Rapid Mixing for the Hardcore Glauber Dynamics and Other Markov Chains in Bounded-Treewidth Graphs

David Eppstein ✉
Department of Computer Science, University of California, Irvine, CA, USA

Daniel Frishberg ✉ 🔗
Department of Computer Science and Software Engineering, California Polytechnic State University, San Luis Obispo, CA, USA

⎯ Abstract ⎯

We give a new rapid mixing result for a natural random walk on the independent sets of a graph G. We show that when G has bounded treewidth, this random walk – known as the *Glauber dynamics* for the *hardcore model* – mixes rapidly for all fixed values of the standard parameter $\lambda > 0$, giving a simple alternative to existing sampling algorithms for these structures. We also show rapid mixing for analogous Markov chains on dominating sets, b-edge covers, b-matchings, maximal independent sets, and maximal b-matchings. (For b-matchings, maximal independent sets, and maximal b-matchings we also require bounded degree.) Our results imply simpler alternatives to known algorithms for the sampling and approximate counting problems in these graphs. We prove our results by applying a divide-and-conquer framework we developed in a previous paper, as an alternative to the *projection-restriction* technique introduced by Jerrum, Son, Tetali, and Vigoda. We extend this prior framework to handle chains for which the application of that framework is not straightforward, strengthening existing results by Dyer, Goldberg, and Jerrum and by Heinrich for the Glauber dynamics on q-colorings of graphs of bounded treewidth and bounded degree.

1 Introduction

The *Glauber dynamics* on independent sets in a graph – motivated in part by modeling systems in statistical physics – is a Markov chain in which one starts at an arbitrary independent set, then repeatedly chooses a vertex at random and, with probability that depends on a fixed parameter $\lambda > 0$, either removes the vertex from the set (if it is in the set), or adds it to the set (if it is not in the set and has no neighbor in the set). This chain, which samples from the *hardcore model* on independent sets, has seen recent rapid mixing results under various conditions. In addition to independent sets, similar dynamics have been studied for a number of other structures – including, for example, q-colorings, matchings, and edge covers (more generally, b-matchings and b-edge covers).

1.1 Our contribution

We prove that the hardcore Glauber dynamics mixes rapidly on graphs of bounded treewidth for all fixed $\lambda > 0$, and that the Glauber dynamics on partial q-colorings (for all $\lambda > 0$) of a graph of bounded treewidth, and on q-colorings of a graph of bounded treewidth and degree, mix rapidly. Marc Heinrich proved the latter result, namely for q-colorings, in a 2020 preprint [10]. Heinrich's result applies to all graphs of bounded treewidth; however, for graphs of bounded treewidth and degree, whose degree is less than quadratic in their treewidth, we improve on Heinrich's upper bound – provided that q is fixed. We also prove that the analogous dynamics on the b-edge covers (when b is bounded) and the dominating sets of a graph of bounded treewidth mix rapidly for all $\lambda > 0$. In a similar vein, we prove that three additional chains – on b-matchings (when $\lambda > 0$), on maximal independent sets, and on maximal b-matchings – mix rapidly in graphs of bounded treewidth and degree.

To prove our results, we apply a framework we introduced in a companion paper [6] that uses the *multicommodity flow* technique (essentially the same as the *canonical paths* technique) for bounding mixing times. (We previously presented this framework in a preprint of the present paper [5].) The framework consists of a set of conditions (which we will define in Section 3.3) that guarantee rapid mixing; these conditions make progress towards unifying prior work on similar Glauber dynamics with prior work on probabilistic graphical models. In that paper [6], we also proved that the *flip walk* on the k-angulations of a convex n-point set mixes in time quasipolynomial in n for all fixed $k \geq 3$, although the special case $k = 3$ was known already to mix rapidly [15]. Thus our framework applies beyond graphical models and graph sampling problems.

1.2 Main results

Our main results are the following (see Section 2 for relevant definitions).

▶ **Theorem 1.** *The hardcore Glauber dynamics mixes in time $n^{O(t)}$ on graphs of treewidth t for all fixed $\lambda > 0$.*

▶ **Theorem 2.** *The (unbiased) Glauber dynamics on q-colorings (when $q \geq \Delta + 2$ is fixed) mixes in time $n^{O(t)}$ on graphs of treewidth t and bounded degree when q is fixed. The Glauber dynamics on partial q-colorings (when $q \geq \Delta + 2$ is fixed) mixes in time $n^{O(t)}$ on graphs of treewidth t for all fixed $\lambda > 0$.*

▶ **Theorem 3.** *The Glauber dynamics on b-edge covers mixes in time $n^{O(t^2)}$ on graphs of treewidth t, for all fixed $b \geq 1$ and fixed $\lambda > 0$. The Glauber dynamics on dominating sets mixes in time $n^{O(t)}$ on graphs of treewidth t for all fixed $\lambda > 0$. The Glauber dynamics on b-matchings mixes in time $n^{O(t)}$ on graphs of treewidth t and bounded degree Δ for all fixed $\lambda > 0$ and fixed $b \geq 1$.*

▶ **Theorem 4.** *There exist Markov chains on maximal independent sets and maximal b-matchings, whose stationary distributions are uniform, that mix in time $n^{O(t)}$ on graphs of treewidth t and bounded degree.*

1.3 The framework: recursive flow construction

A *multicommodity flow* in an undirected graph $G = (V, E)$ with n vertices is a set of n^2 flows, one flow for each ordered pair of vertices (s, t), where each flow sends one unit of a commodity from s to t. More precisely, take each (undirected) edge in E and make two directed copies,

one in each direction; let $E^+ = \bigcup_{\{u,v\} \in E} \{(u,v), (v,u)\}$ denote the set of all these directed copies. A multicommodity flow is a collection of functions $f_{st} : E^+ \to \mathbb{R}_{\geq 0}$ such that each f_{st} is a valid flow function, with s (respectively t) having net out flow (respectively in flow) equal to one, and all other vertices having zero net flow. If a multicommodity flow exists in G with small *congestion* – i.e. one in which no edge carries too much flow – then the natural Markov chain whose states are the vertices of G mixes rapidly.

