

Sampling Edge Covers in 3-Regular Graphs

Ivona Bezáková and William A. Rummmler

Rochester Institute of Technology, Rochester, NY, USA
{ib,war5549}@cs.rit.edu

Abstract. An edge cover C of an undirected graph is a set of edges such that every vertex has an adjacent edge in C . We show that a Glauber dynamics Markov chain for edge covers mixes rapidly for graphs with degrees at most three. Glauber dynamics have been studied extensively in the statistical physics community, with special emphasis on lattice graphs. Our results apply, for example, to the hexagonal lattice. Our proof of rapid mixing introduces a new cycle/path decomposition for the canonical flow argument.

1 Introduction

An edge cover of an undirected graph is a subset of its edges such that every vertex has an adjacent edge in the edge cover (i. e., each vertex is “covered”). We initiate the study of sampling (and the related question of counting) of all edge covers of a graph. Our main motivation for this work is to develop new insights that might help with understanding other related combinatorial problems such as matchings and contingency tables.

Glauber dynamics Markov chains on lattice graphs have been studied extensively in both the computer science and statistical physics communities. These single-site-update Markov chains tend to be natural to design but difficult to analyze, even though they are often believed to mix rapidly. Examples of successful theoretical analysis include the works of Randall and Tetali [1] and Vigoda [2]. Randall and Tetali, building on the result of Diaconis and Saloff-Coste [3], proved rapid mixing of Glauber dynamics Markov chains for domino tilings on the grid graph and lozenge tilings on the triangular lattice (these problems can be viewed as perfect matchings on the grid graph and the hexagonal lattice, respectively). Vigoda showed $O(n \log n)$ mixing time of a Glauber dynamics Markov chain for independent sets of graphs with maximum degree 4.

Edge covers are related to matchings in a similar way as vertex covers are related to independent sets. A vertex cover is a set of vertices such that every edge has an end-point in the set. The complement of a vertex cover is an independent set, i. e., a set of vertices that do not share any edges between them. The relationship between edge covers and matchings is not as straightforward but it is still easy to phrase: the size of the maximum matching plus the size of the minimum edge cover equals the number of vertices of the graph. While there is an obvious relationship between the counts of independent sets and vertex covers, no relationship

is known between the respective counts of edge covers and matchings. The problem of sampling and counting of matchings received a lot of attention [4,5,6,1,7], however, the problem remains unsolved for perfect matchings of arbitrary (i. e., non-bipartite) graphs. A perfect matching can be interpreted as an edge cover where every vertex is covered by exactly one edge. Hence, edge covers may be seen as a natural variation on perfect matchings. Jerrum and Sinclair [5] gave an fpras (fully polynomial randomized approximation scheme) for the problem of sampling all matchings (including matchings that are not perfect) of any graph (bipartite or not); however, the problem of sampling all edge covers appears more challenging due to its similarity to the problem of sampling subgraphs with a given degree sequence. The subgraph problem with bipartite input is also known as the binary, or 0/1, contingency tables.

Technical issues arise once vertices are allowed to have degrees larger than 1. In the matching problem, every vertex has degree one (or at most one) and this property helps with applying the canonical path argument for rapid mixing: namely, one needs to define a way of getting from any matching to any other matching using the transitions of the Markov chain (and this needs to be done so that no transition is “overloaded”). The symmetric difference of any two matchings forms a set of alternating cycles and paths (i. e., the edges strictly alternate between the two matchings). This property helps with transforming one matching into the other matching by handling the cycles and paths one by one. With subgraphs that satisfy a given degree sequence, this becomes more challenging but a fairly natural decomposition into cycles and paths is still possible due to the fact that at every vertex, the two subgraphs have the same degree and thus, the edges from the first subgraph can be paired with the edges of the second subgraph. The pairing notion was introduced by Cooper, Dyer, and Greenhill [8] who were building on ideas by Jerrum and Goldberg, and Kannan, Tetali, and Vempala [9].

In contrast, two different edge covers can have different degrees at the same vertex. In this paper we overcome this problem for graphs with degrees up to 3. Our cycle/path decomposition has some interesting characteristics the previous works did not have. We overcome the disproportionateness of the degrees by allowing pairing of edges from the same edge cover. Thus, the components of our decomposition might be non-alternating – they could consist of arbitrary sequences of edges from the two edge covers, including the possibility of all edges coming from the same edge cover. The alternating property is used to keep the intermediate configurations on a canonical path within the state space. Thus, our non-alternating decomposition needs to be done carefully so that the configurations that arise during the process of changing one edge cover into the other edge cover are edge covers as well.

While our restriction on the degrees is rather weak, we hope that the ideas in this work will bring new perspective to other related problems. For example, the results extend to so-called 1-2-edge covers (where every vertex is covered by one or two edges) for certain classes of graphs with arbitrary degrees; these edge covers can be viewed as close relatives of perfect matchings. We defer this extension to the journal version of this paper.

