \in \mathcal{G}(d)$ has length at most $\frac{3}{4}|E(G \Delta G')|$ and, therefore, $\ell(f') \leq n^2$. \square

Remark 16 (Bipartite case). The proof for the bipartite case is very similar. The only difference is that in the circuit processing procedure there will never be an auxiliary state where one node has degree deficit two. For this to occur there necessarily has to be a simple cycle of odd length in a circuit, see, e.g., Figure 4. This explains the adjusted definition of $\mathcal{G}'(r, c)$ used in the definition of (strong) stability for the bipartite case in Appendix B. Moreover, the exact same encoding and injective mapping arguments can be used.

¹⁴Remember that we do not need to know G and G' in order to determine the set H . It can be found based on L and the transition $t = (Z, Z')$, as described in the proof of Lemma 15.

B Sampling Bipartite Graphs

A similar analysis as in Appendix A.2 holds for the bipartite case, under some slightly adjusted definitions. Throughout this section, we use m and n to denote the number of vertices in the two independent sets of a given bipartition.

For a given tuple (r, c) with $r = (r_1, \dots, r_m)$ and $c = (c_1, \dots, c_n)$, we define $\mathcal{G}(r, c)$ as the set of all bipartite graphs $G = (V \cup U, E)$ with $V = \{1, \dots, m\}$ and $U = \{1, \dots, n\}$, where the nodes in V have degree sequence r , and the nodes in U degree sequence c . The set $\mathcal{G}'(r, c)$ is defined as $\cup_{(r', c')} \mathcal{G}(r', c')$ where we either have $(r', c') = (r, c)$, or, there exist $x \in \{1, \dots, m\}$ and $y \in \{1, \dots, n\}$ such that

$$r'_i = \begin{cases} r_i - 1 & \text{if } i = x, \\ r_i & \text{otherwise,} \end{cases} \quad \text{and} \quad c'_j = \begin{cases} c_j - 1 & \text{if } j = y, \\ c_j & \text{otherwise.} \end{cases}$$

Note that the case in which there is one node with degree deficit two cannot occur here. The definitions of the JS and switch chain, as well as that of the parameter in (1), denoted by $k_{JS}(r, c)$ here, and the notions of stability and strong stability, are very similar. In particular, in the JS chain the only difference is that we now pick $i \in U$ and $j \in V$ uniformly at random instead of picking $i, j \in U \cup V$ uniformly at random.

Crucially, as is stated in Theorem 17 below, the result in Theorem 4 carries over. The flow transformation in Theorem 8 is completely analogous for the bipartite case. The *merging* and *rerouting* procedures can be carried out in exactly the same fashion, and instead of Theorem 9 one can use the corresponding result for the bipartite case, see, e.g., [37, 16]. Theorem 7 is also true for the bipartite case; this is discussed in Remark 16 in Appendix A.2.

Theorem 17 (Bipartite case). *Let \mathcal{D} be a strongly stable family of bipartite degree sequences with respect to some constant k . Then there exists a polynomial $q(n, m)$ such that, for any $0 < \epsilon < 1$, the mixing time τ_{sw} of the switch chain for a graphical sequence $(r, c) \in \mathcal{D}$, with $r = (r_1, \dots, r_m)$ and $c = (c_1, \dots, c_n)$, satisfies*

$$\tau_{sw}(\epsilon) \leq q(n, m)^k \ln(1/\epsilon).$$

Similarly, we can obtain a constant upper bound on $k_{JS}(r, c)$ for certain degree families. We first present a range of sequences that can be considered a bipartite counterpart of Corollary 6.

Corollary 18. *Let $\mathcal{D} = \mathcal{D}(\delta_r, \Delta_r, \delta_c, \Delta_c)$ be the set of all graphical bipartite degree sequences (r, c) on m and n nodes respectively, satisfying*

$$(\Delta_r - \delta_c)^2 \leq 4\delta_c(n - \Delta_r) \quad \text{and} \quad (\Delta_c - \delta_r)^2 \leq 4\delta_r(m - \Delta_c) \quad (12)$$

where δ_r, Δ_r are the minimum and maximum component of r , and δ_c, Δ_c the minimum and maximum component of c respectively. For any $(r, c) \in \mathcal{D}$, we have $k_{JS}(r, c) \leq 8$. Hence the switch chain is rapidly mixing for sequences in \mathcal{D} .

The proof is a bipartite analogue of Lemma 1 in [27]. In particular, we use the notions of co-edges and alternating paths as defined in Appendix A.1.

Proof. Suppose that (r', c') is such that there exist precisely one $s \in V$ and $t \in U$ with degree deficit one, and all other nodes have no deficit. Let G be a realization for the sequence (r', c') if $\mathcal{G}(r', c') \neq \emptyset$.

As both s and t have a deficit, there exist $a \in U$ and $b \in V$ such that $\{s, a\}$ and $\{t, b\}$ are co-edges in G . If either $a = t$ or $b = s$, we are done. Therefore, we may assume that $\{a, b\}$ is a co-edge in G . Now consider the graph $G' = G + \{s, a\} + \{t, b\}$. Let

$$X_1 = \{v \mid \{a, v\} \in G'\} \subseteq V$$

be the set of all neighbors of a , and similarly X_2 the set of all neighbors of b in U . We may assume that the sets X_1 and X_2 form a complete bipartite graph in G' , otherwise there is an alternating path of length 3 and we are done. Define Y_i as the set of nodes with at least one co-edge in X_{i+1} . We may assume that the sets Y_1 and Y_2 form an independent set in G' , otherwise there is an alternating path of length 5 and we are done. Furthermore, let Z_i be the set of nodes not in X_i with a neighbor in Y_{i+1} . Also here, it must be the case that Z_1 and Z_2 form a complete bipartite graph, otherwise there is an alternating path of length 7 and we are done. Moreover, by definition of Y_i , all nodes in Z_i are connected to all nodes in X_{i+1} , and so the sets $K_1 = X_1 \cup Z_1$ and $K_2 = X_2 \cup Z_2$ form a complete bipartite graph. Finally, we let $R_1 = V \setminus (K_1 \cup Y_1)$ and $R_2 = U \setminus (K_2 \cup Y_2)$.

Before we continue, let us introduce some additional notation. For subsets $A \subseteq V$ and $B \subseteq U$ we write $E(A, B)$ for the set of edges with one endpoint in A and one endpoint in B . Moreover, for all sets that were introduced, lower case letters are used to represent the cardinalities of these sets, e.g., $y_1 = |Y_1|$. We will upper and lower bound the quantity $|E(K_1, R_2 \cup Y_2)|$.

Now, by definition all neighbors of nodes in Y_2 are part of K_1 , so that

$$|E(K_1, Y_2)| \geq \delta_c \cdot y_2.$$

Also, by definition all nodes in R_1 are adjacent to all nodes in X_1 , and, in combination with the fact that $|X_1| \geq \delta_c + 1$ (as t has a degree surplus of one), this implies that

$$|E(K_1, R_2)| \geq |E(X_1, R_2)| \geq (\delta_c + 1) \cdot r_2.$$

Combining these results, we find that

$$|E(K_1, R_2 \cup Y_2)| > \delta_c(r_2 + y_2) = \delta_c(n - k_2). \quad (13)$$

We now continue with an upper bound. As K_1 and K_2 form a complete bipartite graph, we have

$$|E(K_1, R_2 \cup Y_2)| \leq k_1(\Delta_r - k_2). \quad (14)$$

Combining (13) and (14) we find

$$\delta_c(n - k_2) < k_1(\Delta_r - k_2). \quad (15)$$

Similarly, by interchanging the roles of U and V , we find

$$\delta_r(m - k_1) < k_2(\Delta_c - k_1). \quad (16)$$

We now make a case distinction depending on the sizes of k_1 and k_2 .

Case 1: $k_1 \leq k_2$. Then (15) implies that

$$\delta_c(n - k_2) < k_2(\Delta_r - k_2),$$

which can be rewritten as

$$f(k_2) := k_2^2 - (\Delta_r + \delta_c)k_2 + \delta_c n < 0.$$

This implies that the function $f(k)$ has two null points, and therefore its discriminant must be strictly greater than zero. However, its discriminant is equal to $(\Delta_r + \delta_c)^2 - 4\delta_c n = (\Delta_r - \delta_c)^2 - 4\delta_c(n - \Delta_r) \leq 0$ which is a contradiction. Recall that the last inequality holds by assumption.

Case 2: $k_2 \leq k_1$. We can repeat the analysis of Case 1 but based on (16) instead of (15). \square

If $\Delta_r = \Delta_c$, $\delta_r = \delta_c$ and $m = n$, the condition in (12) reduces to

$$(\Delta_r - \delta_r)^2 \leq 4\delta_r(m - \Delta_r). \quad (17)$$

In particular, condition (17) is satisfied if $\delta_r \geq m/4$ and $\Delta_r \leq 3m/4 - 1$. This strictly improves the range of $\delta_r \geq m/3$ and $\Delta_r \leq 2m/3$ for the bipartite case recently obtained in [18]. It should be noted that Corollary 18 does not capture the (almost) half-regular case [34, 19].¹⁵ Those cases can be captured by the next result. In [18] it was recently shown that the switch chain is rapidly mixing for the condition in (18). We now provide an alternative proof for this fact by showing that these bipartite degree sequences are strongly stable.

Corollary 19. *Let $\mathcal{D} = \mathcal{D}(\delta_r, \Delta_r, \delta_c, \Delta_c)$ be the set of all graphical bipartite degree sequences (r, c) on m and n nodes respectively, satisfying*

$$(\Delta_c - \delta_c - 1)(\Delta_r - \delta_r - 1) < 1 + \max\{\delta_c(n - \Delta_r), \delta_r(m - \Delta_c)\} \quad (18)$$

where δ_r, Δ_r are the minimum and maximum component of r , and δ_c, Δ_c the minimum and maximum component of c respectively. For any $(r, c) \in \mathcal{D}$, we have $k_{JS}(r, c) \leq 8$. Hence the switch chain is rapidly mixing for sequences in \mathcal{D} .

Proof. The proof bears similarities with the the proof of Corollary 18, especially in the beginning. In fact, let $s, t, G, a, b, G', X_1, X_2, Y_1$ and Y_2 be exactly as in the proof of Corollary 18. Also, let $R_1 = V \setminus (X_1 \cup Y_1 \cup \{b\})$ and $R_2 = U \setminus (X_2 \cup Y_2 \cup \{a\})$ and note that, by the definition of Y_1, R_1 and X_2 form a complete bipartite graph in G' ; the same is true for X_1 and R_2 .

Before we continue, we introduce some additional notation. For subsets $A \subseteq V$ and $B \subseteq U$ we write $\bar{E}(A, B)$ for the set of co-edges in G' with one endpoint in A and one endpoint in B . Like in the proof of Corollary 18, for all sets that were introduced, lower case letters are used to represent the cardinalities of these sets, e.g., $y_1 = |Y_1|$.