The chains we analyze are natural random walks on a *Glauber graph* $\mathcal{M}(G)^1$ – the graph whose vertices are the structures over which the random walk is performed, and whose edges are the pairs of these structures with symmetric distance equal to one. For example, in the case of independent sets, $\mathcal{M}(G)$ has as its vertex set the collection of all independent sets in G, and as its edge set the collection of all (unordered) pairs of independent sets S, S' in G such that $S = S' \cup \{v\}$ for some $v \in V(G)$. Thus each of these random walks is performed on a graph that may be exponentially large with respect to the size of the input graph. In our previous work [6], we showed that when all of a certain set of conditions hold, we can construct a multicommodity flow in $\mathcal{M}(G)$ with congestion polynomial in $n = |V(G)|$, implying that the walk on $\mathcal{M}(G)$ mixes rapidly. The conditions specify that $\mathcal{M}(G)$ can be partitioned into a small number of induced subgraphs, all of which are approximately the same size, with large numbers of edges between pairs of the subgraphs. The conditions require that each of these induced subgraphs can be decomposed into smaller Glauber graphs that are similar in structure to $\mathcal{M}(G)$. This self similarity allows for the recursive construction of a multicommodity flow, by assembling flows on smaller Glauber graphs together into a flow in $\mathcal{M}(G)$ with small congestion.

1.4 Projection-restriction and prior work on the hardcore model

Prior work on rapid mixing of Markov chains on subset systems includes the special case of matroid polytopes. For this case, recent results [2, 1] have partly solved a 30-year-old conjecture of Mihail and Vazirani [16]. Other prior work uses multicommodity flows (and the essentially equivalent *canonical paths* technique) to obtain polynomial mixing upper bounds on structures of exponential size, including matchings and 0/1 knapsack solutions [17, 9]. Madras and Randall [13] used a decomposition of the hardcore model state space to prove rapid mixing under different conditions. We also decompose the state space, but our approach is different and is more similar to Heinrich's [10] application of the *projection-restriction* technique introduced by Jerrum, Son, Tetali, and Vigoda [11]. This technique involves partitioning the state space of a chain into a collection of sub-state spaces, each of which internally has a good *spectral gap* – a property that implies rapid mixing – and all of which are well connected to one another. Heinrich used the vertex separation properties of bounded-treewidth graphs to obtain an inductive argument: the resulting sub-spaces are themselves *Cartesian products* of chains on smaller graphs, and thus mix rapidly. (See Lemma 16.) We partition the state space recursively using the same vertex separation properties, and indeed for the chains on b-matchings and q-colorings in bounded-treewidth, bounded-degree graphs, combining these properties straightforwardly with the existing spectral projection-restriction machinery of [11] suffices for rapid mixing. The main contribution in this paper is to extend the framework to chains for which this application is not straightforward. That is, we give general conditions for constructing a multicommodity flow in the projection chain with small

---

[1] The chains on maximal independent sets and maximal b-matchings are not strictly Glauber dynamics, but we will use the same term for the graph, redefining the edge set as pairs connected by the moves we define in the appendix of the full version.

congestion, giving a good spectral gap in the projection chain. One may then apply induction using the spectral machinery of [11] to obtain rapid mixing in the overall chain; alternatively, one can substitute the flow-based machinery from our companion paper [6] for the spectral technique.

In the case of independent sets, Jerrum, Son, Tetali, and Vigoda [11] applied their technique to a special case of the hardcore model, namely regular trees. However, it was not clear how to generalize this application to bounded-treewidth graphs – since showing the spectral gap of the projection chain is sufficiently large is not straightforward. Martinelli, Sinclair, and Weitz [14] showed that the Glauber dynamics on the hardcore model mixes in $O(n \log n)$ time on the complete $\Delta - 1$-ary tree with n nodes, but they did not address general trees. Berger, Kenyon, Mossel, and Peres [3] showed rapid mixing for q-colorings of regular trees with unbounded degree but also did not address general trees. Our first main technical contribution is to show rapid mixing for general bounded-treewidth graphs by introducing the *hierarchical* version of our framework, in which we construct a flow with small congestion in the projection chain; we show that this construction gives rapid mixing for dominating sets, *partial q-colorings*, and *b*-edge covers in bounded-treewidth graphs. We solve another problem: the technical theorem in [11] as stated requires each of the state spaces in the partition to be a Cartesian product of chains on smaller spaces. For four of our eight chains – those on dominating sets, *b*-edge covers, maximal *b*-matchings, and maximal independent sets – the sub-spaces obtained in the decomposition are not a disjoint union of Cartesian products but may each be a union of Cartesian products, or may be mutually intersecting. In some cases, the sub-spaces may even induce disconnected restriction chains. Our second main contribution is to resolve this problem, using the structure of the state spaces of Glauber dynamics as graphs. We discuss this in the appendix of the full version.

1.5   Paper organization

In Section 2, we give relevant background. In Section 3, we use the chain on independent sets to review the "non-hierarchical" version of our framework (the version we gave in our companion paper) – which works for this chain when treewidth and degree are bounded. In the appendix of the full version we apply it to q-colorings and to *b*-edge covers and *b*-matchings. To fully prove Theorem 1 and Theorem 3, we need to deal with unbounded-degree graphs – our first main technical contribution. In Section 4, we modify the framework to do so, proving Theorem 1 for $\lambda = 1$. We defer some details to the appdenix of the full version, where we also finish the proof of Theorem 2 for $\lambda = 1$. We prove the general case $\lambda > 0$ of Theorems 1 and 2 in the appendix of the full version. We finish the proofs of Theorems 3 and 4 in the appendix of the full version: applying the framework to the relevant chains requires a further refinement of the framework. In all of the above, we prove rapid mixing but defer derivation of specific upper bounds to the appendix of the full version.

2   Preliminaries

2.1   Glauber dynamics

▶ **Definition 5.** *The* hardcore Glauber dynamics *on a graph G is the following chain, defined with respect to a fixed real parameter $\lambda > 0$:*
1. *Let X_0 be an arbitrary independent set in G.*
2. *For $t \geq 0$, select a vertex $v \in V(G)$ uniformly at random.*
3. *If $v \notin X_t$ and $X_t \cup \{v\}$ is not a valid independent set, do nothing.*

4. *Otherwise:*

   Let $X_{t+1} = X_t \cup \{v\}$ with probability $\lambda/(\lambda + 1)$.
   Let $X_{t+1} = X_t \setminus \{v\}$ with probability $1/(\lambda + 1)$.

