

Rapid Mixing of the Down-Up Walk on Matchings of a Fixed Size

Vishesh Jain ✉ 

Department of Mathematics, Statistics, and Computer Science,
University of Illinois Chicago, Chicago, IL, USA

Clayton Mizgerd ✉

Department of Mathematics, Statistics, and Computer Science,
University of Illinois Chicago, Chicago, IL, USA

Abstract

Let $G = (V, E)$ be a graph on n vertices and let $m^*(G)$ denote the size of a maximum matching in G . We show that for any $\delta > 0$ and for any $1 \leq k \leq (1 - \delta)m^*(G)$, the down-up walk on matchings of size k in G mixes in time polynomial in n . Previously, polynomial mixing was not known even for graphs with maximum degree Δ , and our result makes progress on a conjecture of Jain, Perkins, Sah, and Sawhney [STOC, 2022] that the down-up walk mixes in optimal time $O_{\Delta, \delta}(n \log n)$.

In contrast with recent works analyzing mixing of down-up walks in various settings using the spectral independence framework, we bound the spectral gap by constructing and analyzing a suitable multi-commodity flow. In fact, we present constructions demonstrating the limitations of the spectral independence approach in our setting.

2012 ACM Subject Classification Mathematics of computing → Markov-chain Monte Carlo methods

Keywords and phrases Down-up walk, Matchings, MCMC

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2024.63

Category RANDOM

Funding *Vishesh Jain*: NSF CAREER award DMS-2237646

Clayton Mizgerd: NSF award ECCS-2217023

Acknowledgements We thank Huy Tuan Pham for helpful discussions and anonymous referees for several useful suggestions.

1 Introduction

Sampling and counting matchings in graphs is a central and well-studied problem. An early success in this direction is the classical algorithm of Kasteleyn for counting the number of perfect matchings in a planar graph [18]. Starting with the foundational work of Valiant [23], it was established that Kasteleyn’s algorithm is exceptional in the sense that it is #P-hard to (exactly) count the number of perfect matchings, even for restricted classes of input graphs such as bipartite graphs and graphs of bounded degree. In fact, perhaps quite surprisingly, the more general problem of counting matchings of a given size is #P-hard, even restricted to the class of planar graphs [15].

Given the above hardness results, the best one can hope for is fully polynomial-time (possibly randomized) approximation schemes. In particular, in connection with fully polynomial-time randomized approximation schemes (FPRAS) for the number of matchings (possibly of a given size), as well as being an important problem in its own right, much work has been devoted to the problem of approximately sampling from various distributions on matchings of a graph. The celebrated work of Jerrum and Sinclair [16] showed that for the monomer-dimer model at activity λ (i.e. the distribution on matchings where the probability



© Vishesh Jain and Clayton Mizgerd;

licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2024).

Editors: Amit Kumar and Noga Ron-Zewi; Article No. 63; pp. 63:1–63:13



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of a matching M is proportional to $\lambda^{|M|}$; λ is known as the activity), the Glauber dynamics mixes in time polynomial in n and λ . For graphs $G = (V, E)$ of bounded degree and $\lambda = O(1)$, the optimal mixing time $O(|E| \log n)$ was obtained by Chen, Liu, and Vigoda [6].

By combining with a rejection sampling procedure, both of these works give polynomial time algorithms to approximately sample from the uniform distribution on matchings of size $k \leq (1 - \delta)m^*(G)$ for any fixed $\delta > 0$, where $m^*(G)$ denotes the matching number of G i.e. the size of a largest matching in G ; approximately sampling from the uniform distribution on perfect matchings of a graph remains a major open problem, although in the bipartite case, this was famously resolved by Jerrum, Sinclair, and Vigoda [17]. For the class of bounded degree graphs, an algorithm with near-optimal run time was provided by a recent work of Jain, Perkins, Sah, and Sawhney [14]; they gave an algorithm which, given a graph G of maximum degree Δ , an integer $1 \leq k \leq (1 - \delta)m^*(G)$, and a parameter $\varepsilon > 0$, outputs a random matching M of size k in time $\tilde{O}_{\Delta, \delta}(n)$ ¹ such that the total variation distance is less than ε between the distribution on M and the uniform distribution on $\mathcal{M}_k(G)$: the matchings in G of size k .

Despite this progress, the mixing time of perhaps the simplest random walk on $\mathcal{M}_k(G)$ – the so-called down-up walk – is not understood. By the down-up walk for matchings of size k , we refer to the following chain:

1. Denote the state at time t by $M_t \in \mathcal{M}_k$.
2. Choose $e \in M_t$ and $e' \in E$ uniformly at random.²
3. Let $M' := M_t \cup \{e'\} \setminus \{e\}$. If $M' \in \mathcal{M}_k$, then $M_{t+1} = M'$. Else, $M_{t+1} = M_t$.

It is clear that the down-up walk is reversible with respect to the uniform distribution on \mathcal{M}_k so that whenever it is ergodic (this need not be the case; for instance, consider the uniform distribution on perfect matchings of an even cycle), it converges to the uniform distribution on \mathcal{M}_k . It is believed that for any fixed $\delta > 0$ the down-up walk on \mathcal{M}_k mixes in polynomial time for all $1 \leq k \leq (1 - \delta)m^*(G)$.³ In fact, it was conjectured by Jain, Perkins, Sah, and Sawhney [14, Conjecture 1.4] that for graphs G of maximum degree Δ and $1 \leq k \leq (1 - \delta)m^*(G)$, the ε -total-variation mixing time of the down-up walk on $\mathcal{M}_k(G)$ is $O_{\Delta, \delta}(n \log(n/\varepsilon))$, which would be optimal up to the implicit constants.

The main result of this note establishes that the down-up walk on $\mathcal{M}_k(G)$ mixes in polynomial time for all $1 \leq k \leq (1 - \delta)m^*(G)$. While our mixing time is unfortunately not sharp enough to resolve the aforementioned conjecture from [14], our result has the benefit of being applicable to arbitrary graphs (as opposed to graphs of bounded degree).