We will upper and lower bound the quantity $|\bar{E}(Y_1, U)|$. The fact that the degree of each node in Y_1 is at least δ_r , gives us an obvious upper bound:

$$|\bar{E}(Y_1, U)| \leq y_1(n - \delta_r). \quad (19)$$

Since Y_1 and Y_2 form an independent set and, by definition, no node in Y_1 is adjacent to a , we have

$$|\bar{E}(Y_1, U)| = y_1(y_2 + 1) + |\bar{E}(Y_1, U - Y_2 - \{a\})|. \quad (20)$$

So, we need to lower bound the quantity $|\bar{E}(Y_1, U - Y_2 - \{a\})| = |\bar{E}(Y_1, X_2 \cup R_2)|$. Using straightforward upper and lower bounds and the fact that $|\bar{E}(X_1, R_2)| = |\bar{E}(X_1 \cup R_1 \cup \{b\}, X_2)| = 0$,¹⁶ we have

$$\begin{aligned} |\bar{E}(Y_1, X_2 \cup R_2)| &= |\bar{E}(V, X_2 \cup R_2)| - |\bar{E}(X_1 \cup R_1 \cup \{b\}, X_2 \cup R_2)| \\ &\geq (m - \Delta_c)(n - y_2 - 1) - |\bar{E}(X_1 \cup R_1 \cup \{b\}, X_2)| - |\bar{E}(X_1, R_2)| - |\bar{E}(R_1 \cup \{b\}, R_2)| \\ &\geq (m - \Delta_c)(n - y_2 - 1) - (r_1 + 1)r_2. \end{aligned} \quad (21)$$

Now we may combine (19), (20) and (21) to get

$$y_1(n - \delta_r) \geq y_1(y_2 + 1) + (m - \Delta_c)(n - y_2 - 1) - (r_1 + 1)r_2,$$

and by rearranging terms

$$(r_1 + 1)r_2 \geq (y_2 + 1 - n)(y_1 - m + \Delta_c) + \delta_r y_1. \quad (22)$$

Next we derive bounds on the sizes of the different sets involved in the proof so far. First note that a has at least $\delta_c + 1$ neighbors as it has degree surplus of one. Thus, $x_1 \geq \delta_c + 1$. Similarly, $x_2 \geq \delta_r + 1$. To bound y_1 , notice that a node in X_2 has at least $m - \Delta_c$ non-neighbors in V and all of them must be in Y_1 . Thus,

¹⁵A pair (r, c) is *half-regular* if either r of c is regular, and *almost half-regular* if either $\Delta_r \leq \delta_r + 1$ or $\Delta_c \leq \delta_c + 1$.

¹⁶This follows from the definitions of the sets X, Y and R , and the fact that X_1 and X_2 form a complete bipartite graph.

$y_1 \geq m - \Delta_c$. On the other hand, $Y_1 \cup R_1$ contains the non-neighbors of a (except from b) and these can be at most $m - \delta_c - 2$. Thus, $y_1 \leq m - \delta_c - r_1 - 2$. Similarly, $n - \Delta_r \leq y_2 \leq n - \delta_r - r_2 - 2$.

We are going to combine the latter bounds with (22):

$$(n - \delta_r - 2 - y_2)(m - \delta_c - 1 - y_1) \geq (y_2 + 1 - n)(y_1 - m + \Delta_c) + \delta_r y_1.$$

By multiplying everything out and canceling several terms, we get

$$-n\delta_c - n - m\delta_r + \delta_r\delta_c + \delta_r - m + 2\delta_c + 2 + y_1 + \delta_c y_2 + y_2 \geq \Delta_c(y_2 + 1 - n),$$

and by rearranging

$$(n - y_2 - 1)(\Delta_c - \delta_c - 1) + (y_1 - m + \delta_c + 1) \geq \delta_r(m - \delta_c - 1).$$

Next we may use $y_2 \geq n - \Delta_r$ and $y_1 \leq m - \delta_c - 2$ to finally get

$$(\Delta_r - 1)(\Delta_c - \delta_c - 1) - 1 \geq \delta_r(m - \Delta_c) + \delta_r(\Delta_c - \delta_c - 1),$$

which is equivalent to

$$(\Delta_r - \delta_r - 1)(\Delta_c - \delta_c - 1) \geq 1 + \delta_r(m - \Delta_c).$$

Using the exact same arguments, we can work with $|\bar{E}(V, Y_2)|$ to get

$$(\Delta_c - \delta_c - 1)(\Delta_r - \delta_r - 1) \geq 1 + \delta_c(n - \Delta_r).$$

The last two inequalities contradict the choice of (r, c) . □

C The JDM Model: An Example

In this section we provide an example of the joint degree matrix model with two degree classes. We let $V = \{1, \dots, 11\}$ and consider the partition given by $V_1 = \{1, \dots, 6\}$ and $V_2 = \{7, \dots, 11\}$, i.e., $q = 2$. We let

$$c = \begin{pmatrix} 7 & 4 \\ 4 & 8 \end{pmatrix} \quad \text{and} \quad d = (3, 4).$$

This means that the nodes in V_1 have degree three, the nodes in V_2 degree four, and there are in total four edges between the nodes of V_1 and V_2 . In Figure 10 below we give a possible graphical realization of this tuple (c, d) .




Figure 10: An example of a graphical realization for c and d as given above.

D The PAM Model and the Hinge Flip Markov Chain

We start with a general description of the Partition Adjacency Matrix (PAM) model as a generalization of the JDM model. Let $V = \{1, \dots, n\}$ be a given set. An instance of the partition adjacency matrix model is given by a partition $V_1 \cup V_2 \cup \dots \cup V_q$ of V into pairwise disjoint classes. Moreover, we are given a symmetric *partition adjacency matrix* $c = (c_{ij})_{i,j \in [q]}$ of non-negative integers, and a sequence $d = (d_1, \dots, d_n)$ of non-negative integers. We say that the tuple $((V_i)_{i \in [q]}, c, d)$ is graphical if there exists a simple, undirected, labelled graph $G = (V, E)$ on the nodes in V with node $i \in V$ having degree d_i , and so that there are precisely c_{ij} edges between endpoints in V_i and V_j . This is denoted by $E[V_i, V_j] = c_{ij}$. The graph G is called a graphical realization of the tuple. We let $\mathcal{G}((V_i)_{i \in [q]}, c, d)$ denote the set of all graphical realizations of the tuple $((V_i)_{i \in [q]}, c, d)$. We often write $\mathcal{G}(c, d)$ instead of $\mathcal{G}((V_i)_{i \in [q]}, c, d)$ when it is clear what the partition is.

In this work we focus on the case of a partition into two classes V_1 and V_2 , and, without loss of generality, assume that $1 \leq c_{12} \leq |V_1| \cdot |V_2| - 1$.¹⁷ For the case of two classes an initial state can be computed in polynomial time [15].¹⁸ We let $\mathcal{G}'(c, d) = \cup_{(c', d')} \mathcal{G}'(c', d')$ with (c', d') ranging over tuples satisfying

- (i) $\sum_{i=1}^n d_i - d'_i = 0$,
- (ii) $\sum_{i=1}^n |d_i - d'_i| \in \{0, 2, 4\}$,
- (iii) $c'_{12} \in \{c_{12} - 1, c_{12}, c_{12} + 1\}$.

We call elements in $\mathcal{G}'(c, d) \setminus \mathcal{G}(c, d)$ *perturbed (auxiliary) states*. For any $G \in \mathcal{G}'(c, d)$ the *perturbation at node* $v \in V$ is defined as $\alpha_v = d_v - d'_v$ where d' is the degree sequence of G . We say that the node v has a degree surplus if $\alpha_v > 0$ and a degree deficit if $\alpha_v < 0$. Moreover, the total degree surplus is defined as $\sum_{v: \alpha_v > 0} \alpha_v$, and the total degree deficit as $-\sum_{v: \alpha_v < 0} \alpha_v$. Note that

$$\sum_{v: \alpha_v > 0} \alpha_v - \sum_{v: \alpha_v < 0} \alpha_v = \sum_{i=1}^n |d_i - d'_i|.$$

Finally, we say that a tuple (c', d') is *edge-balanced* if $c' = c$ (but possibly $d' \neq d$). From the conditions defining $\mathcal{G}'(c, d)$, we may infer the following properties.

Proposition 20. *For any $G \in \mathcal{G}'(c, d)$, for some tuple (c', d') , it holds that*

- (a) *the perturbation at node v satisfies $\alpha_v \in \{-2, -1, 0, 1, 2\}$ for any $v \in V$,*
- (b) *$\max_{i,j=1,2} |c_{ij} - c'_{ij}| \leq 1$, and $\sum_{1 \leq i < j \leq 2} |c_{ij} - c'_{ij}| \in \{0, 2\}$.*

Proof. If there is some node with degree surplus greater or equal than three, then the total degree deficit is also at least three, which follows from the first condition defining $\mathcal{G}'(c, d)$. This means that $\sum_{i=1}^n |d_i - d'_i| \geq 6$, which violates the second condition defining $\mathcal{G}'(c, d)$. A similar argument holds in case there is some node with degree deficit greater or equal than three. To see that the second property is true, assume without loss of generality that $c'_{11} \geq c_{11} + 2$ (similar arguments hold for c'_{22}). Because of the fact that $c'_{12} \in \{c_{12} - 1, c_{12}, c_{12} + 1\}$, by the third condition defining $\mathcal{G}'(c, d)$, it must be that the total degree surplus of the nodes in V_1 is at least three. This gives a contradiction for similar reasons as before. An analogous argument holds in case $c'_{11} \leq c_{11} - 2$. Finally, the last property is a direct consequence of $\max_{i,j=1,2} |c_{ij} - c'_{ij}| \leq 1$ and the fact that $\sum_{i,j=1,2} |c_{ij} - c'_{ij}|$ is an even number, because of the first property defining $\mathcal{G}'(c, d)$. \square

¹⁷It is not hard to see that the cases $c_{12} \in \{0, |V_1| \cdot |V_2|\}$ reduce to the single class case.

¹⁸For general instances, it is not known if an initial state can be computed in time polynomial in n . It is conjectured to be NP-hard in general [15], see also [12].

Remark 21. As we focus on the case in which V is partitioned into two classes V_1 and V_2 here, we will use some shorthand notation in this section.¹⁹ Given a sequence d the number $\gamma = c_{12}$ uniquely determines the matrix c , and the set $\mathcal{G}(c, d)$ is then denoted by $\mathcal{G}(V_1, V_2, \gamma, d)$. As before, we often leave out V_1 and V_2 from the tuple (for sake of readability). That is, we then write $\mathcal{G}'(\gamma, d)$ instead of $\mathcal{G}'(V_1, V_2, c, d)$. We do this in order to emphasize that the derived results only hold for two classes.

We define the *hinge flip Markov chain*²⁰ $\mathcal{M}(\gamma, d)$ on $\mathcal{G}'(\gamma, d)$ as follows. Suppose the current state of the chain is $G \in \mathcal{G}'(\gamma, d)$:

- With probability $1/2$, do nothing.
- Otherwise, perform a *hinge flip* operation: select an ordered triple i, j, k of nodes uniformly at random. If $\{i, j\} \in E(G)$, $\{j, k\} \notin E(G)$, and $G - \{i, j\} + \{j, k\} \in \mathcal{G}'(\gamma, d)$, then delete $\{i, j\}$ and add $\{j, k\}$.

Note that we can check if $G - \{i, j\} + \{j, k\} \in \mathcal{G}'(\gamma, d)$ in time polynomial in n based on the state G .




Figure 11: Example of a hinge flip operation for the ordered triple i, j, k .

Graphs $G, G' \in \mathcal{G}'(\gamma, d)$ are said to be *adjacent* in \mathcal{M} if G can be obtained from G' with positive probability in one transition of the chain \mathcal{M} . We say that two graphs G, G' are within distance r in \mathcal{M} if there exists a path of at most length r from G to G' in the state space graph of \mathcal{M} . By $\text{dist}(G, \gamma, d)$ we denote the minimum distance of G from an element in $\mathcal{G}(\gamma, d)$. The following parameter is the analog of (1) for the current setting and will be used in a similar manner to define the appropriate variant of strong stability. We define

$$k(\gamma, d) = \max_{G \in \mathcal{G}'(\gamma, d)} \text{dist}(G, \gamma, d). \quad (23)$$

In the PAM model with two degree classes, a family \mathcal{D} of graphical tuples (γ, d) is called *strongly stable* if there exists a constant k such that $k(\gamma, d) \leq k$ for all $(\gamma, d) \in \mathcal{D}$.