Our main result is:

Theorem 1. *Let G be an undirected graph with n vertices and degrees upper bounded by 3 and let Ω be the set of all its edge covers. Then, there exists a Markov chain with uniform stationary distribution over Ω with polynomial mixing time. More precisely,*

$$\tau(\delta) = O(n^6(n + \log \delta^{-1})),$$

where $\tau(\delta)$ is the mixing time of the chain¹ and $\delta > 0$ is the sampling error.

2 Preliminaries

Problem definition. Given is an undirected graph $G = (V_G, E_G)$. An **edge cover** of G is a set of edges $C \subseteq E_G$ such that for every vertex $u \in V_G$, there exists $(u, v) \in C$ for some $v \in V_G$. The **sampling edge covers** problem asks for a randomized polynomial-time algorithm that outputs a uniformly random edge cover of G . The output of the **counting** version of the problem is the number of all edge covers of G .

Markov chain Monte Carlo. In this section we give a very brief overview of the basic Markov chain terminology.

Let (Ω, P) be a Markov chain with state space Ω and transition matrix $P = (p_{i,j})_{|\Omega| \times |\Omega|}$ denoting the probability of going from state $i \in \Omega$ to state $j \in \Omega$. A distribution σ on Ω is called **stationary** if $\omega P = \omega$. A Markov chain is **ergodic** if it is **irreducible** (it is possible to get from every state to every other state with nonzero probability) and **aperiodic** (for every state i , the numbers of steps one can use to get from i to i have greatest common divisor equal to 1). It is well known that for ergodic Markov chains the stationary distribution is unique.

A Markov chain is **reversible** if $\pi(i)p_{i,j} = \pi(j)p_{j,i}$ for every $i, j \in \Omega$ and stationary distribution π . It is not difficult to verify that for an ergodic Markov chain with a symmetric transition matrix P , the stationary distribution is uniform.

The **total variation distance** between two distributions μ, ν on Ω is given by

$$d_{tv}(\mu, \nu) = \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \nu(x)|.$$

The **mixing time** $\tau_x(\delta)$ of the chain starting at state $x \in \Omega$ is defined as

$$\tau_x(\delta) = \min\{t \geq 0 \mid d_{tv}(P^t(x, \cdot), \pi) \leq \delta\},$$

where $P^t(x, \cdot)$ denotes the distribution after t steps of the chain, starting at the state $x \in \Omega$. Moreover, $\tau(\delta) = \max_{x \in \Omega} \tau_x(\delta)$.

¹ We opted for clarity of presentation over the best possible running time estimate.

One of the techniques for bounding mixing time is called the **canonical path/flow technique**. The canonical path technique requires, for every pair of states $I, F \in \Omega$, to define a path between I and F through other states such that the probability of moving between adjacent states is nonzero, i. e., we define $I = \omega_1, \omega_2, \dots, \omega_k = F$ such that $p_{\omega_i, \omega_{i+1}} > 0$ for every $i \in \{1, \dots, k-1\}$. In the canonical flow technique instead of constructing a single path from I to F , we can define a flow from I to F , i. e., we have the option of defining a set of paths $\mathcal{P}_{I,F}$ from I to F , and for every path $p \in \mathcal{P}_{I,F}$ we set its weight $g(p)$ so that $\sum_{p \in \mathcal{P}_{I,F}} g(p) = 1$. Whenever the volume of paths going through any given transition is not too large, the chain mixes in polynomial time. This notion is formalized as follows, see [10,11] for details. Let us consider all possible paths between any pair of states and let g be the weights (flow values) of a concrete canonical flow instance (i. e., whenever a path is not used by the flow, its weight is set to 0). Moreover, let π be the stationary distribution of the chain. The **congestion** through a transition $T = (M, M')$, $p_{M,M'} > 0$ is defined as

$$\rho_g(T) = \frac{1}{\pi(M)p_{M,M'}} \sum_{p \ni T} \pi(I_p)\pi(F_p)g(p)\ell(g),$$

where the sum ranges through all paths that use T , the path p starts at I_p and ends at F_p , and $\ell(g)$ is the maximum length of any path of nonzero weight in g . Then, the overall congestion is

$$\rho_g = \max_{T=(M,M'):p_{M,M'}>0} \rho_g(T)$$

and the mixing time is bounded by

$$\tau_x(\delta) \leq \rho_g(\log \pi(x)^{-1} + \log \delta^{-1}).$$

3 Results

In this section we prove the main theorem of the paper.