► **Theorem 1.** *Let $\delta \in (0, 1)$. For a graph $G = (V, E)$ on n vertices and m edges, and an integer $1 \leq k \leq (1 - \delta)m^*(G)$, the down-up walk on matchings of size k has ε -mixing time $O(n^{4/\delta} m^4 k \log(1/\varepsilon))$.*

► **Remark 2.** Restricted to the class of graphs of maximum degree Δ , our proof gives the improved ε -mixing time bound of $O_{\Delta, \delta}(n^\delta k \log(1/\varepsilon))$ by Equation (7) and Equation (1). We leave it as a very interesting open problem whether the mixing time can be improved to $\tilde{O}_{\Delta, \delta}(n)$ in this case (as was conjectured in [14]).

¹ \tilde{O} hides polylogarithmic factors in n and $1/\varepsilon$.

² Another convention is to choose e' uniformly at random among those edges for which $M_t \cup \{e'\} \setminus \{e\} \in \mathcal{M}_k$; in our setting, this would only have the effect of leading to a constant factor speed-up in the mixing time.

³ Some restriction on the range of k is needed since, as just mentioned, the down-up walk is not even ergodic in general.

Our proof is based on bounding the spectral gap using a carefully constructed flow. It is natural to ask whether the powerful spectral independence framework (developed in [3]) can be used to derive a similar result; in Section 3, we present examples showing that there are serious barriers to this, even for the class of bounded degree graphs. Roughly, the main point is that the condition $k \leq (1 - \delta)m^*(G)$ is not closed under pinnings (even if we take pinnings at random); this is not the case for the parameter range of independent sets considered in [13] and is key to making the spectral independence approach amenable in their setting.

1.1 Related work

For the down-up walk, notice that even the case when G is itself a matching is already interesting; in this case, the down-up walk coincides with the classical and well-studied Bernoulli-Laplace chain to sample from the uniform distribution on $\binom{[n]}{k}$ (e.g. [10, 19]). As discussed earlier, there are polynomial time algorithms, based on the rapid mixing of Glauber dynamics for the monomer-dimer model, to approximately sample from the uniform distribution on $\mathcal{M}_k(G)$, $1 \leq k \leq (1 - \delta)m_k^*(G)$; instead of combining rejection sampling with the Glauber dynamics, one may also combine rejection sampling with a local random walk to sample from the uniform distribution on the union of matchings of size k and $k - 1$ (the rapid mixing of this walk is shown in [8]). For bipartite graphs [17] and planar graphs [1], there are polynomial time algorithms to approximately sample from the uniform distribution on $\mathcal{M}_k(G)$ for all $1 \leq k \leq m^*(G)$.

Perhaps most relevant to this note is recent work of Jain, Michelen, Pham, and Vuong [13] which established optimal mixing of the down-up walk on *independent sets* of a given size $1 \leq k \leq (1 - \delta)\alpha_c(\Delta)n$, for the class of n vertex graphs G with maximum degree Δ (using the spectral independence framework). Here, $\alpha_c(\Delta)$ is a function such that the problem of (approximately) sampling independent sets of size $k > \alpha_c(\Delta)n$ on n vertex graphs with maximum degree Δ is computationally intractable, unless $\text{NP} = \text{RP}$; this was shown by Davies and Perkins [9]. In the same paper, Davies and Perkins showed that by combining the rapid mixing of the Glauber dynamics for the hard-core model in the tree uniqueness regime ([3, 6]) with a rejection sampling step, one can obtain a polynomial time algorithm to approximately sample from the uniform distribution on independent sets of size k provided that $1 \leq k \leq (1 - \delta)\alpha_c(\Delta)n$; this is entirely analogous to how [16, 6] imply polynomial time approximate samplers for the uniform distribution on $\mathcal{M}_k(G)$ for $1 \leq k \leq (1 - \delta)m^*(G)$. Davies and Perkins conjectured [9, Conjecture 5] that the down-up walk for independent sets mixes in polynomial time provided that $1 \leq k \leq (1 - \delta)\alpha_c(\Delta)n$ and this was resolved (in a stronger form) by [13]; our work may be viewed as resolving the analog of the conjecture of Davies and Perkins for matchings.

Finally, we remark that there is a large body of literature in probability concerned with the mixing of analogous walks for product(-like) domains with conservation laws (in our setting, the size of the matching is a conserved quantity); see, e.g., [5, 12] and the references therein. In our setting, the base measure (the natural choice is the monomer-dimer model at a suitable activity) is significantly more complicated and very far from being a product distribution, although we remark works are often able to exploit product structure and other symmetries to obtain rather precise results.

1.2 Organization

In Section 2, we present the proof of Theorem 1. In Section 3, we discuss barriers to a potential spectral independence approach for proving Theorem 1. In each (sub)section, we begin with an overview of the proof and some motivation.

2 Proof of Theorem 1

2.1 Preliminaries

Let P denote the transition matrix of an ergodic Markov chain on the finite state space Ω , which is reversible with respect to the (unique) stationary distribution π . Let $E(P) = \{(x, y) \in \Omega \times \Omega : P(x, y) > 0\}$ denote the “edges” of the transition matrix. Recall that the *Dirichlet form* is defined for $f, g : \Omega \rightarrow \mathbb{R}$ by

$$\mathcal{E}_P(f, g) := \frac{1}{2} \sum_{x, y \in \Omega} \pi(x) P(x, y) (f(x) - f(y))(g(x) - g(y)).$$

The *spectral gap* α is defined to be the largest value such that for all $\varphi : \Omega \rightarrow \mathbb{R}$,

$$\alpha \operatorname{Var}_\pi[\varphi] \leq \mathcal{E}_P(\varphi, \varphi).$$

The (total-variation) *mixing time* is defined by

$$\tau_{\text{mix}} = \max_{x \in \Omega} \min\{t : d(P^t x, \pi)_{\text{TV}} < 1/4\},$$

where $d(\cdot, \cdot)_{\text{TV}}$ denotes the total variation distance between probability distributions.⁴ The following relationship between the spectral gap and the mixing time is standard (see, e.g. [20]):

$$\tau_{\text{mix}} \leq \alpha^{-1} \log\left(\frac{1}{\min_{x \in \Omega} \pi(x)}\right). \quad (1)$$

In order to bound the spectral gap of the down-up walk, we will use the technology of multicommodity flows ([22, 11]).

