
Markov Chains on Orbits of Permutation Groups

Mathias Niepert
Universität Mannheim
mniepert@gmail.com

Abstract

We present a novel approach to detecting and utilizing symmetries in probabilistic graphical models with two main contributions. First, we present a scalable approach to computing generating sets of permutation groups representing the symmetries of graphical models. Second, we introduce orbital Markov chains, a novel family of Markov chains leveraging model symmetries to reduce mixing times. We establish an insightful connection between model symmetries and rapid mixing of orbital Markov chains. Thus, we present the first lifted MCMC algorithm for probabilistic graphical models. Both analytical and empirical results demonstrate the effectiveness and efficiency of the approach.

1 Introduction

Numerous algorithms exploit model symmetries with the goal of reducing the complexity of the computational problems at hand. Examples are procedures for detecting symmetries of first-order theories [7] and propositional formulas [2] in order to avoid the exhaustive exploration of a partially symmetric search space. More recently, symmetry detection approaches have been applied to answer set programming [11] and (integer) linear programming [26, 27, 34, 30]. A considerable amount of attention to approaches utilizing model symmetries has been given by work on “lifted probabilistic inference [36, 9].” Lifted inference is mainly motivated by the large graphical models resulting from statistical relational formalism such as Markov logic networks [38]. The unifying theme of lifted probabilistic inference is that inference on the level of instantiated formulas is avoided and instead lifted to the first-order level. Notable approaches are lifted belief propagation [41, 22], bisimulation-based approximate infer-

ence algorithms [40], first-order knowledge compilation techniques [44, 16], and lifted importance sampling approaches [17]. With the exception of some results for restricted model classes [41, 44, 21], there is a somewhat superficial understanding of the underlying principles of graphical model symmetries and the probabilistic inference algorithms utilizing such symmetries. Moreover, since most of the existing approaches are designed for relational models, the applicability to other types of probabilistic graphical models is limited.

The presented work contributes to a deeper understanding of the interaction between model symmetries and the complexity of inference by establishing a link between the degree of symmetry in graphical models and polynomial approximability. We describe the construction of colored graphs whose automorphism groups are equivalent to those of the graphical models under consideration. We then introduce the main contribution, *orbital Markov chains*, the first general class of Markov chains for lifted inference. Orbital Markov chains combine the compact representation of symmetries with generating sets of permutation groups with highly efficient product replacement algorithms. The link between model symmetries and polynomial mixing times of orbital Markov chains is established via a path coupling argument that is constructed so as to make the coupled chains coalesce whenever their respective states are located in the same equivalence class of the state space. The coupling argument applied to orbital Markov chains opens up novel possibilities of analytically investigating classes of symmetries that lead to polynomial mixing times.

Complementing the analytical insights, we demonstrate empirically that orbital Markov chains converge faster to the true distribution than state of the art Markov chains on well-motivated and established sampling problems such as the problem of sampling independent sets from graphs. We also show that existing graph automorphism algorithms are applicable to compute symmetries of very large graphical models.

2 Background and Related Work

We begin by recalling some basic concepts of group theory and finite Markov chains both of which are crucial for understanding the presented work. In addition, we give a brief overview of related work utilizing symmetries for the design of algorithms for logical and probabilistic inference.

2.1 Group Theory

A symmetry of a discrete object is a structure-preserving bijection on its components. For instance, a symmetry of a graph is a graph automorphism. Symmetries are often represented with permutation groups. A group is an abstract algebraic structure (\mathfrak{G}, \circ) , where \mathfrak{G} is a set closed under a binary associative operation \circ such that there is a identity element and every element has a unique inverse. Often, we refer to the group \mathfrak{G} rather than to the structure (\mathfrak{G}, \circ) . We denote the size of a group \mathfrak{G} as $|\mathfrak{G}|$. A permutation group *acting on* a finite set Ω is a finite set of bijections $\mathfrak{g} : \Omega \rightarrow \Omega$ that form a group.

Let Ω be a finite set and let \mathfrak{G} be a permutation group acting on Ω . If $\alpha \in \Omega$ and $\mathfrak{g} \in \mathfrak{G}$ we write $\alpha^{\mathfrak{g}}$ to denote the image of α under \mathfrak{g} . A cycle $(\alpha_1 \alpha_2 \dots \alpha_n)$ represents the permutation that maps α_1 to α_2 , α_2 to α_3, \dots , and α_n to α_1 . Every permutation can be written as a product of disjoint cycles where each element that does not occur in a cycle is understood as being mapped to itself. We define a relation \sim on Ω with $\alpha \sim \beta$ if and only if there is a permutation $\mathfrak{g} \in \mathfrak{G}$ such that $\alpha^{\mathfrak{g}} = \beta$. The relation partitions Ω into equivalence classes which we call *orbits*. We use the notation $\alpha^{\mathfrak{G}}$ to denote the orbit $\{\alpha^{\mathfrak{g}} \mid \mathfrak{g} \in \mathfrak{G}\}$ containing α . Let $f : \Omega \rightarrow \mathbb{R}$ be a function from Ω into the real numbers and let \mathfrak{G} be a permutation group acting on Ω . We say that \mathfrak{G} is an *automorphism group for* (Ω, f) if and only if for all $\omega \in \Omega$ and all $\mathfrak{g} \in \mathfrak{G}$, $f(\omega) = f(\omega^{\mathfrak{g}})$.