3.1 A Markov Chain on Edge Covers

Let $G = (V, E)$ be a graph and let Ω be the set of all of its edge covers. Let $X_i \in \Omega$ be the current state of the Markov chain. The next state is the result of a simple Glauber dynamics-type move:

1. With probability 1/2, let $X_{i+1} = X_i$.
2. Otherwise, choose an edge $e \in E_G$ uniformly at random.
3. If $e \notin X_i$, let $X_{i+1} = X_i \cup \{e\}$.
4. Else, if $X_i \setminus \{e\}$ is an edge cover of G , let $X_{i+1} = X_i \setminus \{e\}$.
5. Else, let $X_{i+1} = X_i$.

In words, the Markov chain chooses a random edge and if it is not in the edge cover, it adds it. If it is already in the edge cover, it removes it if it can (if the result is still a valid edge cover). This Markov chain is symmetric and thus its stationary distribution π is the uniform distribution on Ω . The lazy step (step 1. of the chain) ensures that the chain is aperiodic, and its irreducibility will follow from the definition of the canonical flow. Thus, the chain is ergodic and its stationary distribution is unique.

3.2 Canonical Flow for This Markov Chain

We will define canonical flow between every pair of states $I, F \in \Omega$. Let $I \oplus F$ be the symmetric difference of I and F . We will gradually change the I -edges in $I \oplus F$ into the F -edges (we will leave the edges shared by both I and F , i. e., the edges in $I \cap F$, unchanged).

Cycle/path decomposition. We will first decompose $I \oplus F$ into a set of cycles and paths. Proofs of rapid mixing for matchings and subgraphs with prescribed degree sequence (in the bipartite case known also as binary contingency tables) use decomposition into **alternating** cycles and paths. Our decomposition will **not** necessarily contain only alternating cycles and paths.

Let $v \in V$ be a vertex and let c_I and c_F be the numbers of I -edges and F -edges, respectively, incident to v in $I \oplus F$. For 3-regular graphs we have $c_I, c_F \leq 3$. We will construct a **pairing of the edges in $I \oplus F$ at v** as follows. Without loss of generality assume that $c_I \leq c_F$. We will pair all the I -edges with the same number of F -edges. Then, we will pair the remaining F edges with other F edges (hence, the pairing will not be alternating). More precisely, for every I -edge e_I incident to v we select an F -edge e_F incident to v and pair them together. We will be left with $c_F - c_I$ edges (all F -edges) and, if $c_F - c_I = 2$, we will take two of these edges and pair them with each other. If $c_F - c_I = 1$, we leave the remaining F -edge unpaired. (Notice that for 3-regular graphs, $|c_F - c_I| \leq 2$. This follows from the fact that I and F are valid edge covers: suppose that $|c_F - c_I| = 3$, then either F contains all edges incident to v and I contains none, or vice versa. However, either case contradicts the assumption that both I and F are edge covers.) We denote the set of all pairings of the edges in $I \oplus F$ at v by $\Psi_{I,F}(v)$.

A **pairing of $I \oplus F$** is a set of valid pairings at every vertex $v \in V$. We will denote the set of all (possible) pairings of $I \oplus F$ by $\Psi_{I,F}$.

The **cycle/path decomposition of $I \oplus F$ with respect to a pairing from $\Psi_{I,F}$** is defined as follows. We take an arbitrary edge in $I \oplus F$, let it be (u_0, u_1) . We find the edge (u_0, u_1) is paired at u_1 with, let it be (u_1, u_2) ; similarly, we find the edge (u_{-1}, u_0) that is paired with (u_0, u_1) at u_0 (if these edges exist – it might also be that (u_0, u_1) does not have a paired edge at u_0 , or u_1 , or both). Next, we find the edge (u_2, u_3) paired with (u_1, u_2) at u_2 , etc. Continuing this process, we either find a path or a cycle. More precisely, if we find a path, it will be $u_{-k_1}, u_{-k_1+1}, \dots, u_{-1}, u_0, u_1, u_2, u_3, \dots, u_{k_2}$ such that (u_{i-1}, u_i) is paired with (u_i, u_{i+1}) at u_i for every $i \in \{-k_1 + 1, \dots, k_2 - 1\}$, and there is no edge paired with (u_{-k_1}, u_{-k_1+1}) at u_{-k_1} , and there is no edge paired with (u_{k_2-1}, u_{k_2})

at u_{k_2} . If we find a cycle, it will be $u_0, u_1, u_2, u_3 \dots, u_k$ such that (u_{i-1}, u_i) is paired with (u_i, u_{i+1}) at u_i for every $i \in \{1, \dots, k-1\}$, and (u_0, u_1) is paired with (u_{k-1}, u_k) at $u_0 = u_k$. Notice that this path or cycle might repeat vertices² We continue finding paths or cycles that use edges that have not yet been used in other paths or cycles.