► **Definition 3.** Consider the undirected graph $H = (\Omega, E(P))$. For $x, y \in \Omega$, let \mathcal{Q}_{xy} denote the set of all simple paths from x to y in H . Let $\mathcal{Q} = \bigcup_{x, y} \mathcal{Q}_{xy}$. A flow is a function $f : \mathcal{Q} \rightarrow \mathbb{R}_{\geq 0}$ such that $\sum_{q \in \mathcal{Q}_{xy}} f(q) = \pi(x)\pi(y)$. Given a flow f , we define its cost by

$$\rho(f) = \max_{(x, y) \in E(P)} \frac{1}{\pi(x)P(x, y)} \sum_{q \ni (x, y)} f(q),$$

and its length by

$$\ell(f) = \max_{q: f(q) > 0} |q|,$$

where $|q|$ denotes the number of edges in the path q .

It was shown in [22, 11] that any flow gives a lower bound on the spectral gap.

► **Theorem 4.** Let P be a reversible ergodic Markov chain and f be a flow. Then the spectral gap α satisfies

$$\alpha \geq 1/(\rho(f)\ell(f)).$$

For each $t \in E(P)$, let $\text{paths}(t) = \{q \ni t : f(q) > 0\}$. A common tool (see, e.g., [22]) for bounding $\rho(f)$ is a *flow encoding*, which is a collection of maps $\eta_t : \text{paths}(t) \rightarrow \Omega$ for all $t \in E(P)$. If all the maps are poly(n)-to-one and the measure π is “fairly tame” (the uniform measure on Ω automatically satisfies this condition), then this gives an inverse polynomial bound on the spectral gap.

⁴ Note that the quantity $1/4$ is fairly arbitrary here. Replacing $1/4$ with ε increases the mixing time by at most a factor of $\log_2(\varepsilon^{-1})$.

2.2 Constructing a flow

Recall that $1 \leq k \leq (1 - \delta)m^*(G)$ and Ω denotes the set of matchings in G of size exactly k . Given two matchings $x, y \in \Omega$, our flow will be constructed by uniformly distributing the demand $\pi(x)\pi(y)$ over a collection of carefully constructed paths. Compared to the construction of a flow in [16], we face two challenges:

- First, since we are not working with perfect matchings (or matchings which are a constant additive size away from being perfect), using just one path to route all the flow for each pair of matchings in the natural fashion results in a flow with exponentially high cost. To get around this issue, we use the (standard) idea of distributing the flow uniformly among essentially all possible paths, as is done for the Bernoulli-Laplace model (see [22]) and also for a random walk on the union of matchings of size k and $k - 1$ [8].
- Second – and this is the main new ingredient in our construction – our state space consists of matchings of a fixed size, whereas all previous walks and flow constructions (e.g. [16, 8]) required working with matchings of at least two adjacent sizes. In order to route flow along such paths while still incurring only polynomial cost, we divide pairs of matchings into a “good” set and a “bad” set depending on the combinatorial structure of the symmetric difference. For the good set, it is fairly simple to construct a flow, incorporating the above idea of distributing the flow uniformly among all possible paths. For a pair in the bad set (x, y) , we show that there is a nearby good pair (\tilde{x}, y) , in the sense that x can be transformed into \tilde{x} using a short path. The fact that we can transform x to a suitable \tilde{x} with a short path is key to bounding the cost of the flow, and this is where we use that $k \leq (1 - \delta)m^*(G)$.

Let $x \oplus y$ denote the symmetric difference $(x \setminus y) \cup (y \setminus x)$. Since x and y are each matchings and so have maximum degree 1, $x \oplus y$ is a disjoint union of paths and even-length cycles. For the sake of analysis, place arbitrary total orders on the set of even length paths in G and the set of even length cycles in G . Associate to each cycle one arbitrary distinguished vertex and to each odd-length path one arbitrary distinguished endpoint. These will all remain fixed for the remainder of the paper.

We partition $\Omega^2 := \Omega \times \Omega$ into $(\Omega^2)_g \cup (\Omega^2)_b$ where

$$(\Omega^2)_b = \{(x, y) : x \oplus y \text{ contains a cycle and no odd-length paths}\},$$

$$(\Omega^2)_g = \Omega^2 \setminus (\Omega^2)_b.$$

We will first describe the collection of paths between the “good pairs” $(\Omega^2)_g$. Later, we will leverage this collection on paths along with an additional idea to obtain a suitable collection of paths between the “bad pairs” $(\Omega^2)_b$.

Good pairs

Let $(x, y) \in (\Omega^2)_g$. The symmetric difference $x \oplus y$ consists of even length paths, even length cycles, odd length paths with more edges in y (which are necessarily x -augmenting paths), and odd length paths with more edges in x (which are necessarily y -augmenting paths). We have an induced ordering on the even paths and the cycles from our total order. Since $|x| = |y|$, $x \oplus y$ contains the same number of x -augmenting and y -augmenting paths; suppose there are $2j$ total odd-length paths. Let σ_x, σ_y be permutations of the sets of x -augmenting, y -augmenting paths respectively. For each such choice of (σ_x, σ_y) , we construct a path as follows from x to y in Ω . Before proceeding to the formal details, let us briefly describe the procedure: we first change x to y along all even paths. We then change x to y along the

first x -augmenting path $\sigma_x(1)$; at the end of such a path, there is an additional y -edge to be added, which gives us the necessary room to switch from x to y along all cycles, while still remaining in Ω . At the end of the cycle processing stage, there is an additional y -edge to be added; we pair this up with switching x to y along the first y -augmenting path $\sigma_y(1)$. Finally, we switch x to y along pairs of x -augmenting and y -augmenting paths $\sigma_x(i), \sigma_y(i)$ in the natural fashion.