Graph-theoretically, the Glauber dynamics is defined as follows: let the *indepdendent set Glauber graph* $\mathcal{M}_{\mathrm{IS}}(G)$ denote the graph whose vertices are identified with the independent sets of a given graph G, and whose edges are the pairs of independent sets whose symmetric difference is one. The hardcore Glauber dynamics is a Markov chain, parameterized by $\lambda > 0$, with state space $\Omega = V(\mathcal{M}_{\mathrm{IS}}(G))$ and probability matrix P, where for $S, S' \in V(\mathcal{M}(G))$ with $S \neq S'$, $P(S, S') = \lambda/(\Delta_{\mathcal{M}}(\lambda + 1))$ when $|S' \setminus S| = 1$, and $P(S, S') = 1/(\Delta_{\mathcal{M}}(\lambda + 1))$ when $|S \setminus S'| = 1$. If $S = S'$, then $P(S, S') = 1 - \sum_{S'' \neq S} P(S, S'')$. (Here $\Delta_{\mathcal{M}}$ is the maximum degree of the Glauber graph – i.e. the maximum number of neighboring states that a state S can have.) The Glauber graph has vertex set Ω and adjacency matrix P – up to the addition of self loops and normalization by degree. (When $\lambda \neq 1$ this graph can still be augmented with suitable weights so that the walk on the graph is the Glauber dynamics.)

2.2 Mixing time

To generate, approximately uniformly at random, an object of a given class – such as an independent set in a given graph – one can conduct a random walk on a graph whose vertices are the objects of interest, and whose edges represent local moves between the objects (or states). It is known that under certain mild conditions satisfied by as all our chains (see the appendix of the full version), the walk converges to the uniform distribution in the limit. The rate of convergence is important: in the case of subset systems such as those we consider, the walk takes place over an exponentially large number of subsets defined over an underlying set of size n. If the convergence, or *mixing time*, of the walk is polynomial in n, then the random walk is said to be *rapidly mixing*. The mixing time is denoted $\tau(\varepsilon)$, where ε denotes the desired precision of convergence to the uniform distribution, and the value of τ at ε is the minimum number of steps in the random walk before convergence is guaranteed. Convergence is measured via the *total variation distance* [19] between the distribution over states induced by the walk at a given time step, and the uniform distribution. One can obtain non-uniform *stationary* distributions by weighting the graph – see the appendix of the full version. See Levin, Peres, and Wilmer [12] for a comprehensive treatment of rapid mixing.

A Markov chain, given a starting state $S \in \Omega$, induces a probability distribution π_t at each time step t. The Glauber dynamics is known, regardless of starting state, to converge in the limit to a *stationary* distribution $\pi^*(S) = \lambda^{|S|}/Z(\mathcal{M}(G))$, where the term $Z(\mathcal{M}(G))$ is simply a normalizing value. When λ is unspecified, assume $\lambda = 1$ (the uniform case). The *mixing time* is defined as follows:

Given an arbitrary $\varepsilon > 0$, the *mixing time*, $\tau(\varepsilon)$, of a Markov chain with state space Ω and stationary distribution π^* is the minimum time t such that, regardless of starting state, we always have $\frac{1}{2} \sum_{S \in \Omega} |\pi(S) - \pi^*(S)| < \varepsilon$. Suppose the chain belongs to a family of Markov chains, the size of whose state space is parameterized by some value n. Here, $|\Omega|$ may be exponential in n. If $\tau(\varepsilon)$ is bounded by a polynomial function in $\log(1/\varepsilon)$ and in n, the chain is said to be *rapidly mixing*. It is common to omit the parameter ε and assume $\varepsilon = 1/4$.

2.3 Treewidth and vertex separators

▶ **Definition 6** ([18]). *A* tree decomposition *of a graph $G = (V, E)$ is a collection of sets $\{X_i\}, i = 1, \ldots, k$, called* bags, *together with a tree T, whose nodes are identified with the bags $\{X_i\}$, such that all of the following hold:*

1. *Every vertex in V lies in some bag, i.e. $\bigcup_{i=1}^{k} X_i = V$.*
2. *For every $(u, v) \in E$, the vertices u and v belong to at least one bag X_i together, i.e. for some i, $u \in X_i$ and $v \in X_i$.*
3. *The collection of all bags containing any given vertex $v \in V$, i.e. $\{X_i \mid v \in X_i\}$ forms a (connected) subtree of T.*

▶ **Definition 7** ([18]). *The* width *of a tree decomposition is one less than the size of the largest bag in the decomposition. The* treewidth *of a graph G is the minimum t such that a tree decomposition of G exists with width t.*

Intuitively, treewidth measures how far away a graph is from being a tree. For example, trees have treewidth one; a graph consisting of a single cycle of size at least three has treewidth two. Treewidth is of interest in large part because many NP-hard problems become tractable on graphs of bounded treewidth. For a full definition of treewidth and a survey of this phenomenon, known as *fixed-parameter tractability*, see [4].

For our purposes, treewidth is of interest due to its relationship to *vertex separators*: a vertex set $X \subseteq V$ in a graph $G = (V, E)$ is called a *vertex separator* if the deletion of X from V leaves the induced subgraph on the remaining vertices disconnected. Say that X is a *balanced separator* if deleting X partitions V into mutually disconnected subsets $A \cup B = V \setminus X$ such that $|V|/3 \leq |A| \leq |B| \leq 2|V|/3$. A graph G is *recursively s-separable* [7] if either (i) $|V(G)| \leq 1$, or (ii) G has a balanced separator X with $|X| \leq s$ and, after deleting X, the resulting subsets A and B induce subgraphs of G that are each recursively s-separable.