Definition of the canonical flow. Let $I, F \in \Omega$ be two states and let $\psi \in \Psi_{I,F}$ be a pairing of $I \oplus F$. We first define a **canonical order of cycles/paths** in the cycle/path decomposition of $I \oplus F$ with respect to ψ . Suppose the edges of G are numbered $1, 2, \dots, |E_G|$, and this numbering is independent of I and F and the pairing ψ . Then we assign a number to every cycle/path identical to its lowest numbered edge, and we order the cycles/paths according to these numbers.

Then, we process each cycle/path following the canonical order, i. e., all the cycles and paths ordered before the current cycle/path have been already processed, and all the cycles and paths ordered after the current cycle/path have not been processed yet. In this context the word “process” means to change the I -edges into the F -edges. Thus, we start at I and after processing every cycle and path, we end up with F .

Processing paths. Let u_1, u_2, \dots, u_k be a path from the decomposition, and, without loss of generality suppose that (u_1, u_2) is a lower-numbered edge than (u_{k-1}, u_k) . We will gradually remove I -edges and add F -edges by dealing with edges (u_i, u_{i+1}) roughly in the order of increasing i . We occasionally need to skip over some i 's and come back to them to guarantee that all vertices are always covered (this problem can also be overcome by designing an additional move of the Markov chain corresponding to swapping one edge for another; however, since our goal is to analyze the Glauber dynamics Markov chain, we opted for this small complication in the canonical flow definition). The exact algorithm is described below (and Lemma 1 will show that at any given time all vertices are covered).

1. Let $i = 1$.
2. While $i < k$ do
3. If (u_i, u_{i+1}) is in F , then
4. Use the “add” transition of the Markov chain to add the edge (u_i, u_{i+1}) .
5. Else (if the edge is in I)
6. If $i + 1 = k$ or if (u_{i+1}, u_{i+2}) is in I , then
7. Use the “remove” transition of the chain to remove the edge (u_i, u_{i+1}) .
8. Else
9. Use the “add” transition to add the edge (u_{i+1}, u_{i+2}) .
10. Then use the “remove” transition to remove the edge (u_i, u_{i+1}) .
11. Increment i by 1.
12. Increment i by 1.

² Technically speaking, it would be more appropriate to use terms “walk” and “tour.” We use “path” and “cycle” to be consistent with earlier papers.

Processing cycles. Cycles are processed similarly. Let $u_1, u_2, \dots, u_{k-1}, u_k = u_1$ be a cycle from the decomposition and let (u_1, u_2) be its lowest numbered F -edge, or let it be the lowest numbered edge if the cycle contains only I -edges. Let the number of (u_2, u_3) be smaller than the number of (u_{k-1}, u_k) (otherwise, we “flip” the cycle, i. e., follow the cycle in the opposite direction, to satisfy this property). We process the cycle as follows.

1. Let $i = 1$ and let $u_{k+1} = u_2$.
2. While $i < k$ do
3. If (u_i, u_{i+1}) is in F , then
4. Use the “add” transition of the Markov chain to add the edge (u_i, u_{i+1}) .
5. Else (if the edge is in I)
6. If (u_{i+1}, u_{i+2}) is in I , then
7. Use the “remove” transition of the chain to remove the edge (u_i, u_{i+1}) .
8. Else
9. If $i + 1 < k$, then
10. Use the “add” transition to add the edge (u_{i+1}, u_{i+2}) .
11. Then use the “remove” transition to remove the edge (u_i, u_{i+1}) .
12. Increment i by 1.
13. Increment i by 1.

The following lemma proves that this process always goes only through valid states of the Markov chain.

Lemma 1. *Let $I, F \in \Omega$ and let $\psi \in \Psi_{I,F}$. Moreover, let $I = S_1, S_2, \dots, S_{\ell-1}, S_\ell = F$ be the sequence of subgraphs of G encountered during the above described process of changing I into F . Then, $S_j \in \Omega$ for every $j \in \{1, 2, \dots, \ell\}$.*

Proof. The proof goes by induction on j . Clearly, the claim holds for S_1 since we assume that $I \in \Omega$. For the inductive case, let us assume that $S_j \in \Omega$ for some $j \geq 1$. We want to show that $S_{j+1} \in \Omega$, too. The described process obtains S_{j+1} either by adding or by removing an edge. If an edge is added to S_j and all vertices were covered in S_j , all vertices must be covered in S_{j+1} as well. Thus, $S_{j+1} \in \Omega$. It remains to deal with the case when an edge is removed from S_j .

Let (u_i, u_{i+1}) be the removed edge and let $C = u_1, u_2, \dots, u_k$ be the cycle or path that contains it. Since (u_i, u_{i+1}) was removed, it belongs to I . To prove that $S_{i+1} \in \Omega$, we need to show that after removing the edge, both u_i and u_{i+1} are still covered by another edge in S_{j+1} (and S_j , since S_j and S_{j+1} differ only in the edge (u_i, u_{i+1})).