Formally, set $M_0 = x$ and proceed as follows:

1. Process all even length paths in order. To process an even path, enumerate the edges $e_1, e_2, \dots, e_{2\ell}$ such that $e_i \cap e_{i+1} \neq \emptyset$ and $e_1 \in y$. This places all odd edges in y and the evens in x . Suppose t steps have been taken. First make the transition $M_{t+1} = M_t \cup e_1 \setminus e_2$,⁵ then $M_{t+2} = M_{t+1} \cup e_3 \setminus e_4$, and continue until $M_{t+\ell} = M_{t+\ell-1} \cup e_{2\ell-1} \setminus e_{2\ell}$. After processing all even paths, if we have reached y , terminate.
2. Process the first x -augmenting path $p = \sigma_x(1)$. Let $e^* \in p$ be the edge incident to the distinguished endpoint. Process $p \setminus e^*$ as an even path as in step (1), leaving only e^* to be added.
3. Process all cycles in order. For a cycle c , let e (respectively e') be the edge in $c \cap x$ (respectively $c \cap y$) incident to the distinguished vertex of c . First, let $M_{t+1} = M_t \cup e^* \setminus e$ to complete the previous path and puncture the cycle. Now, process $c \setminus \{e, e'\}$ as an even path as in step (1). Label $e^* := e'$ and process the next cycle in the same way. At the end of this step, some e^* will remain.
4. Process the first y -augmenting path $p = \sigma_y(1)$. Let $e \in p$ be the edge incident to the distinguished endpoint. Begin with $M_{t+1} = M_t \cup e^* \setminus e$, then process $p \setminus e$ as an even path as in step (1).
5. Process any remaining x -augmenting and y -augmenting paths in pairs $p = \sigma_x(i), p' = \sigma_y(i)$. Let e (respectively e') denote the edges incident to the distinguished endpoints of p (respectively p'). First process $p \setminus e$ as an even path, then exchange $M_{t+1} = M_t \cup e \setminus e'$, then process $p' \setminus e'$ as an even path.

This defines a unique path from x to y for any two permutations σ_x, σ_y , and so gives $(j!)^2$ total paths $x \rightarrow y$. We uniformly distribute the demand $\pi(x)\pi(y) = 1/|\Omega^2|$ by setting $f(q) = 1/(|\Omega|^2(j!)^2)$ for each path q thus defined.

Bad pairs

Given $(x, y) \in (\Omega^2)_b$, we will route the flow through $(\Omega^2)_g$ by choosing some suitable $(\tilde{x}, y) \in (\Omega^2)_g$ and adding a suitable prefix to all paths (as above) from \tilde{x} to y . Since $|x| \leq (1 - \delta)m^*(G)$, x has some augmenting path p of length at most $2\delta^{-1}$. This follows from the pigeonhole principle: for M^* a maximum matching in G , $x \oplus M^*$ is a graph with at most $2m^*(G)$ non-isolated vertices and at least $\delta m^*(G)$ disjoint x -augmenting paths, so that there must be an x -augmenting path of length at most $2\delta^{-1}$. Consider now $x^+ := x \oplus p$; this is a matching of size $k + 1$. We claim that there exists some $e \in x^+$ such that for $\tilde{x} := x^+ \setminus e$ satisfies $(\tilde{x}, y) \in (\Omega^2)_g$. To see this, note that since $|x^+| > |y|$, $x^+ \oplus y$ contains a y -augmenting path p' . If $x^+ \subset p'$, then $x^+ \oplus y$ cannot contain a cycle, as p' is an alternating path between edges in x^+ and in y , and $|x^+| = |y| + 1$, and so $y \subset p'$ as well. Thus $x^+ \oplus y = p'$ is a single path and contains no cycles, so we may choose any $e \in x^+$ and $(\tilde{x}, y) \in (\Omega^2)_g$. Otherwise, by choosing any edge $e \in x^+ \setminus p'$, we guarantee $\tilde{x} \oplus y$ has odd-length paths (in particular, p') and so $(\tilde{x}, y) \in (\Omega^2)_g$.

⁵ As a slight abuse of notation, when m is a matching, we will write $m \cup e$ to mean $m \cup \{e\}$ and similarly $m \setminus e$ instead of $m \setminus \{e\}$.

For every pair $(x, y) \in (\Omega^2)_b$, we make a fixed (but otherwise arbitrary) choice of p (an x -augmenting path of length at most $2\delta^{-1}$) and e as above. For a path $\tilde{q} \in \mathcal{Q}_{\tilde{x}y}$, we define $q \in \mathcal{Q}_{xy}$ as follows.

1. Process p as an x -augmenting path (see previous step 2), leaving some e^* to be added.
2. Make the exchange $M_{t+1} = M_t \cup e^* \setminus e$, arriving at \tilde{x} .
3. Follow the path \tilde{q} .

We assign $f(q) = f(\tilde{q})$ so that the same amount of flow is routed from x to y as from \tilde{x} to y . We remark that choosing an augmenting path of length $O_\delta(1)$ is crucial for bounding the cost of the flow below.