Theorem 22. *Let \mathcal{D} be a family of graphical tuples that is strongly stable with respect to some constant k . Then for every $(\gamma, d) \in \mathcal{D}$, the chain $\mathcal{M}(\gamma, d)$ is irreducible, aperiodic and symmetric, and, hence, has uniform stationary distribution over $\mathcal{G}'(\gamma, d)$. Moreover, $P(G, G')^{-1} \leq n^3$ for all adjacent $G, G' \in \mathcal{G}'(\gamma, d)$, and also the maximum in- and out-degrees of the state space graph of the chain $\mathcal{M}(\gamma, d)$ are bounded by n^3 .*²¹

Proof. The only claim that requires a detailed argument, and uses the assumption of strong stability, is that of the irreducibility of the chain. By definition of strong stability, we always know that every perturbed state is connected to some element in $\mathcal{G}(\gamma, d)$ so it suffices to show that there is a path between any two states in $\mathcal{G}(\gamma, d)$. This follows from the analysis in the remainder of this section. Aperiodicity follows from the holding probability in the description of the chain \mathcal{M} , and symmetry is straightforward. The bound on $P(G, G^{-1})$ follows directly from the description of the chain, as well as the bound on the in- and out-degrees of the state space graph. \square

¹⁹Still, the notation in Proposition 20 is more convenient in Appendix E.

²⁰One could also choose a natural generalization of the JS chain here. Similarly, for the PAM model with one class as studied in Section 3, one could also use a hinge flip based Markov chain. We choose to work with the JS chain in Section 3 as this is the formulation introduced in [29] (and already claimed to be rapidly mixing for various degree sequences).

²¹It might be the case that the chain is always irreducible, even if \mathcal{D} is not strongly stable, but this is not relevant at this point. The assumption of strong stability allows for a shortcut in the proof of irreducibility.

The remainder of this section is devoted to proving Theorem 23 based on ideas introduced in [4]. Throughout this section we always consider tuples (γ, d) coming from strongly stable families.

Theorem 23. *Let \mathcal{D} be a strongly stable family of tuples (γ, d) with respect to some constant k . Then there exist polynomials $p(n), r(n)$ such that for any $(\gamma, d) \in \mathcal{D}$, with $d = (d_1, \dots, d_n)$, there exists an efficient multicommodity flow f for the auxiliary chain $\mathcal{M}(\gamma, d)$ on $\mathcal{G}'(\gamma, d)$ satisfying $\max_e f(e) \leq p(n)/|\mathcal{G}'(\gamma, d)|$ and $\ell(f) \leq r(n)$. Hence, the chain $\mathcal{M}(\gamma, d)$ is rapidly mixing for families of strongly stable tuples.*

We will use the following lemma in order to simplify the proof of Theorem 23. It is the analog of Lemma 14 in Subsection A.2.

Lemma 24. *Let f' be a flow that routes $1/|\mathcal{G}'(\gamma, d)|^2$ units of flow between any pair of states in $\mathcal{G}(\gamma, d)$ in the chain $\mathcal{M}(\gamma, d)$, so that $f'(e) \leq b/|\mathcal{G}'(\gamma, d)|$ for all e in the state space graph of $\mathcal{M}(\gamma, d)$. Then f' can be extended to a flow f that routes $1/|\mathcal{G}'(\gamma, d)|^2$ units of flow between any pair of states in $\mathcal{G}'(\gamma, d)$ with the property that for all e*

$$f(e) \leq q(n) \frac{b}{|\mathcal{G}'(\gamma, d)|}, \quad (24)$$

where $q(\cdot)$ is a polynomial whose degree only depends on $k(\gamma, d) (\leq k)$. Moreover, $\ell(f) \leq \ell(f') + 2k(\gamma, d)$.

Proof (sketch). We extend the flow f' to f as follows. For any $G \in \mathcal{G}'(\gamma, d) \setminus \mathcal{G}(\gamma, d)$ fix some $\phi(G) \in \mathcal{G}(\gamma, d)$ within distance k of G (which exists by assumption of strong stability), and fix some path in the state space graph from G to $\phi(G)$. Moreover, define $\phi(H) = H$ for all $H \in \mathcal{G}(\gamma, d)$. The flow between G and any given $G' \in \mathcal{G}'(\gamma, d)$ is now sent as follows.

First route $1/|\mathcal{G}'(\gamma, d)|^2$ units of flow from G to $\phi(G)$ over the fixed path from G to $\phi(G)$. Then use the flow-carrying paths used to send $1/|\mathcal{G}'(\gamma, d)|^2$ units of flow between $\phi(G)$ and $\phi(G')$ as in the flow f' (note that in general multiple paths might be used for this in the flow f'). Finally, use the reverse of the fixed path from G' to $\phi(G')$ to route $1/|\mathcal{G}'(\gamma, d)|^2$ from $\phi(G')$ to G' . For any $H \in \mathcal{G}(\gamma, d)$, we have $|\phi^{-1}(H)| \leq \text{poly}(n^k)$, as the in- and out-degrees of the nodes in the state space graph of $\mathcal{M}(\gamma, d)$ are polynomially bounded. It can then be shown (left to the reader) that this extension of f' , yielding the flow f , only gives an additional term of at most $\text{poly}(n^k) \frac{b}{|\mathcal{G}'(\gamma, d)|}$ to the congestion of every arc in the state space graph of the chain $\mathcal{M}(\gamma, d)$ in the flow f' . Hence, the extended flow f satisfies (24) for some appropriately chosen polynomial $q(n)$. \square

D.1 Proof of Theorem 23

Because of Lemma 24 it now suffices to show that there exists a flow f' that routes $1/|\mathcal{G}'(\gamma, d)|^2$ units of flow between any two pair of states in $\mathcal{G}(\gamma, d)$, in the state space graph of the chain $\mathcal{M}(\gamma, d)$, with the property that $f'(e) \leq p(n)/|\mathcal{G}'(\gamma, d)|$, and $\ell(f') \leq q(n)$ for some polynomials $p(\cdot), q(\cdot)$ whose degrees may only depend on $k(\gamma, d)$. Note that f' is not a feasible multi-commodity flow as defined in Section 2, but should rather be interpreted as an intermediate auxiliary flow. The proof of Theorem 23 will consist of multiple parts following, conceptually, the proof template in [7] developed for proving rapid mixing of the switch chain for regular graphs. The main difference is that for the so-called *canonical paths* between states we rely on ideas introduced in [4].

D.1.1 Canonical Paths

We first introduce some basic terminology similar to that in [7]. Let V be a set of labeled vertices, let $<_E$ be a fixed total order on the set $\{\{v, w\} : v, w \in V\}$ of edges, and let $<_C$ be a total order on all *circuits* on the complete graph K_V , i.e., $<_C$ is a total order on the closed walks in K_V that visit every edge at most once. We fix for every circuit one of its vertices where the walk begins and ends.

For given $G, G' \in \mathcal{G}(\gamma, d)$, let $H = G \Delta G'$ be their symmetric difference. We refer to the edges in $G \setminus G'$ as *blue*, and the edges in $G' \setminus G$ as *red*. A *pairing* of red and blue edges in H is a bijective mapping that, for each node $v \in V$, maps every red edge adjacent to v , to a blue edge adjacent to v . The set of all pairings is denoted by $\Psi(G, G')$, and, with θ_v the number of red edges adjacent to v (which is the same as the number of blue edges adjacent to v), we have $|\Psi(G, G')| = \prod_{v \in V} \theta_v!$.

Remember that we are considering an instance of the PAM problem with two classes V_1 and V_2 . For a given graphical realization $G \in \mathcal{G}(\gamma, d)$ we say that $e \in E(G)$ is a *cut edge* if it has an endpoint in both V_1 and V_2 . Otherwise we say that e is an *internal edge*, as both endpoints either lie both in the class V_1 or both in class V_2 .

Similar to the approach in [7], the goal is to construct for each pairing $\psi \in \Psi(G, G')$ a canonical path from G to G' that carries a fraction $|\Psi(G, G')|^{-1}$ of the total flow from G to G' in f' . For a given pairing ψ and the total order $<_E$ given above, we first decompose H into the edge-disjoint union of circuits in a canonical way. We start with the lexicographically least edge $w_0 w_1$ in E_H and follow the pairing ψ until we reach the edge $w_k w_0$ that was paired with $w_0 w_1$. This defines the circuit C_1 (which is indeed a closed walk). If $C_1 = E_H$, we are done. Otherwise, we pick the lexicographically least edge in $H \setminus C_1$ and repeat this procedure. We continue generating circuits until $E_H = C_1 \cup \dots \cup C_s$. Note that all circuits have even length and alternate between red and blue edges, and that they are pairwise edge-disjoint.

We form a path from G to G' in the state space graph of the chain $\mathcal{M}(\gamma, d)$ by changing the blue edges of G into the red edges of G' using hinge flip operations. For certain pairings this can be done in a straightforward way, but in general this is not the case. As a warm-up, we first consider a simple case (this case essentially describes how we would process the circuits in case there is only one class).

Warm-up Example. If for every i , the circuit C_i exclusively consist of internal edges, only within V_1 or only within V_2 , or exclusively of cut edges, circuits can be processed according to the ordering $<_C$ as follows. Let $C = x_0 x_1 x_2 \dots x_q x_0$ be a circuit, and assume w.l.o.g. that $x_0 x_1$ is the lexicographically smallest blue edge adjacent to the starting node x_0 of the circuit. The processing of C now consists of performing a sequence of hinge flips on the ordered pairs (x_{i-1}, x_i, x_{i+1}) for $i = 1, \dots, q$ with the convention that $x_{q+1} = x_0$. This is illustrated in Figures 12, 13 and 14 for an example of C as illustrated in Figure 12 on the left.²² We have also indicated the degree surplus and deficit at every step. By assumption, the edges of C either are all internal edges, or all cut edges. Therefore, throughout the processing of C , we never violate the constraint that there should be γ edges between the classes V_1 and V_2 , and, in particular, this implies that every intermediate state is an element of $\mathcal{G}'(\gamma, d)$.




Figure 12: The circuit $C = x_0 x_1 x_2 x_3 x_4 x_5 x_6 x_7 x_8 x_9 x_0$ with $x_0 = x_3$ and $x_5 = x_8$. The blue edges are represented by the solid edges, and the red edges by the dashed edges (left). The edge $x_0 x_1$ is removed and $x_1 x_2$ is added (right).

²²This is similar to the procedure described in Figures 4, 5 and 6 for the JS chain (we give an example here as well for the sake of completeness).




Figure 13: The edge x_2x_3 is removed and x_3x_4 is added (left). The edge x_4x_5 is removed and x_5x_6 is added (right).




Figure 14: The edge x_6x_7 is removed and x_7x_8 is added (left). The edge x_8x_9 is removed and x_9x_0 is added (right).

In general, however, it might happen that circuits contain both cut and internal edges, in which case we cannot use the circuit processing procedure explained above, as the processing of a circuit might result in a graphical realization for which the number of edges between the classes V_1 and V_2 lies outside the set $\{\gamma - 1, \gamma, \gamma + 1\}$. The latter condition is necessary for the intermediate states in the circuit processing procedure to be elements of $\mathcal{G}'(\gamma, d)$, by definition of that set. In order to overcome the issue described above, we will use the ideas in [4], and process a circuit at multiple places *simultaneously* in case there is only one circuit in the canonical decomposition of a pairing, or, process multiple circuits *simultaneously* in case the decomposition yields multiple circuits. At the core of this approach lies (a variation of) the *mountain-climbing problem* [26, 43]. We begin with introducing this problem, and afterwards continue with the description of the circuit processing procedure, based on the solution to the mountain climbing problem.