First we show that u_{i+1} is always covered in S_{j+1} . We have several cases to consider:

- Case 1: $i < k - 1$ and $(u_{i+1}, u_{i+2}) \in I$. Then, after removing (u_i, u_{i+1}) , the vertex u_{i+1} is still covered by the edge (u_{i+1}, u_{i+2}) which must be in both S_j and S_{j+1} because the edges are processed in order (with the minor exception of steps 9-11, but this exception does not apply in this case).

- Case 2: $i < k - 1$ and $(u_{i+1}, u_{i+2}) \in F$. Then, according to the step 9 in the path processing algorithm, or step 10 in the cycle processing algorithm, S_{j-1} and S_j differ in the addition of edge (u_{i+1}, u_{i+2}) , and therefore, the vertex u_{i+1} is covered by that edge in both S_j and S_{j+1} .
- Case 3: $i = k - 1$ and C is a path. Since C is a path, the edge $(u_i, u_{i+1}) = (u_{k-1}, u_k)$ is not paired with any other edge at u_k . However, vertex u_k must be covered by at least one edge in I and at least one edge in F . If there is an edge $e \in I \cap F$ incident to u_k , then, after removing the edge (u_{k-1}, u_k) , vertex u_k will still be covered by e . Otherwise, there is a pair of edges $e_1 \in I$ and $e_2 \in F$ that are paired together at u_k . If e_1, e_2 belong to a different cycle/path than C , then either that cycle/path was processed before C , in which case u_k is covered by e_2 in S_{j+1} , or it will be processed after C , in which case u_k is covered by e_1 in S_{j+1} .

Finally, let us consider the case when e_1, e_2 belong to C . Since (u_{k-1}, u_k) is the last edge in C and $(u_{k-1}, u_k) \in F$, the edges e_1, e_2 must have been processed earlier. Therefore, $e_1 \notin S_j$ and $e_2 \in S_j$, and, therefore, u_k is covered by e_2 in S_j (and S_{j+1}).

- Case 4: $i = k - 1$ and C is a cycle. If $(u_1, u_2) \in F$, then $(u_1, u_2) \in S_j$ since this edge was processed as the first edge of C . Therefore, $u_{i+1} = u_k = u_1$ is covered in S_j (and S_{j+1}). Let us suppose that $(u_1, u_2) \in I$, and, thus, $(u_1, u_2) \notin S_j$. We know that (u_{k-1}, u_k) is paired with (u_1, u_2) at $u_k = u_1$ (forming the cycle C) and both (u_{k-1}, u_k) and (u_1, u_2) are in I . However, the vertex u_k is covered by an edge e in F . Either $e \in I \cap F$, or e must be paired at u_k with an edge $e_2 \in I$ (and e_2 is different from (u_{k-1}, u_k) and (u_1, u_2)) – otherwise we would not have paired together two edges from I . If $e \in I \cap F$, then u_k is always covered by this edge. Otherwise, if e lies on a cycle/path preceding C or on C , then e is in S_j and therefore u_k is covered in S_{j+1} . Alternatively, if e lies on a cycle/path processed after C , then u_k is covered by e_2 in S_{j+1} .

Similar arguments show that u_i is covered in S_{j+1} , we discuss them only briefly:

- Case 1: $i > 1$ and (u_{i-1}, u_i) is in F . Then, u_i is covered by (u_{i-1}, u_i) in S_{j+1} .
- Case 2: $i > 1$ and (u_i, u_{i+1}) is in I , or $i = 1$ and C is a path. In either case, there exists an edge $e \in F$ incident to u_i and it will be either e or the edge e is paired with that will cover u_i in S_{j+1} .
- Case 3: $i = 1$ and C is a cycle. We claim that $(u_{k-1}, u_k) \in I$. Then, the vertex $u_1 = u_k$ is covered by (u_{k-1}, u_k) in S_{j+1} . It remains to show that $(u_{k-1}, u_k) \in I$. By contradiction, suppose this edge is in F . But the edge (u_1, u_2) was chosen as the smallest numbered edge in F , or if there is no edge in F , then the smallest numbered edge in I . Thus, if (u_1, u_2) is in I , then (u_{k-1}, u_k) must be in I as well. \square

Splitting the flow. Finally, we are ready to define the canonical flow between I and F . We consider all pairings from $\Psi_{I,F}$. For a pairing $\psi \in \Psi_{I,F}$, we defined a cycle/path decomposition and we described an algorithm for changing I into F by following the canonical order of cycles/paths. We call the sequence of

intermediate subgraphs of G during the modification phase the **canonical path from I to F with respect to ψ** . Then, the **canonical flow from I to F** splits equally between all canonical paths from I to F , i. e., for each path p that is a canonical path from I to F , the flow value $g(p) = 1/|\Psi_{I,F}|$.