The following is known and easy to prove [7]:

▶ **Lemma 8.** *Every graph with treewidth $t \geq 1$ is recursively s-separable for all $s \geq t + 1$.*

3 $\lambda = 1$: Bounded treewidth and degree

To build up to the proof of Theorem 1, we first show a weaker result: that the unifrom hardcore Glauber dynamics mixes rapidly in graphs of bounded treewidth and degree. Fully proving Theorem 1, even in the unbiased case, requires the non-hierarchical framework. The main technical lemma in this section, Lemma 17, comes from our companion paper. Our contribution in this paper is the application to independent sets in graphs of bounded treewidth and degree – which we strengthen to graphs of bounded treewidth in Section 4.
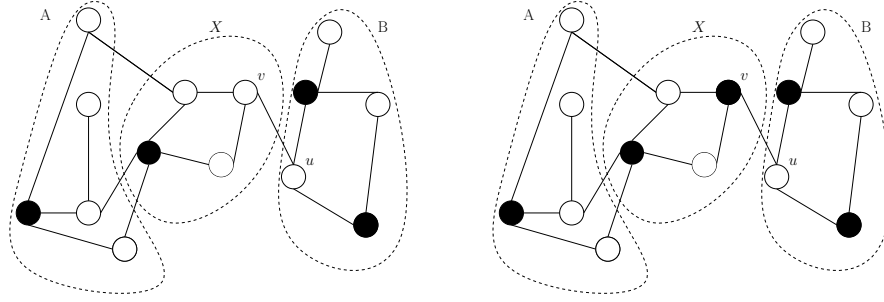
The following is necessary for the Glauber dynamics to sample correctly:

▶ **Lemma 9.** *The independent set Glauber graph is connected.*

Proof. Consider the empty independent set $\emptyset$. Every independent set $S \in V(\mathcal{M}_{\text{IS}}(G))$ has a path of length $|S|$ to $\emptyset$, formed by removing each vertex in S in arbitrary order. ◀

3.1   Partitioning the vertices of $\mathcal{M}_{\text{IS}}(G)$ into classes

The vertices of the Glauber graph $\mathcal{M}_{\text{IS}}(G)$ are subsets of the vertices of an underlying graph G. When G has bounded treewidth, we can choose a small separator X that partitions $V(G) \setminus X$ into two mutually disconnected vertex subsets, A and B, each of which has at most $2|V(G)|/3$ vertices. Consider the problem of sampling an independent set S from G. Given a separator X for G, partition the independent sets in G into equivalence classes as follows:

Figure 1 Two independent sets in a graph G: S (left) and S' (right), belonging to distinct classes. S and S' differ by a move, with the separator X inducing the classes to which the sets belong. S' results from adding v to S. $|S'| < |S|$, since S' excludes those independent sets that contain u.

▶ **Definition 10.** *Let $G = (V, E)$ be a graph. Let $\mathcal{M}_{\text{IS}}(G)$ be the independent set Glauber graph we have defined. Let $X \subseteq V$ be a vertex separator for G. Let $\mathcal{S}_{\text{IS}}(G)$ be the set of equivalence classes of $V(\mathcal{M}_{\text{IS}}(G))$ in which two independent sets S and S' are in the same class if $S \cap X = S' \cap X$. Let $T = S \cap X$, and call the corresponding class $\mathcal{C}_{\text{IS}}(T)$.*

(Technically X is also a parameter for $\mathcal{S}_{\text{IS}}(G)$ and $\mathcal{C}_{\text{IS}}(T)$, but we omit it for ease of notation.)

See Figure 1 for an example of a partitioning.

The *Cartesian product* of two graphs H and J is the graph whose vertex set is $V(H) \times V(J)$ and whose edges are the pairs $((h_1, j_1), (h_2, j_2))$ such that either $h_1 = h_2$ and $(j_1, j_2) \in E(J)$ or else $(h_1, h_2) \in E(H)$ and $j_1 = j_2$.
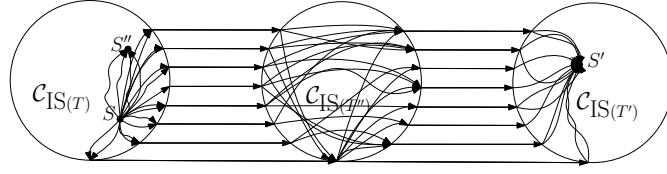
Let A and B be the mutually disconnected vertex subsets into which the removal of X partitions $V(G) \setminus X$. Given a fixed independent subset $T \subseteq X$, identify the independent sets in $\mathcal{C}_{\text{IS}}(T)$ with the pairs of the form (S_A, S_B), where S_A is an independent set in $A \setminus N_A(T)$, and S_B is an independent set in $B \setminus N_B(T)$, where $N_A(T)$ and $N_B(T)$ denote the union of the neighborhoods of vertices in T, in A and B respectively. That is, identify each independent set in $\mathcal{C}_{\text{IS}}(T)$ with a pair of an independent set in A that avoids neighbors of vertices in T, and a similar independent set in B. Consider the two Glauber graphs $\mathcal{M}_{\text{IS}}(A \setminus N_A(T))$ and $\mathcal{M}_{\text{IS}}(B \setminus N_B(T))$, whose vertices are respectively the independent sets in $G[A \setminus N_A(T)]$, and those in $G[B \setminus N_B(T)]$. If two independent sets $S = (S_A, S_B)$ and $S' = (S'_A, S'_B)$ belong to the same class, then a *move* (traversal of an edge in the Glauber graph) exists between S and S' in $\mathcal{M}_{\text{IS}}(G)$ precisely when a move exists between the restrictions of S and S' to either $\mathcal{M}_{\text{IS}}(A \setminus N_A(T))$ or $\mathcal{M}_{\text{IS}}(B \setminus N_B(T))$ (but not both). Therefore, each class induces, in $\mathcal{M}_{\text{IS}}(G)$, a subgraph that is isomorphic to a Cartesian product of two smaller Glauber graphs:

▶ **Lemma 11.** *Given a graph G and a vertex separator X that partitions $V(G)$ into subgraphs A and B, for every class $T \in \mathcal{S}_{\text{IS}}(G)$, $\mathcal{C}_{\text{IS}}(T) \cong \mathcal{M}_{\text{IS}}(A \setminus N_A(T)) \square \mathcal{M}_{\text{IS}}(B \setminus N_B(T))$.*

(Here by the symbol $\cong$ we denote isomorphism, and we identify the class $\mathcal{C}_{\text{IS}}(T)$ with the subgraph it induces in $\mathcal{M}_{\text{IS}}(G)$.)