Intermezzo: mountain climbing problem. We first introduce some notation and terminology. For non-negative integers $a + 1 < b$ we define an $\{a, b\}$ -*mountain* as a function $P : \{a, a + 1, \dots, b\} \rightarrow \mathbb{Z}_{\geq 0}$ with the properties that (i) $P(a) = P(b) = 0$; (ii) $P(i) > 0$ for all $i \in \{a + 1, \dots, b - 1\}$; and (iii) $|P(i + 1) - P(i)| = 1$ for all $i \in \{a, \dots, b - 1\}$. A function $P : \{a, a + 1, \dots, b\} \rightarrow \mathbb{Z}_{\leq 0}$ is called an $\{a, b\}$ -*valley* if the function $-P$ is an $\{a, b\}$ -mountain. We subdivide a mountain into a left side $\{a, \dots, t\}$ and right side $\{t, \dots, b\}$ where t is the smallest integer maximizing the function P . For a valley function P , the left and right side are determined by the smallest integer t minimizing the function P .

Definition 25. A *traversal* of the mountain P on $\{a, \dots, b\}$ is a sequence $(0, t) = (i_1, j_1), \dots, (i_k, j_k) = (t, b)$ with the properties

- (a) $|i_r - i_{r+1}| = |j_r - j_{r+1}| = 1$,
- (b) $P(i_r) + P(j_r) = P(t)$,
- (c) $a \leq i_r \leq t$ and $t \leq j_r \leq b$,

for all $1 \leq r \leq k-1$. We always assume that a traversal is minimal, in the sense that there is no subsequence of $(0, t) = (i_1, j_1), \dots, (i_k, j_k) = (t, b)$ which is also a traversal.

Roughly speaking, we place one person at the far left end of the mountain, and one at the first top. These persons now simultaneously traverse the mountain in such a way that the sum of their heights is always equal, and they always stay on their respective sides of the mountain that they started. The goal of the person on the left is to ascend to the top, whereas the goal of the player at the top is to descend to the far right of the mountain.

Lemma 26 ([4]). *For any mountain or valley function P on $\{a, \dots, b\}$ with first top t , there exists a traversal of P of length at most $O((t-a)(b-t))$, that can be found in time $O((t-a)(b-t))$.*




Figure 15: Example of a mountain function P on the integers in $\{0, \dots, 14\}$ with the first top at $t = 6$. The left side of the mountain is given by $\{0, \dots, 6\}$ and the right side by $\{6, \dots, 14\}$. A traversal of P is given by the sequence $(0, 6), (1, 7), (0, 8), (1, 9), (2, 10), (3, 11), (4, 12), (5, 13), (6, 14)$.

We finish this part with some additional notation that will be used later on. Let $P_j : \{a_j, \dots, b_j\} \rightarrow \mathbb{Z}$ for $j = 1, \dots, p$ be a collection of mountain and valley functions such that $a_1 = 0, a_j = b_{j+1}$ for $j = 1, \dots, p-1$, and every P_j is either a mountain or a valley. We define the *landscape* Q of the functions P_1, \dots, P_l as the function $Q : \{0, 1, \dots, b_l\} \rightarrow \mathbb{Z}$ given by $Q(i) = P_j(i)$ where $j = j(i)$ is such that $i \in \{a_j, \dots, b_j\}$. Note that $Q(0) = Q(b_l) = 0$, and $|Q(i+1) - Q(i)| = 1$ for all $i \in \{0, \dots, b_l - 1\}$. Moreover, for any function $R : \{0, \dots, r\} \rightarrow \mathbb{Z}$ satisfying the latter two conditions, there is a unique collection of mountain and valley functions so that R is the landscape of those functions. We call functions satisfying these conditions *landscape functions*.

General Case. We first partition every circuit into a collection of so-called sections, which in turn will be grouped into so-called segments. Let C_1, \dots, C_s be the canonical circuit decomposition of the symmetric difference $G \Delta G'$ for some pairing ψ , and assume w.l.o.g. that $C_i <_C C_j$ whenever $i < j$. We write $C_i = x_0^i x_1^i \dots x_{q_i}^i x_0^i$ where $x_0^i x_1^i$ is the lexicographically smallest blue edge adjacent to the starting point x_0^i of the circuit C_i that contains $q_i + 1$ edges (and where $x_0 = x_{q_i+1}$). For any i , we define the function

$$l_i(r) = \begin{cases} -1 & \text{if } \{x_{r-2}^i, x_{r-1}^i\} \text{ is cut edge and } \{x_{r-1}^i, x_r^i\} \text{ is internal edge,} \\ 1 & \text{if } \{x_{r-2}^i, x_{r-1}^i\} \text{ is internal edge and } \{x_{r-1}^i, x_r^i\} \text{ is cut edge,} \\ 0 & \text{otherwise,} \end{cases}$$

for $r = 2, 4, \dots, q_i + 1$. The function l_i indicates what happens to number of cut edges of a graphical realization when we perform a hinge flip on a pair of consecutive edges $\{x_{r-2}^i, x_{r-1}^i\}$ and $\{x_{r-1}^i, x_r^i\}$ on the circuit C_i .

Decomposition into segments. We subdivide every circuit C_i into a sequence of (not necessarily closed) walks of even length, called *sections*. Let $Z_i = \{r : l_i(r) \neq 0\} = \{z_1, \dots, z_{u_i}\} \subseteq \{2, 4, \dots, q_i + 1\}$ be the set of indices that represent a change in cut edges along the circuit, where we assume that $z_1 \leq z_2 \leq \dots \leq z_{u_i}$. We define $C_i^1 = x_0^i x_1^i \dots x_{z_1}^i$ and $C_i^j = x_{z_{j-1}}^i \dots x_{z_j}^i$ for $j = 2, \dots, u_i - 1$. If $l_i(q_i + 1) \neq 0$ this procedure partitions the circuit C_i completely, with $C_i^{u_i}$ being the last section. Otherwise, we define $C_i^{u_i+1} = x_{z_{u_i}}^i \dots x_0^i$ as the

final section, which is the remainder of the circuit C_i . We define U_i as the total number of obtained sections, which is either u_i or $u_i + 1$. Note that when $Z_i = \emptyset$, the whole circuit will form one section $C_i = C_i^1$. Also note that a section always starts with a blue edge. We extend the function l_i to sections in the following way:

$$l_i(C_i^j) = \sum_{r=2,4,\dots,z_j} l_i(r) = \begin{cases} -1 & \text{if } l_i(z_j) = -1, \\ 1 & \text{if } l_i(z_j) = 1, \\ 0 & \text{otherwise,} \end{cases}$$

for $j = 1, \dots, U_i$. Note that $l(C_i^j) \in \{-1, 1\}$ for $j = 1, \dots, u_i$, and zero for $j = u_i + 1$ if this term is present. An example is given in Figure 16.




Figure 16: The circuit $C_1 = x_0x_1 \dots x_{15}x_0$ with $q_1 = 15$. The blue edges are represented by the solid edges, and the red edges by the dashed edges. A label c on an edge indicates that it is a cut edge (all others are internal edges). We have $C_1^1 = x_0x_1x_2$ with $l_1(C_1^1) = -1$; $C_1^2 = x_2x_3x_4x_5x_6$ with $l_1(C_1^2) = 1$; $C_1^3 = x_6x_7x_8x_9x_{10}$ with $l_1(C_1^3) = -1$; $C_1^4 = x_{10}x_{11}x_{12}x_{13}x_{14}$ with $l_1(C_1^4) = -1$; and $C_1^5 = x_{14}x_{15}x_0$ with $l_1(C_1^5) = 0$ (note that $U_1 = 5$ in this example).

We now continue by grouping the union of all sections into segments in a similar flavor. For sake of readability, we rename the sections $C_1^1, \dots, C_1^{U_1}, C_2^1, \dots, C_2^{U_2}, \dots, C_s^1, \dots, C_s^{U_s}$ as D_1, \dots, D_U in the obvious way, where $U = \sum_{i=1}^s U_i$, and we define $l(D_k) = l_i(C_i^j)$ if C_i^j was renamed D_k . We define $W = \{k : l(D_k) \neq 0\} = \{w_1, \dots, w_B\}$ as the set of sections representing a change in cut edges along a circuit, where we assume that $w_1 \leq \dots \leq w_B$. We define the *segment* $S_1 = (D_1, \dots, D_{w_1})$, and $S_i = (D_{w_{i-1}+1}, \dots, D_{w_i})$ for $i = 2, \dots, w_B - 1$. If $l(D_U) \neq 0$, i.e., when $w_B = U$, this procedure completely groups the collection of sections into segments. Otherwise, we redefine the last segment as $S_B = (D_{w_{B-1}+1}, \dots, D_U)$. We can extend the function l to segments in the following way:

$$l(S_i) = \sum_{j=w_{i-1}+1}^{w_i} l(D_j) = \begin{cases} -1 & \text{if } l(D_{w_i}) = -1, \\ 1 & \text{if } l(D_{w_i}) = 1, \end{cases}$$

for $i = 1, \dots, B - 1$, and $l(S_B) = \sum_{j=w_{B-1}+1}^U l(D_j)$. Note that

$$l(S_i) \in \{-1, 1\} \quad \text{for } i = 1, \dots, B. \tag{25}$$

An example of a decomposition into segments is given in Figures 19 and 20 later on. Roughly speaking, a segment is a maximal collection of edges that could be processed, using hinge flips operations as in the warm-up example, until the number of cut-edges changes. In particular, the first segment represents precisely the point up to where we could carry out the same processing steps as in the warm-up example until the number of cut edges will have changed for the first time. Note that a segment might contain sections

²³Unless in the special case that there is only one segment S_1 covering all circuits, then $l(S_1) = 0$. This happens, e.g., in the situation of the warm-up example.

from multiple circuits, in particular, it might consist of a final section of a circuit J_1 , then some full circuits J_2, \dots, J_h (which all form a section on their own) and then the first section of some circuit J_{h+1} . The function l is then zero on the last section of J_1 and all circuits (sections) J_2, \dots, J_h , and non-zero on the section of J_{h+1} .

Unwinding/rewinding of a segment. The *unwinding* of a section $D = x_f \dots x_g$ consists of performing a number of hinge flip operations, that represent transitions in the Markov chain $\mathcal{M}'(\gamma, d)$. That is we perform a sequence of hinge flip operations replacing the (blue) edges $\{x_{r-2}, x_{r-1}\}$ by (red) edges $\{x_{r-1}, x_r\}$ for $r = f + 2, \dots, g$, in increasing order of r . Sometimes, we need to temporarily undo the unwinding of a section, in which case we perform a sequence of hinge flip operations replacing the (red) edges $\{x_{r-1}, x_r\}$ by (blue) edges $\{x_{r-2}, x_{r-1}\}$ for $r = f + 2, \dots, g$, in decreasing order of r this time. That is, we reverse the operations done during the unwinding. This is called *rewinding* a section. We say that a circuit is (currently) processed if all its sections have been unwound, and it is (currently) unprocessed if at least one section has not been unwound.

The unwinding of a segment $S_i = (D_{a_i}, \dots, D_{a_i+1})$ consists of unwinding the sections $D_{a_i}, \dots, D_{a_i+1}$ in increasing order. The rewinding of S_i consists of rewinding the section $D_{a_i}, \dots, D_{a_i+1}$ in decreasing order.