2.3 Flow encoding

For $t \in E(P)$, we now bound $f(t) := \sum_{q \ni t} f(q)$ using the method of flow encodings. Recall that $\text{paths}(t) = \{q \ni t : f(q) > 0\}$. Fix some transition $t = (z, z') \in E(P)$. We will partition $\text{paths}(t)$ into three sets and bound the contribution to $f(t)$ from each of the three using a ‘‘partial flow encoding’’. We have the ‘‘good’’ paths $\text{paths}_g(t)$ consisting of paths $q \in \text{paths}(t)$ whose endpoints are in $(\Omega^2)_g$. Recall that the paths in $(\Omega^2)_b$ consist of two phases: the prefix from $x \rightarrow \tilde{x}$, and then a good path from $\tilde{x} \rightarrow y$. Denote by $\text{paths}_a(t)$ those paths which use the transition t in the prefix $x \rightarrow \tilde{x}$, and by $\text{paths}_b(t)$ those paths which use the transition t in the path from $\tilde{x} \rightarrow y$. We will frequently need the set of short paths in G

$$\mathcal{P}_\delta := \left\{ (v_1 v_2 \cdots v_\ell) : \{v_i, v_{i+1}\} \in E(G), \ell \leq 2/\delta \right\}.$$

We will construct $\Omega \times \mathcal{P}_\delta$ -valued functions η_g, η_b, η_a on these subsets of $\text{paths}(t)$.

Construction of η_g

We first construct $\eta_g : \text{paths}_g(t) \rightarrow \Omega \times \mathcal{P}_\delta$. Let $t = (z, z') \in E(P)$. For a path q , let q^-, q^+ be the endpoints. Let $m = q^- \oplus q^+ \oplus (z \cup z')$. It is easily checked that m is a matching of size $k - 1$ and that $m' \oplus (z \cup z') = q^- \oplus q^+$ (the same construction is used in [16]). For consistency, we further map m into an element of $\Omega \times \mathcal{P}_\delta$ using a fixed (but otherwise arbitrary) m' -augmenting path $p \in \mathcal{P}_\delta$. The existence of a short m -augmenting path is guaranteed by the fact that $|m| < (1 - \delta)m^*(G)$. Formally, we have

$$\begin{aligned} \eta_g : \text{paths}_g(t) &\rightarrow \Omega \times \mathcal{P}_\delta \\ q &\mapsto (q^- \oplus q^+ \oplus (z \cup z') \oplus p, p). \end{aligned}$$

Note that given the image $(m \oplus p, p)$, we take $(m' \oplus p) \oplus p$ to recover m , which then recovers $q^- \oplus q^+$ as before. We can now reindex the sum

$$\sum_{q \in \text{paths}_g(t)} f(q) = \sum_{(m,p) \in \Omega \times \mathcal{P}_\delta} \sum_{q \in \eta_g^{-1}(m,p)} f(q).$$

The endpoints of every $q \in \eta_g^{-1}(m, p)$ have the same symmetric difference as noted above. Let this symmetric difference have $2j$ odd-length paths. By our construction of the flow, $f(q) = |\Omega|^{-2} (j!)^{-2}$ for all $q \in \eta_g^{-1}(m, p)$. We now count how many paths use the transition (z, z') based on which G -paths it is processing. This requires some case analysis, but ultimately, is based on blending the analysis of the flow encoding for the Glauber dynamics for the monomer-dimer model in [16] with the flow encoding for the Bernoulli-Laplace model [22].

- Case I:** If (z, z') is processing even length paths, then we can use our total order on even length paths to identify which parts of each cycle belong to q^- and to q^+ . We know that we have not yet begun processing odd paths or cycles, so the parts in z belong to q^- and those outside z belong to q^+ . We thus know q^- and q^+ and so there are exactly $(j!)^2$ paths using (z, z') .
- Case II:** If (z, z') finishes processing an odd-length G -path and begins processing a cycle, then we know the odd G -path is the first such processed, and by the same reasoning as before, we can deduce the endpoints q^- and q^+ . We also know $\sigma_x(1)$ is the path intersecting $z \oplus z'$. The remainder of σ_x and the entirety of σ_y is free, so there are $(j!)^2/j$ paths using (z, z') .
- Case III:** If (z, z') is processing entirely cycles, then z has a perfect matching on $j + 1$ of the odd-length paths and the interior edges on $j - 1$ of the odd-length paths. One of these G -paths has already been augmented, and then there will be a path for every ordering of the remaining G -paths. Thus there are $(j + 1)j!(j - 1)! = (1 + 1/j)(j!)^2$ paths.
- Case IV:** If (z, z') finishes a cycle and begins an odd-length path, then z has a perfect matching on $j + 1$ odd paths (one of which is being de-augmented in (z, z')) and the interior edges on $j - 1$ odd paths. The path touched by $z \oplus z'$ is $\sigma_y(1)$. We must choose one of the remaining j perfectly matched paths to be $\sigma_x(1)$, and then the remainder of σ_x, σ_y are free on the sets of $j - 1$ interior, perfect paths respectively. There are thus $j((j - 1)!)^2 = (j!)^2/j$ paths using (z, z') .
- Case V:** If (z, z') is augmenting an odd-length path, then of the other $2j - 1$ paths, z is perfect on j and interior on $j - 1$. Suppose $2r$ paths have already been processed. Then we may choose the already-augmented paths ($\binom{j}{r}$ choices), the already-de-augmented paths ($\binom{j-1}{r}$ choices), the order for each ($(r!)^2$ choices), and the order for the remaining augmentations and de-augmentations ($(j - r)!(j - 1 - r)!$ choices). This gives $j!(j - 1)!$ paths through (z, z') that have already processed r pairs of G -paths. We now sum over $0 \leq r \leq j - 1$ to get $j(j!)(j - 1)! = (j!)^2$ total paths through (z, z') .
- Case VI:** If (z, z') is de-augmenting an odd-length path, then by the same logic but with $j - 1$ perfect paths and j interior paths, we again have $(j!)^2$ total paths.
- Case VII:** If (z, z') finishes augmenting one odd path and begins de-augmenting the next, then we know these G -paths occur adjacently in the path q . Then by the same reasoning as Case V but with $j - 1$ perfect paths and $j - 1$ interior paths, we will get $(j!)^2/j$ total paths through (z, z') .