3.2 Rapid mixing for the hardcore Glauber dynamics when G has bounded treewidth and degree

As described in Section 3.1, we use a small vertex separator X in G to give a decomposition of $\mathcal{M}_{\text{IS}}(G)$ into subgraphs, each of which has a Cartesian product structure – in which both factor graphs in the product are themselves Glauber graphs. Since Cartesian products preserve flow congestion upper bounds (see Lemma 16), this decomposition provides a crucial inductive structure. We analyze this structure in this section.

■ **Figure 2** A schematic view of three classes in the independent set Glauber graph $\mathcal{M}_{\mathrm{IS}}(G)$. The large circles denote classes under the partition described in Section 3.1. The curved arrows illustrate the construction of a flow in $\mathcal{M}_{\mathrm{IS}}(G)$ from an independent set $S \in \mathcal{C}_{\mathrm{IS}}(T)$ to another independent set $S'' \in \mathcal{C}_{\mathrm{IS}}(T)$ – and also to an independent set $S' \in \mathcal{C}_{\mathrm{IS}}(T')$. Here, $\mathcal{C}_{\mathrm{IS}}(T)$ and $\mathcal{C}_{\mathrm{IS}}(T'')$ are adjacent classes in $\mathcal{M}_{\mathrm{IS}}(G)$, connected by a large number of edges, and similarly $\mathcal{C}_{\mathrm{IS}}(T')$ and $\mathcal{C}_{\mathrm{IS}}(T'')$ are adjacent. In Section 3.2 we formalize this flow.

▶ **Lemma 12.** *Let G be a graph with bounded treewidth t and bounded degree Δ, let $\mathcal{M}_{\mathrm{IS}}(G)$ be as we have defined, and let $\mathcal{S}_{\mathrm{IS}}(G)$ be as in Definition 10 with respect to a small balanced separator X with $|X| \leq t + 1$. Then:*

1. *The number of classes in $\mathcal{S}_{\mathrm{IS}}(G)$ is $O(1)$.*
2. *For every pair of classes $\mathcal{C}_{\mathrm{IS}}(T), \mathcal{C}_{\mathrm{IS}}(T') \in \mathcal{S}_{\mathrm{IS}}(G)$, $|\mathcal{C}_{\mathrm{IS}}(T)| = \Theta(1)|\mathcal{C}_{\mathrm{IS}}(T')|$.*
3. *Let $\mathcal{C}_{\mathrm{IS}}(T), \mathcal{C}_{\mathrm{IS}}(T') \in \mathcal{S}_{\mathrm{IS}}(G)$ be two classes. No independent set $S \in \mathcal{C}_{\mathrm{IS}}(T)$ has more than one move to an independent set $S' \in \mathcal{C}_{\mathrm{IS}}(T')$.*
4. *Let $\mathcal{C}_{\mathrm{IS}}(T), \mathcal{C}_{\mathrm{IS}}(T') \in \mathcal{S}_{\mathrm{IS}}(G)$ be two classes. Suppose there exists at least one move between an independent set in $\mathcal{C}_{\mathrm{IS}}(T)$ and an independent set in $\mathcal{C}_{\mathrm{IS}}(T')$. Then there exist at least $\Omega(1)|\mathcal{C}_{\mathrm{IS}}(T)|$ moves between independent sets in $\mathcal{C}_{\mathrm{IS}}(T)$ and independent sets in $\mathcal{C}_{\mathrm{IS}}(T')$.*

Proof. Claim 1 follows from the fact that $|\mathcal{S}_{\mathrm{IS}}(G)| \leq 2^{|X|} \leq 2^{t+1} = O(1)$, where the first inequality is true because each class is identified with a subset of the vertices in X. The proofs of claims 2 through 4 are in the appendix of the full version. ◀

We will use Lemma 12 to prove the following, applying the framework from our previous paper, in Section 3.3:

▶ **Lemma 13.** *Given a graph G with bounded treewidth and degree, the natural random walk on the independent set Glauber graph $\mathcal{M}_{\mathrm{IS}}(G)$ has mixing time $\tau(\varepsilon) = O(n^c \log 1/\varepsilon)$, where $c = O(1)$.*

To prove Theorem 1, however, we need to get rid of the assumption that degree is bounded. We address this issue in Section 4.

3.3 Abstraction into framework conditions

The observations in Lemma 12 correspond to a set of conditions we gave in our previous work [6]. These conditions are, given a connected graph $\mathcal{M}(G)$, on some set of combinatorial structures over an underlying graph G with n vertices:

1. The vertices of $\mathcal{M}(G)$ can be partitioned into a set $\mathcal{S}$ of classes, where $|\mathcal{S}| = O(1)$.
2. The ratio of the sizes of any two classes in $\mathcal{S}$ is $\Theta(1)$.
3. Given two classes $\mathcal{C}(T), \mathcal{C}(T') \in \mathcal{S}$, no vertex in $\mathcal{C}(T)$ has more than $O(1)$ edges to vertices in $\mathcal{C}(T')$.
4. For every pair of classes that share at least one edge, the number of edges between the two classes is $\Theta(1)$ times the size of each of the two classes.
5. Each class in $\mathcal{S}$ is the Cartesian product of two graphs $\mathcal{M}(G_1)$ and $\mathcal{M}(G_2)$, each of which can be recursively partitioned in the same way as $\mathcal{M}(G)$.
6. The recursive partitioning mentioned in Condition 5 reaches the base case (graphs with one or zero vertices) in $O(\log n)$ steps.

Conditions 1 through 4 correspond respectively to Lemma 12; Condition 5 corresponds to Lemma 11. Condition 6 corresponds to the observation at the end of the proof sketch of Lemma 13.

We introduce some facts that we previously used to prove that these conditions suffice for rapid mixing, via *expansion*, then review a sketch of the proof; we will build on these techniques in Section 4 (our main contribution) and in the appendices.