Figure 17: A section $D = x_0x_1 \dots x_6$. The blue edges are represented by the solid edges. The unwinding consists of performing first a hinge flip with $\{x_0, x_1\}$ to $\{x_1, x_2\}$; then $\{x_2, x_3\}$ to $\{x_3, x_4\}$; and finally $\{x_4, x_5\}$ to $\{x_5, x_6\}$. The rewinding consist of first a hinge flip with $\{x_5, x_6\}$ to $\{x_4, x_5\}$; then $\{x_3, x_4\}$ to $\{x_2, x_3\}$; and finally $\{x_1, x_2\}$ to $\{x_0, x_1\}$

Landscape processing. Remember that B is the number of segments obtained from the decomposition of circuits into segments. We define the function $P : \{0, 1, \dots, B\} \rightarrow \mathbb{Z}$ by $P(0) = 0$ and $P(i) = \sum_{j=1}^i l(S_j)$ for $i = 1, \dots, B$.

Lemma 27. *The function P is a landscape function.*

Proof. We have to check that $P(0) = P(B) = 0$ and that $|P(i + 1) - P(i)| = 1$ for all $i = 0, \dots, B - 1$, see the description of the mountain climbing problem. We have $P(0)$ by definition. Moreover, since both graphical realizations G and G' contain γ cut edges, it holds that $P(B) = \sum_{i=1}^B l(S_i) = 0$. Finally, using (25) and the definition of P , it follows that

$$|P(i + 1) - P(i)| = \left| \sum_{j=1}^{i+1} l(S_j) - \sum_{j=1}^i l(S_j) \right| = |l(S_{i+1})| = 1$$

for all $i = 1, \dots, B - 1$. □

Based on the segments S_1, \dots, S_B , we define the canonical path from G to G' in the state space graph of the chain $\mathcal{G}'(\gamma, d)$ that replaces all the blue edges in $G \Delta G'$ with the red edges in $G \Delta G'$. By Lemma 27 we know P is a landscape function and therefore there is a unique decomposition into mountain and valley functions P_1, \dots, P_p so that P is the landscape function for this collection, where every function is of the

form $P_j : \{a_j, \dots, b_j\} \rightarrow \mathbb{Z}$ with $a_1 = 0, b_j = a_{j+1}$ for $j = 1, \dots, p-1$, and $b_p = B$.²⁴ The *processing of a mountain/valley* P_j means that all segments $S_{a_{j+1}}, \dots, S_{b_j}$ will be unwound (it might be that during this procedure segments are temporarily rewound). This processing will rely on a traversal of the mountain, see Definition 25. We say that the segments $S_{a_{j+1}}, \dots, S_{c_j}$ are on the left side of the mountain, and the segments $S_{c_{j+1}}, \dots, S_{b_j}$ on the right side. Let $P = P_j$ for some j and assume that P is a mountain function. For sake of notation, we write $a = a_j$ and $b = b_j$, and $t = t_j$ where t_j is the first top of the mountain.

Now, fix some traversal $(a, t) = (r_1, s_1), \dots, (r_k, s_k) = (t, b)$ of P . For $c = 1, \dots, k-1$ in increasing order, do the following:

- if $r_{c+1} > r_c$ and $s_{c+1} > s_c$: first unwind segment $S_{r_{c+1}}$, then unwind segment $S_{s_{c+1}}$;
- if $r_{c+1} > r_c$ and $s_{c+1} < s_c$: first unwind segment $S_{r_{c+1}}$, then rewind segment S_{s_c} ;
- if $r_{c+1} < r_c$ and $s_{c+1} > s_c$: first rewind segment S_{r_c} , then unwind segment $S_{s_{c+1}}$;
- if $r_{c+1} < r_c$ and $s_{c+1} < s_c$: first rewind segment S_{r_c} , then rewind segment S_{s_c} .

This describes the processing of a mountain based on a traversal. Note that after the processing of a mountain, indeed all its segments have been unwound. If P is a valley function, we can use essentially the same procedure performed on $-P$. The *processing of a landscape* is done by processing the mountains/valleys P_1, \dots, P_p in increasing order.

This procedure generates a sequence $G = Z_1, Z_2, \dots, Z_l = G'$ of graphical realizations transforming G into G' where any two consecutive realizations differ by a hinge flip operation. The following lemma shows that this sequence indeed defines a (canonical) path from G to G' in the state space graph of $\mathcal{M}(\gamma, d)$, for a given pairing ψ . This lemma is essentially the motivation for the definition of $\mathcal{G}'(\gamma, d)$.

Lemma 28. *Let $Z = Z_i$ be a graphical realization on the constructed path from G to G' for pairing ψ , with degree sequence d' and γ' cut edges. Then properties (i), (ii) and (iii) defining $\mathcal{G}'(\gamma, d)$ are satisfied.*

Moreover, there exists a polynomial $r(\cdot)$ such that the length of any constructed (canonical) path carrying flow is at most $r(n)$.

Proof (sketch). Since hinge flip operations never add or remove edges, property (i) is clearly satisfied. Since the operations (1) – (4) given above unwind and rewind at most two segments, and by construction of the trajectories describing the traversal, the property (ii) is also satisfied. Finally, the cases (1) – (4), in combination with the second property of a traversal as in Definition 25, guarantee that property (iii) is satisfied. To see that all canonical paths have polynomial length, note that the traversal has polynomial length, and also every individual segment has polynomial length. \square

D.1.2 Encoding

We continue with defining the notion of an *encoding* that will be used in the next section to bound the congestion of an edge in the state space graph of $\mathcal{M}(\gamma, d)$. Let $\tau = (Z, Z')$ be a given transition of the Markov chain. Suppose that a canonical path from G to G' for some pairing $\psi \in \Psi(G, G')$, with canonical circuit decomposition $\{C_1, \dots, C_s\}$, uses the transition τ . We define $L_\tau(G, G') = (G \Delta G') \Delta Z$. An example is given in Figures 19, 20, 21 and 22.

Lemma 29. *Given $\tau = (Z, Z')$, ψ , L , if there is some pair (G, G') so that $L = L_\tau(G, G')$, then there are at most $\frac{1}{8}n^4$ such pairs.*

²⁴The function P_1 can be found by determining the first $j > 0$ so that $P(j) = 0$. The sign of $P(1)$ determines if it is a mountain or a valley. The remaining mountains and valleys can be found similarly.

Proof. For any pair (G, G') , let P be the landscape function of this canonical path between G and G' using the transition τ , and P_1, \dots, P_p its decomposition into mountain and valley functions. Let $T_{\tau, \psi}(G, G') \in \{C_1, \dots, C_s\}$ be the circuit containing the first node of the first segment of the right part of the mountain/valley P_j containing the transition τ . Moreover, if τ is used in the processing of a segment on the left side of the mountain P_j containing τ , let $\sigma_\psi(G, G')$ be the circuit containing the last node of the segment with highest index on the right side of the mountain that is currently unwound. If τ lies on the right side of the mountain, we let $\sigma_\psi(G, G')$ be the circuit containing the last node of the segment with highest index on the left side of the mountain that is currently unwound.




Figure 18: The dashed vertical lines sketch the ranges of the circuits T_ψ , σ_ψ and Γ . For every other circuit, contained in one of the four regions represented below the landscape, we know whether it has currently been processed or not.

We claim that, given $T_\psi, \sigma_\psi \in \{C_1, \dots, C_s\}$, it can be argued that there are at most 8 pairs (G, G') so that $T_\psi = T_\psi(G, G')$, $\sigma_\psi = \sigma_\psi(G, G')$. This can be seen as follows. Note that we can infer for all other circuits in $\{C_1, \dots, C_s\} \setminus \{T_\psi, \sigma_\psi, \Gamma\}$ which edges belong to G and which to G' using the (global) circuit ordering. To see this, assume that $\Gamma \leq_C T_\psi \leq_C \sigma_\psi$ (the only other case $\sigma_\psi \leq_C T_\psi \leq_C \Gamma$ is similar). Because the landscapes of the canonical paths always respect the circuit ordering, we know that all circuits in the canonical decomposition of ψ appearing before Γ have been unwound at this point. All circuits lying strictly between Γ and T_ψ are not unwound. The circuits strictly between T_ψ and σ_ψ again have been unwound, and finally, all circuits appearing after σ_ψ have not been unwound. By comparison with Z , it is uniquely determined with edges on these circuits belong to G and which to G' . For the remaining three circuits T_ψ , σ_ψ and γ there are for every circuit two possible configurations of the edges of G and G' , since every circuit alternates between edges of G and G' .²⁵ Hence, there are at most $2^3 = 8$ possible pairs (G, G') with the desired properties given T_ψ and σ_ψ .

Finally, note that for any pairing ψ , there are at most $\frac{1}{4} \binom{n}{2}$ circuits in the canonical circuit decomposition $\{C_1, \dots, C_s\}$ of the pairing ψ , as every circuit has length at least four. Hence, for both T_ψ and σ_ψ there are at most $\frac{1}{4} \binom{n}{2}$ possible choices. Since Γ is uniquely determined by the transition τ , this implies that there are at most

$$8 \cdot \frac{1}{4} \binom{n}{2} \cdot \frac{1}{4} \binom{n}{2} \leq \frac{n^4}{8}$$

possible pairs (G, G') with $L = L_\tau(G, G')$.²⁶ □

²⁵Note that we cannot use the transition τ to infer which edges belong to G and G' on the circuit Γ , as we do not know (i.e., we do not encode) whether we are unwinding or rewinding the segment containing τ .

²⁶A canonical path uses every transition at most once, which follows from the fact that we assumed that a traversal is always minimal, see Definition 25.




Figure 19: Symmetric difference $H = G \Delta G'$ where the solid edges represent the (blue) edges G and the dashed edges the (red) edges of G' . From left to right the circuit are numbered $C_1 = a_1 a_2 a_3 a_4 a_1$, $C_2 = x_0 \cdots x_{15} x_0$ and $C_3 = b_1 b_2 b_3 b_4 b_1$, and assume that this is also the order in which they are processed. Cut edges are indicated with the label c .




Figure 20: The landscape, consisting of two valleys, corresponding to the symmetric difference in Figure 19. The segments are given by $S_1 = (a_1 a_2 a_3 a_4 a_1, x_0 x_1 x_2)$, $S_2 = (x_2 x_3 x_4 x_5 x_6)$, $S_3 = (x_6 x_7 x_8 x_9 x_{10})$, $S_4 = (x_{10} x_{11} x_{12} x_{13} x_{14})$, $S_5 = (x_{14} x_{15} x_{16}, b_1 b_2)$, and $S_6 = (b_3 b_4 b_1)$.




Figure 21: The transition $\tau = (Z, Z')$ that is the hinge flip operation that removes the edge $\{x_{10}, x_{11}\}$ and adds the edge $\{x_{11}, x_{12}\}$ as part of the unwinding of S_4 . Note that the segments S_1 and as S_2 , forming the first valley, have been processed already. Also, the first segment S_3 of the left part of the second valley, as well as the segment S_5 being the first segment of the right part of the second valley, have been processed already. The edges in $(E(G) \cup E(G')) \setminus E(H)$ are left out.




Figure 22: The encoding $L = L_t(G, G') = (G \Delta G') \Delta Z$ for the symmetric difference in Figure 19 and transition as in Figure 21, where again the edges in $(E(G) \cup E(G')) \setminus E(H)$ are left out.

D.1.3 Bounding the Congestion

For a tuple (G, G', ψ) , let $p_\psi(G, G')$ denote the canonical path from G to G' for pairing ψ . Let

$$\mathcal{L}_\tau = \cup_{(G, G', \psi) \in \mathcal{F}_\tau} L_\tau(G, G')$$

be the union of all (distinct) encodings L_τ , where $\mathcal{F}_\tau = \{(G, G', \psi) : \tau \in p_\psi(G, G')\}$ is the set of all tuples (G, G', ψ) such that the canonical path from G to G' under pairing ψ uses the transition τ . A crucial observation is the following.