3.3 Analyzing the Congestion

In this section we bound $\rho(T)$ for every transition T . In words, we are estimating the number of canonical paths (weighted by $1/|\Psi_{I,F}|$) using the transition T . We will use the standard **encoding** technique to bound this number.

Encoding. Let $T = (M, M')$ be a transition from a state M to a state M' . Let I and F be a pair of states and $\psi \in \Psi_{I,F}$ be a pairing of $I \oplus F$ such that the canonical path from I to F associated with ψ uses the transition T . We will encode this canonical path by an “almost state” $E = I \oplus F \oplus M$ and a pairing ψ' of edges in $M \oplus E$ that we will specify later. (Technically, we defined pairings only between a pair of states from Ω , but the same definition easily applies to any subgraphs of G .) The following lemmas show that E has all except possibly up to four vertices covered by edges.

Lemma 2. *Let $I, F \in \Omega$ and $\psi \in \Psi_{I,F}$, let C_1, C_2, \dots, C_ℓ be the components of the cycle/path decomposition of $I \oplus F$ with respect to ψ , and suppose the components are processed in the given order. Let X_i be the state right after processing the i -th component, and let $Y_i = I \oplus F \oplus X_i$. Moreover, let $c_I(v)$, $c_F(v)$, $c_{X_i}(v)$, and $c_{Y_i}(v)$ be the degrees of a vertex v in $I \setminus F$, $F \setminus I$, $X_i \setminus Y_i$, and $Y_i \setminus X_i$, respectively. Then,*

- $\{c_I(v), c_F(v)\} = \{c_{X_i}(v), c_{Y_i}(v)\}$ for all vertices $v \in V$, and
- $Y_i \in \Omega$ for every i .

Proof. The proof proceeds by induction on i . The base case, $i = 0$, follows from the fact that $X_0 = I$ and $Y_0 = F$. For the inductive case, we assume the statement holds for $i \geq 0$ and we will prove it for $i + 1$.

It will be useful to notice that $0 \leq c_I(v), c_F(v) \leq 2$. This follows from the fact that the graph has maximum degree 3: then, if I and F share an edge adjacent to v , then the degree of v in $I \setminus F$ (and in $F \setminus I$) is at most 2. On the other hand, if there is no common edge adjacent to v , then the degree of v in I must be smaller than 3 because there must exist an edge covering v in F . Thus, $c_I(v) \leq 2$, and, similarly, $c_F(v) \leq 2$.

By the definition of the canonical path with respect to ψ , the state X_i differs from the state X_{i+1} only in the edges that belong to the component C_i . Namely, X_i contains the I -edges of C_i but not the F -edges, and X_{i+1} contains only the F -edges but not the I -edges of C_i . We will show that for every vertex v , we have $\{c_{X_i}(v), c_{Y_i}(v)\} = \{c_{X_{i+1}}(v), c_{Y_{i+1}}(v)\}$, and that v is covered in both X_{i+1} and Y_{i+1} , thus proving the lemma.

Clearly, $c_{X_i}(v) = c_{X_{i+1}}(v)$ and $c_{Y_i}(v) = c_{Y_{i+1}}(v)$ for every vertex v that is not on C_i , thus the claim holds for such vertices. Suppose that v is on C_i . There are several cases:

- Case 1: v is adjacent to two edges on C_i that are paired at v , one edge is in I , and one is in F . Then, we know that $1 \leq c_I, c_F \leq 2$. Thus, if $c_I = c_F$, we have $c_{X_i} = c_{Y_i} = c_{X_{i+1}} = c_{Y_{i+1}}$ and the claim follows. Otherwise, without loss of generality assume that $c_{X_i} = 1$ and $c_{Y_i} = 2$. Then, $c_{X_{i+1}} = 2$ and $c_{Y_{i+1}} = 1$ and the claim holds.
- Case 2: v is adjacent to two edges on C_i that are paired at v , without loss of generality assume that both are from I . Then, we know that $c_{X_i} = 2$ and $c_{Y_i} = 0$ (otherwise the two I -edges would not have been paired). Thus, $c_{X_{i+1}} = 0$ and $c_{Y_{i+1}} = 2$, and $\{c_{X_i}(v), c_{Y_i}(v)\} = \{c_{X_{i+1}}(v), c_{Y_{i+1}}(v)\}$. It remains to show that v is covered in X_{i+1} . Since v was, by the inductive hypothesis, covered in both X_i and Y_i , the only possibility for this to happen is that X_i and Y_i share an edge adjacent to v . This edge is then in X_{i+1} as well and it covers v .
- Case 3: v is adjacent to a single edge on C_i , without loss of generality suppose it is an edge from I . Then, either $c_{X_i} = 1$ and $c_{Y_i} = 0$, or $c_{X_i} = 2$ and $c_{Y_i} = 1$. In the second case, $c_{X_{i+1}} = 1$ and $c_{Y_{i+1}} = 2$ and the claim holds. For the first case we have $c_{X_{i+1}} = 0$ and $c_{Y_{i+1}} = 1$ and we need to argue that v is covered in X_{i+1} . The argument is analogous to Case 2: there exists an edge shared by both X_i and Y_i (and therefore X_{i+1} and Y_{i+1}) that covers v . \square