The *edge expansion* (or simply the *expansion*) of a graph $G = (V, E)$ is the quantity $h(G) := \min_{S \subseteq V : |S| \leq |V|/2} \frac{|\{(u,v) \in E | u \in S, v \notin S\}|}{|S|}$, i.e. the minimum quotient of the number of edges in the cut by the number of vertices on the smaller side of the cut. The *vertex expansion* is the quantity $h(G) := \min_{S \subseteq V : |S| \leq |V|/2} \frac{|\{v \in V \setminus S | \exists u \in S, (u,v) \in E\}|}{|S|}$, i.e. the minimum quotient of the number of neighbors of a set S with $|S| \leq |V|/2$ that are not in S.

Mixing can be bounded from above via a lower bound on expansion [19] when the degree of a Glauber graph is small (linear in the case of our chains):

▶ **Lemma 14.** *Given a graph $\mathcal{M} = (\mathcal{V}, \mathcal{E})$, consider the Markov chain whose state space is $\mathcal{V}$ and whose transitions are of the form $P(x,x) = 1/2, P(x,y) = 0 \forall (x,y) \notin \mathcal{E}, P(x,y) = 1/(2\Delta) \forall (x,y) \in \mathcal{E}$, where Δ is the maximum degree of $\mathcal{M}$. The mixing time of this Markov chain is at most*

$$\tau(\varepsilon) = O\left(\frac{\Delta^2}{(h(\mathcal{M}))^2} \cdot \ln \frac{|\mathcal{V}|}{\varepsilon}\right).$$

Expansion, in turn, can be bounded from below via an upper bound on the *congestion* of a multicommodity flow. Given a multicommodity flow $f = \{f_{st} | s, t \in V \times V\}$ in a graph $G = (V, E)$, define the congestion of f as the quantity $\rho = \max_{(u,v) \in E^+} \sum_{s,t \in V \times V} f_{st}(u, v)$, i.e. the maximum amount of flow sent across an edge.

▶ **Lemma 15** ([19]). *For every graph $G = (V, E)$ and for every flow function f defined over G and having congestion ρ, $h(G) \geq 1/(2\rho)$.*

▶ **Lemma 16.** *Given graphs H and J, let G be the Cartesian product $H \square J$. Suppose multicommodity flows exist in H and J with congestion at most ρ_H and ρ_J respectively. Then there exists a multicommodity flow in G with congestion at most $\max\{\rho_H, \rho_J\}$.*

We proved Lemma 16 in [6], although an analogous result for expansion is known [8]. We also proved the following in [6]. We review the lemma and give a modified proof sketch here, in terms more intuitive for the chains we are analyzing in this paper. We will modify the technique in Section 4.

▶ **Lemma 17.** *Given a graph $\mathcal{M}(G)$ satisfying the conditions in Section 3.3, the expansion of $\mathcal{M}(G)$ is $\Omega(1/n^c)$, where $c = O(1)$.*

Proof Sketch. Partition $\mathcal{M}(G)$ into classes as in Definition 10. By Lemma 11, each class $\mathcal{C}(T) \in \mathcal{S}(G)$ is isomorphic to the Cartesian product $\mathcal{M}(A \setminus N_A(T)) \square \mathcal{M}(B \setminus N_B(T))$. We make an inductive argument, in which the inductive hypothesis assumes that for each such Cartesian product, the graphs $\mathcal{M}(A \setminus N_A(T))$ and $\mathcal{M}(B \setminus N_B(T))$ have multicommodity flows f_A and f_B with congestion $\rho_A \leq c^{\log |V(G)| - 1}, \rho_B \leq c^{\log |V(G)| - 1}$ respectively, for some constant c. By Lemma 16, $\mathcal{C}(T)$ then has a flow f_T with congestion $\rho_T \leq c^{\log |V(G)| - 1}$. The inductive step is to combine the f_T flows for all of the classes, giving a flow f in $\mathcal{M}(G)$ with small congestion. We need to route flow between every $S, S' \in V(\mathcal{M}(G))$. If S and S' belong to the same class $\mathcal{C}(T)$, simply use the same flow that S uses to send its unit to S' in f_T. If $S \in \mathcal{C}(T)$ and $S' \in \mathcal{C}(T') \neq \mathcal{C}(T)$ belong to different classes, we find a sequence

of intermediate classes through which to route flow from $\mathcal{C}(T)$ to $\mathcal{C}(T')$. See Figure 2. We specified in our companion paper [6] how to route the flow through each intermediate class so that the congestion across each edges between a pair of classes is at most $O(1)$, then made use of the existing flows within each class guaranteed by the inductive hypothesis to bound the resulting amount of flow within each class $\mathcal{C}(T)$. We showed that the latter is at most $O(1) \cdot \rho_T$, giving overall congestion $O(1)^l$, where l is the number of induction levels. Since X is a balanced separator we have $l = O(\log n)$; the lemma now follows from Lemma 14. ◀

The full proof of Lemma 17 is in our companion paper [6]. We will use the phrase "non-hierarchical framework" to describe this set of conditions – which apply to the chains we study when the underlying graph G has bounded treewidth and degree. Although Jerrum, Son, Tetali, and Vigoda [11] did not consider bounded-treewidth graphs generally, these conditions do allow their projection-restriction technique to be applied. In effect, Lemma 17 and its proof, which we gave in our previous work, characterize a sufficient set of conditions for applying Jerrum, Son, Tetali, and Vigoda's technique: specifically, one can, instead of routing flow internally through each intermediate class, simply treat the construction above as a flow in the projection graph, concluding that the projection graph has a good spectral gap – then apply [11]. The first main technical contribution of this paper is in Section 4, in which we give an alternative set of conditions – which we will call our "hierarchical framework" – that allows us to handle underlying graphs of unbounded degree (though treewidth still must be bounded), and to handle chains other than the hardcore model. This will allow us to complete the proofs of Theorems 1, 2, and 3.