In all cases, we have at most $(1 + 1/j)(j!)^2 \leq 2(j!)^2$ paths in $\eta_g^{-1}(m, p)$ that use the transition (z, z') . As each has weight $f(q) = |\Omega|^{-2}(j!)^{-2}$, this means the inner sum is at most $2/|\Omega|^2$ and so we can bound

$$\sum_{q \in \text{paths}_g(t)} f(q) = \sum_{(m,p) \in \Omega \times \mathcal{P}_\delta} \sum_{q \in \eta_g^{-1}(m,p)} f(q) \leq \sum_{(m,p) \in \Omega \times \mathcal{P}_\delta} \frac{2}{|\Omega|^2} = \frac{2|\mathcal{P}_\delta|}{|\Omega|}. \quad (2)$$

Construction of η_b

Recall that $\text{paths}_b(t)$ are those paths that use the transition t as part of following a “good” path from \tilde{x} to y . Thus we may instead choose the good path q that is routed through t , and then count how many starting points x could route through q^- (the starting point of q) to get to q^+ (the ending point of q). By the construction of our paths, this requires $x \oplus q^-$ to be a single short augmenting G -path in \mathcal{P}_δ together with a single edge in $E(G)$. These

together will uniquely determine the total path $x \rightarrow q^+$. Thus each good path through t is used in at most $|\mathcal{P}_\delta| |E(G)|$ bad paths, and the value of f is unchanged by the prefix, so we may bound using Equation (2)

$$\sum_{q \in \text{paths}_b(t)} f(q) \leq \sum_{\tilde{q} \in \text{paths}_g(t)} |\mathcal{P}_\delta| |E(G)| f(\tilde{q}) \leq \frac{2|\mathcal{P}_\delta|^2 |E(G)|}{|\Omega|}. \quad (3)$$

Construction of η_a

Finally, for $t = (z, z^-)$ and $q \in \text{paths}_a(t)$, let p be the G -path that is augmented during the prefix. Then, define the function

$$\begin{aligned} \eta_a : \text{paths}_a(t) &\rightarrow \Omega \times \mathcal{P}_\delta \\ q &\mapsto (q^+, p). \end{aligned}$$

Suppose $q \in \eta_a^{-1}(m, p)$ for some matching $m \in \Omega$ and G -path $p \in \mathcal{P}_\delta$ and let $t = (z, z^-)$. Then we know that $q^+ = m$, and we know that q^- consists of $z \setminus p$ and the interior alternating edges of p . Thus all paths in $\eta_a^{-1}(m, p)$ have the same endpoints, and so their total flow is at most the net flow between those two points, which is $|\Omega|^{-2}$. We can then calculate

$$\sum_{q \in \text{paths}_a(t)} f(q) = \sum_{(m,p) \in \Omega \times \mathcal{P}_\delta} \sum_{q \in \eta_a^{-1}(m,p)} f(q) \leq \sum_{(m,p) \in \Omega \times \mathcal{P}_\delta} \frac{1}{|\Omega|^2} = \frac{|\mathcal{P}_\delta|}{|\Omega|}. \quad (4)$$

Bounding the cost of the flow

Using Equation (2), Equation (3), and Equation (4), for any transition $t \in E(P)$,

$$\sum_{q \in \text{paths}(t)} f(q) = \sum_{q \in \text{paths}_g(t)} f(q) + \sum_{q \in \text{paths}_b(t)} f(q) + \sum_{q \in \text{paths}_a(t)} f(q) \leq \frac{3|\mathcal{P}_\delta| + 2|\mathcal{P}_\delta|^2 |E(G)|}{|\Omega|}. \quad (5)$$

Since $\pi(z) = 1/|\Omega|$ for all $z \in \Omega$ and $P(z, z') \geq 1/(k|E(G)|)$ (since the possible transitions from z consist of removing one of k edges and adding one of $|E(G)|$ edges), we get that

$$\rho(f) \leq |\Omega| \cdot k|E(G)| \cdot \frac{3|\mathcal{P}_\delta| + 2|\mathcal{P}_\delta|^2 |E(G)|}{|\Omega|} \leq 3k|E(G)|^2 |\mathcal{P}_\delta|^2.$$

Finally, let Δ denote the maximum degree of G and note that $|\mathcal{P}_\delta| \leq 2n\Delta^{2/\delta-1}$ to get that

$$\rho(f) \leq 12k|E(G)|^2 \cdot n^2 \Delta^{4/\delta-2}. \quad (6)$$

2.4 Rapid mixing

We will use Theorem 4 to bound the spectral gap via the flow f defined in Section 2.2. Note that the down-up walk is reversible with respect to the uniform distribution, aperiodic since $P(x, x) > 0$, and irreducible (for instance, by using the paths used in our flow f). Therefore, the assumptions of Theorem 4 are satisfied. To bound the maximum length $\ell(f)$ of any path used in our flow, note that by construction, any edge in G is included in at most three exchanges. Hence, $\ell(f) \leq 3|E(G)|$. Combining this with Equation (6), we see that the spectral gap α of the down-up walk satisfies

$$\alpha^{-1} \leq 36k|E(G)|^3 \cdot n^2 \Delta^{4/\delta-2}. \quad (7)$$

Finally, the mixing time bound in Theorem 1 follows from Equation (1) by noting that $\log |\Omega| \leq \log 2^{|E(G)|} \leq |E(G)|$.

3 Barriers to the spectral independence approach

For a distribution π on $\binom{[n]}{k}$, we define the (signed) pairwise influence matrix $M_\pi \in \mathbb{R}^{n \times n}$ by

$$M_\pi(i, j) = \begin{cases} 0 & \text{if } j = i, \\ \mathbb{P}_\pi[j | i] - \mathbb{P}_\pi[j | \bar{i}] & \text{otherwise.} \end{cases}$$

where $\mathbb{P}_\pi[i] = \mathbb{P}_{S \sim \pi}[i \in S]$ and $\mathbb{P}_\pi[\bar{i}] = \mathbb{P}_{S \sim \pi}[i \notin S]$. We say that π is η -spectrally independent (at link \emptyset) if $\lambda_{\max}(M_\pi) \leq \eta$ and that π is η - ℓ_∞ -independent (at link \emptyset) if $\max_{i \in [n]} \sum_{j=1}^n |M_\pi(i, j)| \leq \eta$. Note that the latter condition implies the former.