Lemma 30. *If $L_\tau(G, G') = (G \Delta G') \Delta Z$ for transition $\tau = (Z, Z')$ used by a canonical path between G and G' , then $L \in \mathcal{G}'(\gamma, d)$. This implies that*

$$|\mathcal{L}_\tau| \leq |\mathcal{G}'(\gamma, d)|. \quad (26)$$

Proof (sketch). We check that the properties (i), (ii) and (iii) defining the set $\mathcal{G}'(\gamma, d)$ are satisfied by L . Note that $L \Delta Z = G \Delta G$. As every individual hinge flip operations adds and removes an arc from the symmetric difference, it follows that L and Z have the same number of edges. This proves property (i). Also, if Z has a perturbation of $\alpha_v \in \{-2, -1, 0, -1, -2\}$ (see Proposition 20) at node v , then L has a perturbation of $-\alpha_v$ at node v , which shows that property (ii) is satisfied for L . Finally, with $\beta \in \{-1, 0, 1\}$, if Z contains $\gamma - \beta$ cut edges, then L contains $\gamma + \beta$ cut edges (using implicitly that G and G' contain the same number of cut edges). This implies that property (iii) is satisfied. \square

Moreover, with $H = G \Delta G'$ and $L = L_t(G, G')$, the pairing ψ has the property that it pairs up the edges of $E(H) \setminus E(L)$ and $E(H) \cap E(L)$ in such a way that for every node v each edge in $E(H) \setminus E(L)$ that is incident to v is paired up with an edge in $E(H) \cap E(L)$ that is incident to v , except for at most four pairs.²⁷ Let $\Psi'(L)$ be the set of all pairings with this property. Remember that we do not need to know G and G' in order to determine the set $H = L_\tau \Delta Z$. Note that not every such pairing has to correspond to a tuple (G, G', ψ) for which $t \in p_\psi(G, G')$. Using a counting argument,²⁸ we can upper bound $|\Psi'(L)|$ in terms of $|\Psi(H)|$. In particular, there exists a polynomial $q(n)$ such that

$$|\Psi'(L)| \leq q(n) \cdot |\Psi(H)|. \quad (27)$$

²⁷These are the nodes x_0, x_{12}, b_1 and b_3 in Figure 22.

²⁸This can be done similarly as the argument used towards the end of Appendix A.2.

²⁹A very rough choice is $q(n) = n^{20}$.

Putting everything together, we have

$$\begin{aligned}
|\mathcal{G}'(\gamma, d)|^2 f'(\tau) &= \sum_{(G, G')} \sum_{\psi \in \Psi(G, G')} \mathbf{1}(e \in p_\psi(H)) |\Psi(H)|^{-1} \\
&\leq \frac{1}{8} n^4 \sum_{L \in \mathcal{L}_\tau} \sum_{\psi' \in \Psi'(L)} |\Psi(H)|^{-1} \quad (\text{using Lemma 15}) \\
&\leq \frac{1}{8} n^4 \cdot q(n) \sum_{L \in \mathcal{L}_\tau} 1 \quad (\text{using (27)}) \\
&\leq \frac{1}{8} n^4 \cdot q(n) \cdot |\mathcal{G}'(\gamma, d)| \quad (\text{using (26)})
\end{aligned} \tag{28}$$

The usage of Lemma 29 for the first inequality works as follows. Every tuple $(G, G', \psi) \in \mathcal{F}_t$ with encoding $L_t(G, G')$ generates a unique tuple in $\{L_t(G, G')\} \times \Psi'(L_t(G, G'))$. But since, by Lemma 29, there are at most $\frac{1}{8}n^4$ pairs (G, G') with $L = L_\tau(G, G')$ for given L, τ and ψ , we have that $\frac{1}{8}n^4 \sum_{L \in \mathcal{L}_\tau} |\{L\} \times \Psi'(L)| = \frac{1}{8}n^4 \sum_{L \in \mathcal{L}_\tau} \sum_{\psi' \in \Psi'(L)} 1$ is an upper bound on the number of canonical paths that use τ .

By rearranging (28) we get the upper bound for f' required in Lemma 24. We already observed that the length of any canonical path is polynomially bounded as well. This then completes the proof of Theorem 23.

E Strongly Stable Families for the PAM Model

In Appendix D we have shown that the hinge flip Markov chain for PAM instances with two classes is rapidly mixing on $\mathcal{G}'(\gamma, d)$ in case (γ, d) comes from a family of strongly stable tuples. In this section we give two explicit families of sequences that are strongly stable. When dealing with a family of instances, even when this is not explicitly mentioned, we only consider the tuples (c, d) for which there is at least one graphical realization.

Theorem 31 (Regular classes). *Let \mathcal{D} be the family of instances of the joint degree matrix model, i.e., where for every tuple (V_1, V_2, γ, d) it holds that $1 \leq \beta_1, \beta_2 \leq |V| - 1$, and $1 \leq \gamma \leq |V_1||V_2| - 1$, where β_1 and β_2 are the common degrees in the classes V_1 and V_2 , respectively. The family \mathcal{D} is strongly stable for $k = 6$, and, hence, the hinge flip chain is rapidly mixing for all tuples in \mathcal{D} .*

Proof. We first show that this family is strongly stable for $k = 6$. For convenience, we will work with the notation $\mathcal{G}'(c, d)$ instead of $\mathcal{G}'(\gamma, d)$. Remember that

$$c_{ii} = \left(\sum_{j \in V_i} d_j \right) - \gamma$$

for $i = 1, 2$ is the number of internal edges that V_i has in any graphical realization in $\mathcal{G}(\gamma, d)$, and that $\gamma = c_{12} = c_{21}$. For sake of readability, we define the notion of a cancellation hinge flip. For either $i = 1$ or $i = 2$, suppose nodes $v, w \in V_i$, are such that v has a degree deficit of at least one, and w a degree surplus of at least one. Then w has a neighbor $z \in V$ that is not a neighbor of v (using that v and w have the same degree β_i). The hinge flip operation that removes the edge $\{z, w\}$ and adds the edge $\{z, v\}$ is called a *cancellation flip on v and w* . Note that the number of internal edges in V_1 and V_2 as well as the number of cut edges does not change with such an operation.³⁰ Moreover, we say that an edge $\{a, b\}$ is a non-edge of a graphical realization G if $\{a, b\} \notin E(G)$.

³⁰That is, either z lies in the other class, in which case the cancellation flip removes and adds a cut edge, or, z lies in the same class as v and w in which case an internal edge in V_i is removed and added.

Let $G \in \mathcal{G}'(c, d)$ for some tuple (c', d') as in the definition of $\mathcal{G}'(c, d)$ at the start of Section D. We first show that with at most four hinge flip operations, we can obtain a perturbed auxiliary state $G^* \in \mathcal{G}'(c, d)$ for which its tuple (c^*, d^*) is edge-balanced. That is, it satisfies $c^* = c$. Remember that the value c'_{12} uniquely determines the matrix c' , and, by assumption of $\mathcal{G}'(c, d)$, we have $c'_{12} \in \{c_{12} - 1, c_{12}, c_{12} + 1\}$. We can therefore distinguish the following cases.

- **Case 1:** $c'_{12} = c_{12} + 1$. Then, by Proposition 20, either $c'_{11} = c_{11} - 1$ and $c'_{22} = c_{22}$, or, $c'_{22} = c_{22} - 1$ and $c'_{11} = c_{11}$. Assume without loss of generality that we are in the first case. Then it holds that

$$\sum_{j \in V_1} d'_j = \left(\sum_{j \in V_1} d_j \right) - 1 \quad \text{and} \quad \sum_{j \in V_2} d'_j = \left(\sum_{j \in V_2} d_j \right) + 1. \quad (29)$$

Moreover, there must be at least one node $v_2 \in V_2$ with a degree surplus (of either one or two), and there is at least one non-edge $\{a, b\}$ with both endpoints in V_1 . If v_2 is adjacent to either a or b , we can perform a hinge flip to make the realization G edge-balanced, so assume this is not the case. Also, if the total deficit of a and b is -2 , there must be a node in V_1 with degree surplus, otherwise (29) is violated. Then we can perform a cancellation flip in V_1 to remove the deficit at either a or b . Hence, we may assume without loss of generality that a does not have a degree deficit.

- **Case A:** v_2 has a neighbor $v_1 \in V_1$. If v_1 has a degree surplus we can perform a cancellation flip in V_1 to remove it, which must exist by (29). So assume v_1 has no degree surplus. As node a has no deficit, and is not connected to v_2 , whereas v_1 is, there must be some neighbor p of a which is not a neighbor of v_1 . This holds since v_1 and b have the same degree β_1 in the sequence d . Then the path $v_2 - v_1 - p - a - b$ alternates between edges and non-edges of G , and with two hinge flips we can obtain an edge-balanced realization in $\mathcal{G}'(\gamma, d)$.




Figure 23: Sketch of first case with subcase A.

- **Case B:** v_2 has no neighbors in V_1 . We know that there is at least one cut edge $\{q, r\}$ in the realization G , since $c'_{12} = c_{12} + 1$. If q has a degree surplus, we are in the situation of Case A. Otherwise v_2 has a neighbor u which is not a neighbor of q , since q and v_2 have the same degree β_2 in the sequence d . We can then perform the hinge flip that removes $\{v_2, u\}$ and adds $\{u, q\}$. If q now has a degree surplus, we are in Case A. Otherwise, in case this hinge flip cancelled out a degree deficit at q , there must be at least one other node in V_2 with a degree surplus, because of (29). We can then perform the same step again, which will now result in a degree surplus at q . This is true since the node q cannot have a deficit of -2 , since (29) would then imply that the total degree surplus of nodes in V_2 is at least three, which violates the second property defining $\mathcal{G}'(c, d)$. That is, we can always reduce to the situation in Case A.

Summarizing, we can always find an edge-balanced realization G^* using at most four hinge flip operations in case $c'_{12} = c_{12} + 1$.

- **Case 2:** $c'_{12} = c_{12} - 1$. Using complementation, it can be seen that this case is similar to Case 1. That is, we consider the tuple (\bar{c}, \bar{d}) in which all nodes in V_1 have degree $|V| - \beta_1$, all nodes in V_2 have degree $|V| - \beta_2$, and where all feasible graphical realizations have $\bar{c}_{12} = |V_1||V_2| - c_{12}$ cut edges. The case $c_{12} = c_{12} - 1$ then corresponds to the case $\bar{c}'_{12} = \bar{c}_{12} + 1$.
- **Case 3:** $c'_{12} = c_{12}$. If also $c_{11} = c'_{11}$ we are done. Otherwise, suppose that $c_{11} = c'_{11} + 1$. Then it must be that $c_{22} = c'_{22} - 1$, as $c'_{12} = c_{12}$, and it holds that

$$\sum_{j \in V_1} d'_j = \left(\sum_{j \in V_1} d_j \right) + 2 \quad \text{and} \quad \sum_{j \in V_2} d'_j = \left(\sum_{j \in V_2} d_j \right) - 2. \quad (30)$$

Then there is at least one edge $\{a, b\}$ in the graphical realization with $a, b \in V_1$. Moreover, we may assume that a has a degree surplus. If not, then there is at least one other node u with a degree surplus because of (30). Performing a cancellation flip then gives the node a a degree surplus (it could not be that a had a degree deficit, as this would imply, in combination with (30), that the total degree surplus of nodes in V_1 is at least three).




Figure 24: Sketch of last situation in Case 3.