Lemma 3. *Let $I, F \in \Omega$ and $\psi \in \Psi_{I,F}$. Let $T = (M, M')$ be a transition used on the ψ -defined canonical path from I to F and let $C = u_1, u_2, \dots, u_k$ be the component of the cycle/path decomposition of $I \oplus F$ that uses T . Let (u_t, u_{t+1}) be the edge in $M \oplus M'$. Then, for every vertex $v \in V \setminus \{u_1, u_{t-1}, u_t, u_{t+1}\}$, we have that $\{c_I(v), c_F(v)\} = \{c_M(v), c_E(v)\}$ and v is covered in $E = I \oplus F \oplus M$.*

The proof of Lemma 3 follows from Lemma 2. We defer the details to the journal version of this paper.

Let us denote the set of “almost states” as Ω' , i. e., Ω' is the set of subgraphs of G where at most four vertices are not covered by an edge of the subgraph. The following lemma states a polynomial relationship between the sizes of Ω and Ω' – this relationship will be important for establishing a polynomial bound on the congestion.

Lemma 4. $|\Omega'| \leq (m+1)^4 |\Omega|$, where m is the number of edges of G .

Proof. Let Ω'_k be the subset of Ω' containing graphs with exactly k vertices uncovered. We will give an injection h from Ω'_k to $\Omega \times E_G^k$. Let $H \in \Omega'_k$ and let v_1, v_2, \dots, v_k be its uncovered vertices. For every $i \in \{1, 2, \dots, k\}$, take the smallest numbered edge adjacent to v_i , let it be (v_i, v'_i) . Add all these edges to H , creating a graph H' . Thus, $h(H) = (H', (v_1, v'_1), (v_2, v'_2), \dots, (v_k, v'_k))$. Clearly, this map is injective since $H' \setminus \{(v_i, v'_i) \mid i \in \{1, 2, \dots, k\}\}$ results in H . Therefore, $|\Omega'_k| \leq m^k |\Omega|$. Summing over $k \in \{0, 1, 2, 3, 4\}$, we get that $|\Omega'| \leq (m^4 + m^3 + m^2 + m + 1) |\Omega| \leq (m+1)^4 |\Omega|$. \square

Specifying the pairing from $\Psi_{E,M}$. By the definition of E we have that $I \oplus F = M \oplus E$. Thus, we need to pair the same set of edges in $M \oplus E$ as in

$I \oplus F$. Recall that the pairing $\psi \in \Psi_{I,F}$ is given. We construct the corresponding pairing ψ' by pairing up the same edges as those paired up in ψ . However, now we need to argue that this pairing is “legal,” i. e., that for every vertex, there is at most one unpaired edge (this property follows directly from the fact that it is satisfied for ψ), and if there is a pairing of two edges from E , then no two edges from M are paired together at the same vertex, and vice versa. We allow exceptions at at most four vertices. This statement is summarized in the following lemma and its proof follows from the arguments made in proofs of Lemmas 2 and 3.

Lemma 5. *Let ψ' be constructed as described above and let $C = u_1, u_2, \dots, u_k$ be the component of the cycle/path decomposition of $I \oplus F$ that uses T . Let (u_s, u_{s+1}) be the edge in $M \oplus M'$ and let $v \in V \setminus \{u_1, u_{s-1}, u_s, u_{s+1}\}$. Finally, recall that $c_M(v)$ is the degree of v in $M \setminus E$ and $c_E(v)$ is its degree in $E \setminus M$. Then, the pairing ψ' pairs $\min\{c_M(v), c_E(v)\}$ edges from M with edges from E at v , plus it pairs the remaining edges so that at most one is left unpaired.*

Notice that if we know $E, T = (M, M')$, and the pairing ψ' , we can reconstruct I, F , and ψ . Since ψ' defines a cycle/path decomposition on $E \oplus M = I \oplus F$ and $\psi' = \psi$, we know which components of the decomposition have been already processed, which have not been processed, and which one is partially processed. Therefore, we can uniquely determine which edges of $E \oplus M$ belong to I and which belong to F .

Bounding the congestion. Suppose that $T = (M, M')$ where M and M' are almost identical valid edge covers with the only difference that M does not contain (u, v) and M' does contain this edge, or vice versa.