We begin by noting that for the class of bounded degree graphs, for k bounded away from the matching number, the uniform distribution on matchings of size k is $O(1)$ - ℓ_∞ -independent.

► **Proposition 5.** *Let $G = (V, E)$ be a graph on n vertices with maximum degree Δ . Let $\delta > 0$ and for $1 \leq k \leq (1 - \delta)m^*(G)$, let π be the uniform distribution on matchings of G of size k . Then π is $O_{\delta, \Delta}(1)$ - ℓ_∞ -independent (at link \emptyset).*

Proof. For $k = o(n)$, this is implied by a coupling argument (e.g. [21]). For $k = \Omega(n)$, the proof follows from the same argument as in the proof of [13, Theorem 8]: the differences are that we compare to the monomer-dimer model at activity $\lambda = O_{\delta, \Delta}(1)$ using [14, Lemma 4.1], replace [13, Theorem 9] by [6, Theorem 2.10], and replace [13, Theorem 15] by a suitable multivariate zero-free region for the matching polynomial (e.g. [7]). We omit further details. ◀

Given Proposition 5, one might hope to obtain an inverse polynomial bound on the spectral gap of the down-up walk (at least for the class of bounded degree graphs) using the powerful spectral independence framework as is done, for instance, in the case of the down-up walk on independent sets of a fixed size in [13]; we refer the reader to [3] for an introduction to this framework. In order to do this, we need to show that the distribution remains $O_{\delta, \Delta}(1)$ -spectrally independent under any pinning. In our situation, a pinning τ is a matching of size $\ell < k$. We would then consider $\Omega_\tau = \{m \in \Omega : \tau \subset m\}$ under the distribution induced by π (uniform, in our case), and show $O_{\delta, \Delta}(1)$ -spectral independence of this space. We note that there is a more powerful “average-case” version of this argument, which (roughly) allows us to consider typical pinning obtained by starting from some fixed matching of size k and pinning a random subset of $k - \ell$ edges to be included (see [4, 2]). We present barriers to this approach.

- We observe that such an approach cannot work for the down-up walk. Indeed, if it were to work, then one would also be able to show that the down-up walk has inverse polynomial spectral gap for the induced uniform distribution on size ℓ matchings obtained by starting with an arbitrary matching of size k and pinning a uniform subset of $k - \ell$ edges to belong to the matching. However, as we discuss below, it is easy to construct an example where even for polynomially large ℓ , with high probability, the down-up walk is not even ergodic (Claim 7).
- In the above example, the failure of ergodicity may be circumvented by using an $O(1)$ -step down-up walk. However, it is still the case that proving mixing of the $O(1)$ -step down-up walk using the (average) spectral independence framework necessitates proving mixing for the $O(1)$ -step down-up walk for the aforementioned induced distributions on matchings of size ℓ . We present a construction (Claim 6) showing that these induced distributions can correspond to the uniform distribution on size ℓ matchings in pretty arbitrary graphs

with matching number $\ell(1 + o(1))$; hence, there does not seem to be a way to use this machinery without basically showing that $O(1)$ -step down-up walks mix rapidly for (almost) maximum matchings in arbitrary bounded-degree graphs, which is a major open problem.

Our examples will follow the same general template. To set up some notation, given a graph $G = (V, E)$ and a pinning τ (a matching τ in G), define the *residual graph* G_τ to be the induced subgraph

$$G_\tau = G[V(G_\tau)],$$

where

$$V(G_\tau) = V(G) \setminus \bigcup_{e \in \tau} e.$$

Sampling from the uniform distribution on matchings in G of size k , conditioned on pinning τ to be in the matching, is equivalent to finding a matching of size $k - |\tau|$ in G_τ .

We are now ready to construct our examples. Fix some $0 < \delta < 1/5$ the desired gap from maximality, as in the statement of Theorem 1. We define a graph $G = (V, E)$ where $|V| = n$ as follows: G consists of $\delta n/2$ disjoint copies of P_9 (the path with 10 vertices and 9 edges) and an arbitrary graph G' on the remaining $(1 - 5\delta)n$ vertices such that G' has a perfect matching. Let M be the matching given by taking the union of a perfect matching M' in G' with the interior alternating edges on each P_9 ; note that $|M| = n/2 - \delta n/2 = (1 - \delta)m^*(G)$. We will be considering pinning a uniform random subset of M of a fixed size.

▷ **Claim 6.** For a random pinning τ of size $(1 - \lambda)|M|$, the ratio

$$\mathbb{E} \left[1 - \frac{|M| - |\tau|}{m^*(G_\tau)} \right] = O(\delta\lambda^4).$$

The implication of this claim is that, while we started with the uniform distribution on matchings of size at most $(1 - \delta)$ of the maximum matching in G , we now need to deal with the uniform distribution on matchings of size at least $(1 - \delta\lambda^4)$ times the maximum matching in G_τ .

Proof. Let τ be a random pinning of size $(1 - \lambda)|M|$. Let X_τ be the number of P_9 s that τ does not intersect. Then by linearity of expectation,

$$\mathbb{E}[X_\tau] = \frac{\delta n}{2} \Pr[\tau \text{ avoids a fixed } P_9] = O(\delta n \lambda^5). \tag{8}$$

The key observation here is that once we have pinned any edge in $M \cap P_9$ for some copy of P_9 , we have split P_9 into two even paths and are demanding a maximum matching on each of those. Hence, by construction, we see that $m^*(G_\tau) = |M| - |\tau| + X_\tau$. We now compute

$$\begin{aligned} \mathbb{E} \left[1 - \frac{|M| - |\tau|}{m^*(G_\tau)} \right] &= 1 - \mathbb{E} \left[\frac{|M| - |\tau|}{|M| - |\tau| + X_\tau} \right] \leq 1 - \frac{|M| - |\tau|}{|M| - |\tau| + \mathbb{E}[X_\tau]} \\ &\leq \frac{\mathbb{E}[X_\tau]}{\lambda|M|} = O(\delta\lambda^4), \end{aligned}$$

where the first line uses Jensen's inequality and the second line uses Equation (8). ◀

In the above construction, take G' to be a disjoint union of $(1 - 5\delta)n/4$ copies of C_4 (the 4-cycle) and accordingly, take M' to be a union of perfect matchings on each C_4 .