Now, if there is a non-edge of the form $\{b, v_2\}$ for some $v_2 \in V_2$, we can perform a hinge flip operation removing $\{a, b\}$ and adding $\{b, v_2\}$ in order to end up in Case 1. Otherwise, assume that b is adjacent to all $v_2 \in V_2$. As b is also adjacent to a , and a has a degree surplus of at least one,³¹ it follows that $\beta_1 \geq |V_2|$. Now, by the assumption that $c_{12} \leq |V_1||V_2| - 1$, there is at least one non-edge $\{p, q\}$ with $p \in V_1$ and $q \in V_2$. As p is not adjacent to q , but has degree at least $\beta_1 \geq |V_2|$, it must be that p is adjacent to some $r \in V_1$. If r has a degree surplus, then we can perform a hinge flip that removes $\{p, r\}$ and adds $\{p, q\}$ in order to end up in the situation of Case 1. Otherwise, node a , which has a degree surplus, has some neighbor w which is not a neighbor of r . This implies the path $a - w - r - p - q$ alternates between edges and non-edges of G . Performing two hinge flips then brings us in the situation of Case 1.

We have shown that with at most four hinge flips we can always obtain some $G^* \in \mathcal{G}'(c, d)$ that is edge-balanced. This implies that

$$\sum_{j \in V_1} d_j^* = \sum_{j \in V_1} d_j \quad \text{and} \quad \sum_{j \in V_2} d_j^* = \sum_{j \in V_2} d_j. \quad (31)$$

Now, if $v \in V_1$ has a degree surplus, there must be some $w \in V_1$ that has a degree deficit, because of (31). We can then perform a cancellation flip to decrease the sum of the total degree deficit and degree surplus. A similar statement is true if $v \in V_2$ has a degree surplus. By performing this step at most twice, we obtain a realization $H \in \mathcal{G}(c, d)$. That is, with at most six hinge flip operations in total we can transform G into a graphical realization in $\mathcal{G}(c, d)$. This shows that \mathcal{D} is strongly stable for $k = 6$. \square

³¹That is, b can have a degree surplus of at most one. A degree surplus of two at b would only give a bound of $\beta_1 \geq |V_2| - 1$.

We next show that certain families of sparse instances are strongly stable as well. Sparsity here refers to the fact that the maximum degree in a class is significantly smaller than the size of the class, as well as the fact that the number of cut edges is (much) smaller than the total number of edges in a graphical realization.

Theorem 32 (Sparse irregular families). *Let $0 < \alpha < 1/2$ be fixed, and let \mathcal{D}_α be the family of tuples (V_1, V_2, c, d) for which*

- (i) $|V_1|, |V_2| \geq \alpha n$,
- (ii) $2 \leq d_i \leq \sqrt{\frac{\alpha n}{4}}$ for $i \in V_1 \cup V_2$, and,
- (iii) $1 \leq c_{12} \leq \frac{\alpha n}{2}$.

The class \mathcal{D}_α is strongly stable for $k = 9$, and, hence, the hinge flip chain is rapidly mixing for all tuples in \mathcal{D}_α .

Proof. We proceed in a similar fashion as the proof of Corollary 31. We will use the notion of an alternating path. For a given graph $H = (W, E)$, an alternating path (x_1, \dots, x_q) is an odd sequence of nodes so that $\{x_i, x_{i+1}\} \in E$ for i even, and $\{x_i, x_{i+1}\} \notin E$ for i odd.

Lemma 33 (following from [28]). *Let $\delta = (\delta_1, \dots, \delta_r)$ be a degree sequence with $1 \leq \delta_i \leq \sqrt{r/2}$ for all $i = 1, \dots, r$. Fix $x, y \in [r]$ and let $H = ([r], E)$ be a graphical realization of the degree sequence δ' where $\delta'_x = \delta_x + 1$, $\delta'_y = \delta_y - 1$ and $\delta'_i = \delta_i$ for all $i \in [r] \setminus \{x, y\}$. Then there exists an alternating path of length at most four starting at x and ending at y .*

Now, let $G \in \mathcal{G}'(c, d)$ for some tuple (c', d') as in the definition of $\mathcal{G}'(c, d)$ at the start of Section D. We show that with at most four hinge flip operations, we can transform G into a perturbed auxiliary state $G^* \in \mathcal{G}'(c, d)$ for which its tuple (c^*, d^*) is edge-balanced.

- **Case 1:** $c'_{12} = c_{12} + 1$. Then, by Proposition 20, either $c'_{11} = c_{11} - 1$ and $c'_{22} = c_{22}$, or, $c'_{22} = c_{22} - 1$ and $c'_{11} = c_{11}$. Assume without loss of generality that we are in the first case. Then it holds that

$$\sum_{j \in V_1} d'_j = \left(\sum_{j \in V_1} d_j \right) - 1 \quad \text{and} \quad \sum_{j \in V_2} d'_j = \left(\sum_{j \in V_2} d_j \right) + 1. \quad (32)$$

Moreover, there must be at least one node $v_2 \in V_2$ with a degree surplus (of either one or two).

- **Case A:** v_2 has a neighbor $v_1 \in V_1$. By assumptions *i*) and *ii*) it must be that there is some $b \in V_1$ so that v_1 is not a neighbor of b . Then we can perform the hinge flip that removes $\{v_2, v_1\}$ and adds $\{v_1, b\}$, resulting in an edge-balanced realization.
- **Case B:** v_2 has no neighbors in V_1 . Since $c_{12} \geq 1$ by assumption *iii*), there is some edge $\{a, b\}$ in G with $a \in V_1$ and $b \in V_2 \setminus \{v_2\}$. We consider the induced subgraph H on the nodes in V_2 and use δ' to denote its degree sequence. We next apply Lemma 33 with $v_2 = x$ and $b = y$. To see that this is possible, note that $c_{12} \leq \alpha n/2$ by assumption *iii*). This implies, in combination with assumption *i*) that at least $\alpha n/2$ nodes in V_1 have degree at least one in H . Thus, we may apply Lemma 33 with $r = |V(H)| \geq \alpha n/2$, since $d_i \leq \sqrt{\alpha n/4}$ by assumption *ii*) (which is less or equal than $\sqrt{r/2}$ if $r \geq \alpha n/2$). Hence, there exists an alternating path of length at most four starting at v_2 and ending at b . Then by performing two hinge flips, resulting in the removal of $\{v_2, f\}, \{g, h\}$ and addition of $\{f, g\}, \{h, b\}$, we are in the situation of Case A where b now plays the role of v_2 .




Figure 25: Sketch of subcase B in Case 1.

- **Case 2:** $c'_{12} = c_{12} - 1$. Suppose without loss of generality that $c'_{11} = c_{11} + 1$ and $c'_{22} = c_{22}$, and note that there must be at least one node $v_1 \in V_1$ with a degree surplus. If v_1 has a neighbor in V_1 we can perform a hinge flip operation to obtain an edge-balanced realization as desired (as this neighbor has at least one non-neighbor in V_2 by assumptions *i*) and *ii*). Therefore, assume that all neighbors of v_1 lie in V_2 . Pick some neighbor $b \in V_2$ of v_1 . Since all nodes in V_1 have degree at least one, and $c_{12} \leq \alpha n/2$, it must be that b is not adjacent to some node $a \in V_1$ that has degree at least one in the induced subgraph H on the nodes of V_1 , as $d_b \leq \sqrt{\alpha n/4}$ and $|V(H)| \geq \alpha n/2$. Let c be some neighbor of a in V_1 that exists by assumption that a has degree at least one in H . Also, c is not adjacent to some $d \in V_2$ for similar reasons as that b was not adjacent to a . This means that with two hinge flip operations we can obtain an edge-balanced realization.




Figure 26: Sketch of Case 2.

- **Case 3:** $c'_{12} = c_{12}$. If also $c_{11} = c'_{11}$ we are done. Otherwise, suppose that $c_{11} = c'_{11} + 1$. Then it must be that $c_{22} = c'_{22} - 1$, as $c'_{12} = c_{12}$, and it holds that

$$\sum_{j \in V_1} d'_j = \left(\sum_{j \in V_1} d_j \right) + 2 \quad \text{and} \quad \sum_{j \in V_2} d'_j = \left(\sum_{j \in V_2} d_j \right) - 2. \quad (33)$$

There is at least one node $v_1 \in V_1$ with a degree surplus. If v_1 has a neighbor $a \in V_1$, which in turn has a non-neighbor $b \in V_2$ by assumptions *i*) and *ii*), then we can reduce to the situation of Case 1 by performing a hinge flip removing $\{v_1, a\}$ and adding $\{a, b\}$. Therefore, assume that all neighbors of v_1 are in V_2 . We can then pick some neighbor $b \in V_2$ and perform a similar step as in Case 2 to find an alternating path (v_1, b, a, c) with $v_1, a, c \in V_1$ and $b \in V_2$. Then we can perform two hinge flip operations to reduce to Case 1 again.

We have shown that with at most three hinge flips we can always obtain some $G^* \in \mathcal{G}'(c, d)$ that is edge-balanced. This implies that

$$\sum_{j \in V_1} d_j^* = \sum_{j \in V_1} d_j \quad \text{and} \quad \sum_{j \in V_2} d_j^* = \sum_{j \in V_2} d_j. \quad (34)$$

Now, if $v \in V_1$ has a degree surplus, there must be some $w \in V_1$ that has a degree deficit, because of (34). Moreover, if all neighbors in of v are in V_2 , we can transfer the degree deficit to some node with degree at

least one in the subgraph induced on V_1 at the cost of one hinge flip operation (similar as the analysis in Case 2). That is, we may assume that v has a neighbor in V_1 . Then, using similar arguments as in Case 1.B, it follows that there exists an alternating path from v to w in V_1 , which allows us to decrease the total degree surplus/deficit using two hinge flip operations. This can be repeated to obtain a feasible realization $H \in \mathcal{G}(\gamma, d)$.

Overall, this can be done using at most nine hinge flip operations. \square

In particular, Theorem 32 directly implies the following which, to the best of our knowledge, is the first result of its kind for the PAM model.

Corollary 34. *Let \mathcal{D}_α be as in Theorem 32. Then there is an fully polynomial almost uniform generator for tuples in the family \mathcal{D}_α .*

F Switch Markov Chain for Two Regular Degree Classes

In this section we will use an embedding argument similar to that in the proof of Theorem 8 to show the switch chain is rapidly mixing in case both classes are regular, i.e., for instances that are essentially JDM instances with two degree classes. Recall the *restricted switch Markov chain* for sampling graphical realizations of $\mathcal{G}(c, d)$ defined in Section 2: Let G be the current state of the *switch chain*.

- With probability $1/2$, do nothing.
- Otherwise, perform a *switch move*: select two edges $\{a, b\}$ and $\{x, y\}$ uniformly at random, and select a perfect matching M on nodes $\{x, y, a, b\}$ uniformly at random (there are three possible options). If $M \cap E(G) = \emptyset$ and $G + M - \{a, b\} \cup \{x, y\} \in \mathcal{G}(c, d)$, then delete $\{a, b\}, \{x, y\}$ from $E(G)$ and add the edges of M . This local operation is called a *switch*.

As usual, graphs $G, G' \in \mathcal{G}(c, d)$ are said to be *switch adjacent* if G can be obtained from G' with positive probability in one transition of this chain. The *switch-distance* $\text{dist}_{\mathcal{G}(c,d)}(G, G')$ is the length of a shortest path between G and G' in the state space graph of the switch chain. Also recall, $P(G, G')^{-1} \leq n^4$ for all adjacent $G, G' \in \mathcal{G}(c, d)$, and also the maximum in- and out-degrees of the state space graph of the switch chain are bounded by n^4 .