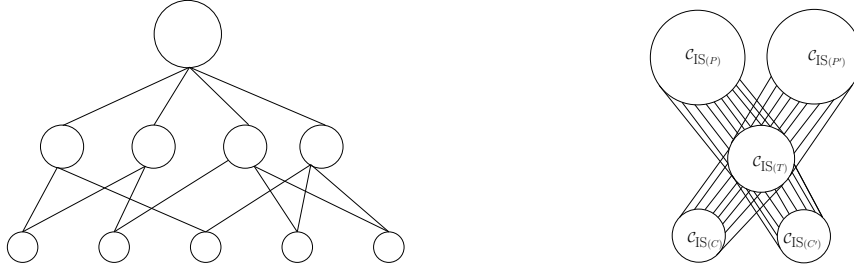## 4 $\lambda = 1$: Unbounded degree

4.1 Hierarchical framework

We now sketch "hierarchical" framework conditions that guarantee rapid mixing in the case of unbounded degree (when treewidth is bounded). Several of the chains we consider satisfy these conditions so long as the treewidth of the underlying graph is bounded. This is the first main technical contribution in this paper. In the original framework, we assumed that the classes were approximately the same size. Although all of the graphs to which we apply this hierarchical framework satisfy this condition in graphs with bounded treewidth and degree, this is not the case when the degree is unbounded. Fortunately, in the case of independent sets, partial q-colorings, dominating sets, and b-edge covers, we solve this problem with some modifications to the framework.

4.2 Independent sets

In the proof of Lemma 17, the assumption that the classes were approximately the same size allowed every class $\mathcal{C}_{\mathrm{IS}}(T)$ to route flow for all pairs of vertices without being too congested, because $\mathcal{C}_{\mathrm{IS}}(T)$ is sufficiently large. Once we discard this assumption, we need to specify explicitly the path through which a given $\mathcal{C}_{\mathrm{IS}}(T)$ routes flow to each $\mathcal{C}_{\mathrm{IS}}(T')$. We construct a flow in which for every such $\mathcal{C}_{\mathrm{IS}}(T), \mathcal{C}_{\mathrm{IS}}(T')$, every intermediate class $\mathcal{C}_{\mathrm{IS}}(T'')$ that handles flow between sets $S \in \mathcal{C}_{\mathrm{IS}}(T)$ and $S' \in \mathcal{C}_{\mathrm{IS}}(T')$ is larger than one of $\mathcal{C}_{\mathrm{IS}}(T)$ or $\mathcal{C}_{\mathrm{IS}}(T')$. We then bound the number of pairs of sets, relative to $|\mathcal{C}_{\mathrm{IS}}(T'')|$, for which $\mathcal{C}_{\mathrm{IS}}(T'')$ carries flow.

To accomplish this, we observe that for any $\mathcal{C}_{\mathrm{IS}}(T), \mathcal{C}_{\mathrm{IS}}(T')$ such that there exists *one* move between an independent set in $\mathcal{C}_{\mathrm{IS}}(T')$ and an independent set in $\mathcal{C}_{\mathrm{IS}}(T)$, either *every* independent set in $\mathcal{C}_{\mathrm{IS}}(T')$ has a move to some independent set in $\mathcal{C}_{\mathrm{IS}}(T)$, or vice versa. This move consists of dropping some vertex v from $T' \subseteq X$ to obtain T, i.e. $T = T' \setminus \{v\}$.

Figure 3 Left: a schematic representation of the classes in the independent set Glauber graph and edges between them when the degree is unbounded. Right: a class $\mathcal{C}_{\text{IS}}(T)$, with two parents, $\mathcal{C}_{\text{IS}}(P)$ and $\mathcal{C}_{\text{IS}}(P')$, and two children, $\mathcal{C}_{\text{IS}}(C)$ and $\mathcal{C}_{\text{IS}}(C')$. (Classes with larger cardinality are drawn larger.) The parallel edges depict the fact that a child class always has every one of its vertices adjacent to a vertex in a given parent class, and that the edges between any given pair of classes are vertex-disjoint.

We call $\mathcal{C}_{\text{IS}}(T)$ a *parent* of $\mathcal{C}_{\text{IS}}(T')$, and $\mathcal{C}_{\text{IS}}(T')$ a *child* of $\mathcal{C}_{\text{IS}}(T)$. See Figure 3. Since the set of edges connecting vertices in $\mathcal{C}_{\text{IS}}(T)$ with vertices in $\mathcal{C}_{\text{IS}}(T')$ forms a matching, this implies that $|\mathcal{C}_{\text{IS}}(T)| \geq |\mathcal{C}_{\text{IS}}(T')|$. In fact, whenever $T \subseteq T'$, we have $|\mathcal{C}_{\text{IS}}(T)| \geq |\mathcal{C}_{\text{IS}}(T')|$. We route flow between any pair of classes $\mathcal{C}_{\text{IS}}(T)$ and $\mathcal{C}_{\text{IS}}(T')$ along a path through a "least common ancestor". Since for every class $\mathcal{C}_{\text{IS}}(T'')$ on this path, either $|\mathcal{C}_{\text{IS}}(T'')| \geq |\mathcal{C}_{\text{IS}}(T)|$ or $|\mathcal{C}_{\text{IS}}(T'')| \geq |\mathcal{C}_{\text{IS}}(T')|$, we obtain a bound on congestion that we make precise in the appendix of the full version.

In the proof sketch of Lemma 17 (Section 3.2), for every pair $S \in \mathcal{C}(T), S' \in \mathcal{C}(T') \neq \mathcal{C}(T)$, we found a sequence of classes $\mathcal{C}(T) = \mathcal{C}(T_1), \mathcal{C}(T_2), \ldots, \mathcal{C}(T_{k-1}), \mathcal{C}(T_k) = \mathcal{C}(T')$ through which to route the $S - S'$ flow. As discussed in Section 4, when the degree is unbounded, the classes are no longer nearly the same size – so if this sequence is chosen carelessly, some $\mathcal{C}(T_i)$ may carry flow for too many $S - S'$ pairs. We therefore choose the sequences carefully: the parent-child relationships induce a partial order $\prec$ on the classes with a unique maximal element, where $\mathcal{C}(T) \succ \mathcal{C}(T')$ implies $|\mathcal{C}(T)| \geq |\mathcal{C}(T')|$. We choose our sequence so that for some i with $1 \leq i \leq k$, $|\mathcal{C}(T_1)| \leq |\mathcal{C}(T_2)| \leq \cdots \leq |\mathcal{C}(T_i)| \geq |\mathcal{C}(T_{i+1})| \geq \cdots |\mathcal{C}(T_{k-1})| \geq |\mathcal{C}(T_k)|$.