63:12 Rapid Mixing of the Down-Up Walk on Matchings of a Fixed Size

▷ Claim 7. For a random pinning τ of size $n/2 - n^{2/3}$, with high-probability, the down-up walk on matchings of size $|M| - \tau$ on the induced graph G_τ is not ergodic.

Proof. It suffices to show that for a random pinning τ of size $n/2 - n^{2/3}$, with high probability, (i) τ intersects every P_9 , (ii) τ fails to intersect some C_4 .

For (i), by a union bound and direct computation, we get that

$$\mathbb{P}[\tau \text{ avoids some } P_9 \cap M] \leq \frac{\delta n}{2} \mathbb{P}[\tau \text{ avoids a fixed } P_9 \cap M] = O(\delta n^{-1/3}).$$

For (ii), we get that

$$\begin{aligned} \mathbb{P}[\tau \text{ intersects all } C_4] &= \mathbb{P}[\tau^c \text{ contains no } C_4 \cap M] \\ &\leq n^{O(1)} \mathbb{P}[\tau^c \text{ does not contain a fixed } C_4 \cap M]^{(1-5\delta)n/4} \\ &\leq \exp(-\Theta(n^{1/3})), \end{aligned}$$

where the second follows by comparing probabilities between the independent model of density $\Theta(n^{-1/3})$ and the slice model and the last line follows by direct computation. The union bound now shows that with high probability, (i) and (ii) simultaneously hold. ◁

References

- 1 Yeganeh Alimohammadi, Nima Anari, Kirankumar Shiragur, and Thuy-Duong Vuong. Fractionally log-concave and sector-stable polynomials: counting planar matchings and more. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 433–446, 2021.
- 2 Nima Anari, Vishesh Jain, Frederic Koehler, Huy Tuan Pham, and Thuy-Duong Vuong. Universality of spectral independence with applications to fast mixing in spin glasses. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5029–5056. SIAM, 2024.
- 3 Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. In *IEEE Symposium on Foundations of Computer Science*, 2020.
- 4 Nima Anari, Yang P Liu, and Thuy-Duong Vuong. Optimal sublinear sampling of spanning trees and determinantal point processes via average-case entropic independence. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 123–134. IEEE, 2022.
- 5 Pietro Caputo. Spectral gap inequalities in product spaces with conservation laws. In *Stochastic analysis on large scale interacting systems*, volume 39, pages 53–89. Mathematical Society of Japan, 2004.
- 6 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Optimal mixing of glauber dynamics: Entropy factorization via high-dimensional expansion. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1537–1550, 2021.
- 7 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Spectral independence via stability and applications to holant-type problems. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 149–160. IEEE, 2022.
- 8 Paul Dagum, Michael Luby, Milena Mihail, and U Vazirani. Polytopes, permanents and graphs with large factors. In *[Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science*, pages 412–421. IEEE Computer Society, 1988.
- 9 Ewan Davies and Will Perkins. Approximately counting independent sets of a given size in bounded-degree graphs. In *48th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 198, pages 62:1–62:18, 2021.

- 10 Persi Diaconis and Mehrdad Shahshahani. Time to reach stationarity in the bernoulli–laplace diffusion model. *SIAM Journal on Mathematical Analysis*, 18(1):208–218, 1987.
- 11 Persi Diaconis and Daniel Stroock. Geometric bounds for eigenvalues of markov chains. *The annals of applied probability*, pages 36–61, 1991.
- 12 Yuval Filmus, Ryan O’Donnell, and Xinyu Wu. Log-sobolev inequality for the multislice, with applications. *Electronic Journal of Probability*, 27:1–30, 2022.
- 13 Vishesh Jain, Marcus Michelen, Huy Tuan Pham, and Thuy-Duong Vuong. Optimal mixing of the down-up walk on independent sets of a given size. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1665–1681. IEEE, 2023.
- 14 Vishesh Jain, Will Perkins, Ashwin Sah, and Mehtaab Sawhney. Approximate counting and sampling via local central limit theorems. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1473–1486, 2022.
- 15 Mark Jerrum. Two-dimensional monomer-dimer systems are computationally intractable. *Journal of Statistical Physics*, 48:121–134, 1987.
- 16 Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18(6):1149–1178, 1989.
- 17 Mark Jerrum, Alistair Sinclair, and Eric Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *Journal of the ACM (JACM)*, 51(4):671–697, 2004.
- 18 Pieter Kasteleyn. Graph theory and crystal physics. *Graph Theory and Theoretical Physics*, pages 43–110, 1967.
- 19 Tzong-Yow Lee and Horng-Tzer Yau. Logarithmic sobolev inequality for some models of random walks. *The Annals of Probability*, 26(4):1855–1873, 1998.
- 20 David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.
- 21 Kuikui Liu. From coupling to spectral independence and blackbox comparison with the down-up walk. *arXiv preprint*, 2021. [arXiv:2103.11609](https://arxiv.org/abs/2103.11609).
- 22 Alistair Sinclair. Improved bounds for mixing rates of markov chains and multicommodity flow. *Combinatorics, probability and Computing*, 1(4):351–370, 1992.
- 23 Leslie G Valiant. The complexity of computing the permanent. *Theoretical computer science*, 8(2):189–201, 1979.