Notice that $\pi(M) = \pi(I) = \pi(F) = 1/|\Omega|$, and $p_{M,M'} = 1/(2m)$. Also note that the length of any above described canonical path is upper-bounded by m since every edge gets changed (added or removed) at most once. Finally, as we argued above, every path through T can be encoded by a pair of $E \in \Omega'$ and $\psi' \in \Psi_{M,E}$ where the definition of the set of all pairings between two states has been naturally extended to include any two subgraphs of G . Therefore,

$$\begin{aligned} \rho_g(T) &= \frac{1}{\pi(M)p_{M,M'}} \sum_{p \ni T} \pi(I_p)\pi(F_p)g(p)\ell(g) \leq \frac{2m^2}{|\Omega|} \sum_{p \ni T} \frac{1}{|\Psi_{I_p, F_p}|} \\ &\leq \frac{2m^2}{|\Omega|} \sum_{E \in \Omega'} \sum_{\psi' \in \Psi_{M,E}} \frac{1}{|\Psi_{I_p, F_p}|} \leq \frac{2m^2}{|\Omega|} \sum_{E \in \Omega'} 3^4 \leq 162m^2(m+1)^4 = O(m^6), \end{aligned}$$

where the second to last inequality follows from Lemma 5. Namely, we know that the number of pairings of $M \oplus E$ at any given vertex (with a possible exception of up to four vertices) is the same as the number of pairings of $I \oplus F$ at the same vertex. Thus, for every such vertex v , we have $|\Psi_{I_p, F_p}(v)| = |\Psi_{M,E}(v)|$. At the four special vertices, the ratio $|\Psi_{M,E}(v)|/|\Psi_{I_p, F_p}(v)|$ could be arbitrary; however, the largest number of pairings at a vertex v of degree 3 is 3 (if all edges adjacent to v are in M and no edge is in E), the smallest number of pairings is

1. Thus, $|\Psi_{M,E}(v)|/|\Psi_{I_p,F_p}(v)| \leq 3$ and therefore $|\Psi_{M,E}| = \prod_{v \in V_G} |\Psi_{M,E}(v)| \leq 3^4 \prod_{v \in V_G} |\Psi_{I_p,F_p}(v)| = 3^4 |\Psi_{I_p,F_p}|$.

Then, the mixing time is bounded by

$$\tau_x(\delta) \leq \rho_g(\log \pi(x)^{-1} + \log \delta^{-1}) = O(m^6(m + \log \delta^{-1})),$$

where the last equality follows from the fact that the total number of edge covers is upper-bounded by the total number of subgraphs of G . Thus, $\pi(x) = 1/|\Omega|$ and $|\Omega| \leq 2^m$. Also notice that for constant-degree graphs $m = O(n)$. This finishes the proof of Theorem 1.

References

1. Cooper, C., Dyer, M., Greenhill, C.: Sampling regular graphs and a peer- to-peer network. *Combinatorics, Probability and Computing* 16, 557–593 (2007)
2. Diaconis, P., Saloff-Coste, L.: Comparison Theorems for Reversible Markov Chains. *The Annals of Applied Probability* 3(3), 696–730 (1993)
3. Diaconis, P., Stroock, D.: Geometric Bounds for Eigenvalues of Markov Chains. *The Annals of Applied Probability* 1(1), 36–61 (1991)
4. Jerrum, M., Sinclair, A.: Approximating the permanent. *SIAM Journal on Computing* 18, 1149–1178 (1989)
5. Jerrum, M., Sinclair, A.: Approximate counting, uniform generation and rapidly mixing Markov chains. *Information and Computation* 82, 93–133 (1989)
6. Jerrum, M., Sinclair, A., Vigoda, E.: A polynomial-time approximation algorithm for the permanent of a matrix with non-negative entries. *J. ACM* 51(4), 671–697 (2004)
7. Kannan, R., Tetali, P., Vempala, S.: Simple Markov-chain algorithms for generating bipartite graphs and tournaments. *Random Structures and Algorithms (RSA)* 14(4), 293–308 (1999)
8. Luby, M., Randall, D., Sinclair, A.: Markov chain algorithms for planar lattice structures. *SIAM Journal on Computing* 31, 167–192 (2001)
9. Randall, D., Tetali, P.: Analyzing Glauber dynamics by comparison of Markov chains. *Journal of Mathematical Physics* 41, 1598–1615 (2000)
10. Sinclair, A.: Improved Bounds for Mixing Rates of Marked Chains and Multicommodity Flow. *Combinatorics, Probability and Computing* 1, 351–370 (1992)
11. Valiant, L.G.: The complexity of computing the permanent. *Theoretical Computer Science* 8, 189–201 (1979)
12. Vigoda, E.: A note on the Glauber dynamics for sampling independent sets. *Electronic Journal of Combinatorics* 8(1) (2001); Research paper 8