4.3   Hierarchical Framework Conditions

The conditions are as follows. Conditions 2 through 4 are new and concern the partial order described above; Condition 1 and Conditions 5 through 7 are as in Section 3.3.
1. The vertices of $\mathcal{M}(G)$ can be partitioned into a set $\mathcal{S}$ of classes, where $|\mathcal{S}| = O(1)$.
2. There exists a partial order $\prec$ on the classes in $\mathcal{S}$, such that whenever $\mathcal{C}(T), \mathcal{C}(T') \in \mathcal{S}$ and $\mathcal{C}(T) \succ \mathcal{C}(T')$, we have $|\mathcal{C}(T)| \geq |\mathcal{C}(T')|$.
3. The partial order $\prec$ has a unique maximal element.
4. Whenever an edge exists between vertices in $\mathcal{C}(T)$ and $\mathcal{C}(T')$ with $\mathcal{C}(T) \succ \mathcal{C}(T')$, the number of such edges is $|\mathcal{C}(T')|$.
5. For every pair of classes $\mathcal{C}(T)$ and $\mathcal{C}(T')$ that share an edge, the maximum degree, in $\mathcal{C}(T)$, of a vertex in $\mathcal{C}(T')$, is $O(1)$, and the maximum degree, in $\mathcal{C}(T')$, of a vertex in $\mathcal{C}(T)$, is $O(1)$.
6. Each class in $\mathcal{S}$ is the Cartesian product of two graphs $\mathcal{M}(G_1)$ and $\mathcal{M}(G_2)$, each of which can be recursively partitioned in the same way as $\mathcal{M}(G)$.
7. The recursive partitioning mentioned in Condition 6 reaches the base case (graphs with one or zero vertices) in $O(\log n)$ levels of recursion.

▶ **Lemma 18.** *Given a graph $\mathcal{M}(G)$ satisfying the conditions in Section 4.3, the expansion of $\mathcal{M}(G)$ is $\Omega(1/n^c)$, where $c = O(1)$.*

We defer the proof of Lemma 18 to the appendix of the full version.

───── **References** ─────

1 Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials II: High-dimensional walks and an FPRAS for counting bases of a matroid. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)*, New York, NY, USA, 2019. Association for Computing Machinery. `doi:10.1145/3313276.3316385`.

2 Nima Anari, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials, entropy, and a deterministic approximation algorithm for counting bases of matroids. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 35–46, 2018. `doi:10.1109/FOCS.2018.00013`.

3 Noam Berger, Claire Kenyon, Elchanan Mossel, and Yuval Peres. Glauber dynamics on trees and hyperbolic graphs. *Probability Theory and Related Fields*, 131:311–340, 2005. `doi:s00440-004-0369-4`.

4 Hans L. Bodlaender. Fixed-parameter tractability of treewidth and pathwidth. In *The Multivariate Algorithmic Revolution and Beyond: Essays Dedicated to Michael R. Fellows on the Occasion of His 60th Birthday*, volume 7370 of *Lecture Notes in Computer Science*, pages 196–227. Springer, 2012. `doi:10.1007/978-3-642-30891-8_12`.

5 David Eppstein and Daniel Frishberg. Rapid mixing of the hardcore Glauber dynamics and other Markov chains in bounded-treewidth graphs. *arXiv preprint*, 2021. `doi:10.48550/arXiv.2111.03898`.

6 David Eppstein and Daniel Frishberg. Improved Mixing for the Convex Polygon Triangulation Flip Walk. In *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, volume 261, pages 56:1–56:17, 2023. `doi:10.4230/LIPIcs.ICALP.2023.56`.

7 Jeff Erickson. Computational topology: Treewidth. Lecture Notes, 2009. URL: `http://jeffe.cs.illinois.edu/teaching/comptop/2009/notes/treewidth.pdf`.

8 F. Graham and P. Tetali. Isoperimetric inequalities for cartesian products of graphs. *Comb. Probab. Comput.*, 7:141–148, 1998.

9 Venkatesan Guruswami. Rapidly mixing markov chains: A comparison of techniques (a survey). *arXiv*, 2016. `arXiv:1603.01512`.

10 Marc Heinrich. Glauber dynamics for colourings of chordal graphs and graphs of bounded treewidth, 2020. `arXiv:2010.16158`.

11 Mark Jerrum, Jung-Bae Son, Prasad Tetali, and Eric Vigoda. Elementary bounds on Poincaré and log-Sobolev constants for decomposable Markov chains. *The Annals of Applied Probability*, 14(4):1741–1765, 2004. URL: `http://www.jstor.org/stable/4140446`.

12 David A Levin, Elizabeth Wilmer, and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Society, 2009. URL: `https://bookstore.ams.org/mbk-58`.

13 Neal Madras and Dana Randall. Markov chain decomposition for convergence rate analysis. *The Annals of Applied Probability*, 12(2):581–606, 2002. `doi:10.1214/aoap/1026915617`.

14 Fabio Martinelli, Alistair Sinclair, and Dror Weitz. Fast mixing for independent sets, colorings, and other models on trees. *Random Structures & Algorithms*, 31(2):134–172, 2007. `doi:10.1002/rsa.20132`.

15 Lisa McShine and P. Tetali. On the mixing time of the triangulation walk and other catalan structures. In *Randomization Methods in Algorithm Design*, 1997.

16 Milena Mihail and Umesh Vazirani. On the expansion of 0-1 polytopes. *Journal of Combinatorial Theory, Series B*, 1989.

17 Ben Morris and Alistair Sinclair. Random walks on truncated cubes and sampling 0-1 knapsack solutions. *SIAM journal on computing*, 34(1):195–226, 2004.

18 Neil Robertson and P.D Seymour. Graph minors. ii. algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, 1986. `doi:10.1016/0196-6774(86)90023-4`.

19 Alistair Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1(4):351–370, 1992. `doi:10.1017/S0963548300000390`.