While this chain is known to be irreducible for the instances of the JDM model [1, 11], in general this is not true [15]. To the best of our knowledge, there is no clear understanding for which pairs c and d it is irreducible in general. Nevertheless, we present the following meta-result for the rapid mixing of the switch chain, which in particular applies in case the degrees are component-wise regular (Theorem 10).

Theorem 35. *Let \mathcal{D} be a strongly stable family of tuples (γ, d) with respect to some constant k , and suppose there exists a function $p_0 : \mathbb{N} \rightarrow \mathbb{N}$ with the property that, for any fixed $x \in \mathbb{N}$: if $(\gamma, d) \in \mathcal{D}$, and $G, G' \in \mathcal{G}(\gamma, d)$ so that $|E(G) \Delta E(G')| \leq x$, the switch-distance satisfies $\text{dist}_{\mathcal{G}(\gamma,d)}(G, G') \leq p_0(x)$. Then the switch chain is rapidly mixing for all tuples in the family \mathcal{D} with respect to the uniform stationary distribution over $\mathcal{G}(\gamma, d)$.*

Proof (sketch). First note that by definition of the function p_0 the switch chain is irreducible. Moreover, it is not hard to see that the switch chain is aperiodic and symmetric as well. This implies that it has a unique stationary distribution which is the uniform distribution over $\mathcal{G}(\gamma, d)$. Moreover, by assumption of strong stability, we know that the auxiliary chain $\mathcal{M}(\gamma, d)$ is rapidly mixing. In particular, from Theorem 23, we know there exists a flow f that routes $1/|\mathcal{G}'(\gamma, d)|^2$ units of flow between any pair of states in $\mathcal{G}(\gamma, d)$ in the

state space graph of the chain $\mathcal{M}(\gamma, d)$, with the property that $f(e) \leq p(n)/|\mathcal{G}'(\gamma, d)|$, and $\ell(f) \leq r(n)$, for some polynomials $p(\cdot), r(\cdot)$ whose degrees may only depend on $k(\gamma, d)$.

One can then use exactly the same embedding argument as in the proof of Theorem 8. The existence of the function p_0 , together with the notion of strong stability as defined in Appendix D, are sufficient for reproducing all the arguments. \square

Theorem 10. *Let \mathcal{D} be the family of instances of the joint degree matrix model with two degree classes. Then the switch chain is rapidly mixing for instances in \mathcal{D} .*

Proof. Strong stability was shown in the previous section in Theorem 31. Moreover, from the proof of Lemma 7 in [1] it follows that for any two graphs $H, H' \in \mathcal{G}(\gamma, d)$, H can be transformed into H' using at most $\frac{3}{2}|E(H)\Delta E(H')|$ switches of the restricted switch chain. That is, we may take $p_0(x) = \frac{3}{2}x$. Then the statement follows from Theorem 35. \square

References

- [1] Georgios Amanatidis, Bradley Green, and Milena Mihail. Graphic realizations of joint-degree matrices. *CoRR*, abs/1701.03856, 2008.
- [2] Georgios Amanatidis, Bradley Green, and Milena Mihail. Connected realizations of joint-degree matrices. *Discrete Applied Mathematics*, 2018.
- [3] Mohsen Bayati, Jeong Han Kim, and Amin Saberi. A sequential algorithm for generating random graphs. *Algorithmica*, 58(4):860–910, 2010. *A preliminary conference version appeared in APPROX-RANDOM 2007.*
- [4] Nayantara Bhatnagar, Dana Randall, Vijay V. Vazirani, and Eric Vigoda. Random bichromatic matchings. *Algorithmica*, 50(4):418–445, Apr 2008. *A preliminary conference version appeared in LATIN 2006.*
- [5] Joseph Blitzstein and Persi Diaconis. A sequential importance sampling algorithm for generating random graphs with prescribed degrees. *Internet mathematics*, 6(4):489–522, 2011.
- [6] Corrie Jacobien Carstens and Pieter Kleer. Speeding up switch Markov chains for sampling bipartite graphs with given degree sequence. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018, August 20-22, 2018 - Princeton, NJ, USA*, volume 116 of *LIPICs*, pages 36:1–36:18. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.
- [7] Colin Cooper, Martin E. Dyer, and Catherine S. Greenhill. Sampling regular graphs and a peer-to-peer network. *Combinatorics, Probability & Computing*, 16(4):557–593, 2007. *A preliminary conference version appeared in SODA 2005.*
- [8] Colin Cooper, Martin E. Dyer, and Catherine S. Greenhill. Corrigendum: Sampling regular graphs and a peer-to-peer network. *CoRR*, abs/1203.6111, 2012.
- [9] Colin Cooper, Martin E. Dyer, Catherine S. Greenhill, and Andrew J. Handley. The flip Markov chain for connected regular graphs. *Discrete Applied Mathematics*, 2018. *A preliminary conference version appeared in PODC 2009.*

- [10] Éva Czabarka. Partition adjacency matrices in network modeling. *Technical report on GRAPHS project FA9550-12-1-040*, 2014.
- [11] Éva Czabarka, Aaron Dutle, Péter L. Erdős, and István Miklós. On realizations of a joint degree matrix. *Discrete Applied Mathematics*, 181:283–288, 2015.
- [12] Éva Czabarka, Laszlo A Szekely, Zoltan Toroczkai, and Shanise Walker. An algebraic Monte-Carlo algorithm for the bipartite partition adjacency matrix realization problem. *CoRR*, abs/1708.08242, 2017.
- [13] Martin Dyer, Leslie Ann Goldberg, Mark Jerrum, and Russell Martin. Markov chain comparison. *Probab. Surveys*, 3:89–111, 2006.
- [14] Martin Dyer, Mark Jerrum, and Haiko Müller. On the switch Markov chain for perfect matchings. *J. ACM*, 64(2):12:1–12:33, April 2017. *A preliminary conference version appeared in SODA 2016.*
- [15] Péter L. Erdős, Stephen G. Hartke, Leo van Iersel, and István Miklós. Graph realizations constrained by skeleton graphs. *Electr. J. Comb.*, 24(2):P2.47, 2017.
- [16] Péter L Erdős, Zoltán Király, and István Miklós. On the swap-distances of different realizations of a graphical degree sequence. *Combinatorics, Probability and Computing*, 22(3):366–383, 2013.
- [17] Péter L. Erdős, Sándor Z. Kiss, István Miklós, and Lajos Soukup. Approximate counting of graphical realizations. *PLOS ONE*, 10(7):1–20, 2015.
- [18] Péter L. Erdős, Tamás Róbert Mezei, István Miklós, and Dániel Soltész. Efficiently sampling the realizations of bounded, irregular degree sequences of bipartite and directed graphs. *PLOS ONE*, 13(8):1–20, 08 2018.
- [19] Péter L Erdős, István Miklós, and Zoltán Toroczkai. A decomposition based proof for fast mixing of a Markov chain over balanced realizations of a joint degree matrix. *SIAM Journal on Discrete Mathematics*, 29(1):481–499, 2015.
- [20] Tomás Feder, Adam Guetz, Milena Mihail, and Amin Saberi. A local switch Markov chain on given degree graphs with application in connectivity of peer-to-peer networks. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 69–76. IEEE, 2006.
- [21] Pu Gao and Nicholas C. Wormald. Uniform generation of random graphs with power-law degree sequences. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1741–1758. SIAM, 2018.
- [22] Minas Gjoka, Balint Tillman, and Athina Markopoulou. Construction of simple graphs with a target joint degree matrix and beyond. In *2015 IEEE Conference on Computer Communications, INFOCOM 2015, Kowloon, Hong Kong, April 26 - May 1, 2015*, pages 1553–1561. IEEE, 2015.
- [23] Catherine S. Greenhill. A polynomial bound on the mixing time of a Markov chain for sampling regular directed graphs. *Electronic Journal of Combinatorics*, 18(1), 2011.
- [24] Catherine S. Greenhill. The switch Markov chain for sampling irregular graphs. In *Proceedings of the Twenty-sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '15*, pages 1564–1572, Philadelphia, PA, USA, 2015.
- [25] Catherine S. Greenhill and Matteo Sfragara. The switch Markov chain for sampling irregular graphs and digraphs. *Theor. Comput. Sci.*, 719:1–20, 2018.

- [26] Tatsuo Homma. A theorem on continuous functions. *Kodai Mathematical Seminar Reports*, 4(1):13–16, 1952.
- [27] Mark Jerrum, Brendan D. McKay, and Alistair Sinclair. *When is a Graphical Sequence Stable?* Dep. of Computer Science: Internal report. University of Edinburgh, Department of Computer Science, 1989.
- [28] Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18(6):1149–1178, 1989. *A preliminary conference version appeared in STOC 1988.*
- [29] Mark Jerrum and Alistair Sinclair. Fast uniform generation of regular graphs. *Theoretical Computer Science*, 73(1):91 – 100, 1990.
- [30] Ravi Kannan, Prasad Tetali, and Santosh Vempala. Simple Markov-chain algorithms for generating bipartite graphs and tournaments. *Random Structures and Algorithms*, 14(4):293–308, 1999. *A preliminary conference version appeared in SODA 1997.*
- [31] David Asher Levin, Yuval Peres, and Elizabeth Lee Wilmer. *Markov chains and mixing times*. American Mathematical Soc., 2009.
- [32] Priya Mahadevan, Dmitri V. Krioukov, Kevin R. Fall, and Amin Vahdat. Systematic topology analysis and generation using degree correlations. In *Proceedings of the ACM SIGCOMM 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Pisa, Italy, September 11-15, 2006*, pages 135–146. ACM, 2006.
- [33] Brendan D. McKay and Nicholas C. Wormald. Asymptotic enumeration by degree sequence of graphs with degree $o(n^{1/2})$. *Combinatorica*, 11(4):369–382, 1991.
- [34] István Miklós, Péter L. Erdős, and Lajos Soukup. Towards random uniform sampling of bipartite graphs with given degree sequence. *Electronic Journal of Combinatorics*, 20(1), 2013.
- [35] Benjamin P. Olding and Patrick J. Wolfe. Inference for graphs and networks: Adapting classical tools to modern data. In *Data Analysis for Network Cyber-Security*, pages 1–31. World Scientific, 2014.
- [36] A.N. Patrinos and S.L. Hakimi. Relations between graphs and integer-pair sequences. *Discrete Mathematics*, 15(4):347–358, 1976.
- [37] A. Ramachandra Rao, Rabindranath Jana, and Suraj Bandyopadhyay. A Markov chain Monte Carlo method for generating random $(0, 1)$ -matrices with given marginals. *Sankhyā: The Indian Journal of Statistics, Series A*, 58(2):225–242, 1996.
- [38] Julián Salas and Vicenç Torra. Improving the characterization of P-stability for applications in network privacy. *Discrete Applied Mathematics*, 206:109 – 114, 2016.
- [39] Alistair Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1:351–370, 1992. *A preliminary conference version appeared in LATIN 1992.*
- [40] Alistair Sinclair and Mark Jerrum. Approximate counting, uniform generation and rapidly mixing Markov chains. *Inf. Comput.*, 82(1):93–133, 1989.
- [41] Isabelle Stanton and Ali Pinar. Constructing and sampling graphs with a prescribed joint degree distribution. *ACM Journal of Experimental Algorithmics*, 17(1), 2011.

- [42] Richard Taylor. Constrained switchings in graphs. In Kevin L. McAvaney, editor, *Combinatorial Mathematics VIII*, pages 314–336, Berlin, Heidelberg, 1981. Springer Berlin Heidelberg.
- [43] James V Whittaker. A mountain-climbing problem. *Canad. J. Math*, 18(873-882):4, 1966.