2.2 Finite Markov chains

Given a finite set Ω a *finite Markov chain* defines a random walk (X_0, X_1, \dots) on elements of Ω with the property that the conditional distribution of X_{n+1} given (X_0, X_1, \dots, X_n) depends only on X_n . For all $x, y \in \Omega$ $P(x, y)$ is the chain's probability to transition from x to y , and $P^t(x, y) = P_x^t(y)$ the probability of being in state y after t steps if the chain starts at x . A Markov chain is *irreducible* if for all $x, y \in \Omega$ there exists a t such that $P^t(x, y) > 0$ and *aperiodic* if for all $x \in \Omega$, $\gcd\{t \geq 1 \mid P^t(x, x) > 0\} = 1$. A chain that is both irreducible and aperiodic converges to its unique stationary distribution.

The total variation distance d_{tv} of the Markov chain

from its stationary distribution π at time t with initial state x is defined by

$$d_{\text{tv}}(P_x^t, \pi) = \frac{1}{2} \sum_{y \in \Omega} |P^t(x, y) - \pi(y)|.$$

For $\varepsilon > 0$, let $\tau_x(\varepsilon)$ denote the least value T such that $d_{\text{tv}}(P_x^t, \pi) \leq \varepsilon$ for all $t \geq T$. The *mixing time* $\tau(\varepsilon)$ is defined by $\tau(\varepsilon) = \max\{\tau_x(\varepsilon) \mid x \in \Omega\}$. We say that a Markov chain is *rapidly mixing* if the mixing time is bounded by a polynomial in n and $\log(\varepsilon^{-1})$, where n is the size of each configuration in Ω .

2.3 Symmetries in Logic and Probability

Algorithms that leverage model symmetries to solve computationally challenging problems more efficiently exist in several fields. Most of the work is related to the computation of symmetry breaking predicates to improve SAT solver performance [7, 2]. The construction of our symmetry detection approach is largely derived from that of symmetry detection in propositional theories [7, 2]. More recently, similar symmetry detection approaches have been put to work for answer set programming [11] and integer linear programming [34]. Poole introduced the notion of lifted probabilistic inference as a variation of variable elimination taking advantage of the symmetries in graphical models resulting from probabilistic relational formalisms [36]. Following Poole's work, several algorithms for lifted probabilistic inference were developed such as lifted and counting belief propagation [41, 22], bi-simulation-based approximate inference [40], general purpose MCMC algorithm for relational models [29] and, more recently, first-order knowledge compilation techniques [44, 16]. In contrast to existing methods, we present an approach that is applicable to a much larger class of graphical models.

3 Symmetries in Graphical Models

Similar to the method of symmetry detection in propositional formulas [7, 2, 8] we can, for a large class of probabilistic graphical models, construct a colored undirected graph whose automorphism group is equivalent to the permutation group representing the model's symmetries. We describe the approach for sets of partially weighted propositional formulas since Markov logic networks, factor graphs, and the weighted model counting framework can be represented using sets of (partially) weighted formulas [38, 44, 16]. For the sake of readability, we describe the colored graph construction for partially weighted *clauses*. Using a more involved transformation, however, we can extend it to sets of partially weighted formulas. Let $\mathcal{S} = \{(f_i, w_i)\}, 1 \leq i \leq n$, be a set of

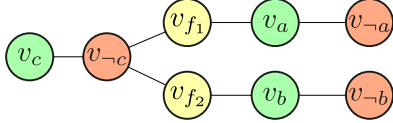


Figure 1: The colored graph resulting from the set of weighted clauses of Example 3.1.

partially weighted clauses with $w_i \in \mathbb{R}$ if f_i is weighted and $w_i = \infty$ otherwise. We define an *automorphism of \mathcal{S}* as a permutation mapping (a) unnegated variables to unnegated variables, (b) negated variables to negated variables, and (c) clauses to clauses, respectively, such that this permutation maps \mathcal{S} to an identical set of partially weighted clauses. The set of these permutations forms the automorphism group of \mathcal{S} .

The construction of the colored undirected graph $G(\mathcal{S}) = (V, E)$ proceeds as follows. For each variable a occurring in \mathcal{S} we add two nodes v_a and v_{-a} modeling the unnegated and negated variable, respectively, to V and the edge $\{v_a, v_{-a}\}$ to E . We assign color 0 (1) to nodes corresponding to negated (unnegated) variables. This coloring precludes permutations that map a negated variable to an unnegated one or vice versa. We introduce a distinct color c_∞ for unweighted clauses and a color c_w for each distinct weight w occurring in \mathcal{S} . For each clause f_i with weight $w_i = w$ we add a node v_{f_i} with color c_w to V . For each unweighted clause f_i we add a node v_{f_i} with color c_∞ to V . Finally, we add edges between each clause node v_{f_i} and the nodes of the negated and unnegated variables occurring in f_i . Please note that we can incorporate evidence by introducing two novel and distinct colors representing *true* and *false* variable nodes.

Example 3.1. Let $\{f_1 := (a \vee \neg c, 0.5), f_2 := (b \vee \neg c, 0.5)\}$ be a set of weighted clauses. We introduce 6 variable nodes $v_a, v_b, v_c, v_{-a}, v_{-b}, v_{-c}$ where the former three have color 1 (green) and the latter three color 0 (red). We connect the nodes v_a and v_{-a} ; v_b and v_{-b} ; and v_c and v_{-c} . We then introduce two new clause nodes v_{f_1}, v_{f_2} both with color 2 (yellow) since they have the same weight. We finally connect the variable nodes with the clause nodes they occur in. Figure 1 depicts the resulting colored graph. A generating set of $\text{Aut}(G(\mathcal{S}))$, the automorphism group of this particular colored graph, is $\{(v_a v_b)(v_{-a} v_{-b})(v_{f_1} v_{f_2})\}$.

The following theorem states the relationship between the automorphisms of \mathcal{S} and the colored graph $G(\mathcal{S})$.

Theorem 3.2. *Let $\mathcal{S} = \{(f_i, w_i)\}, 1 \leq i \leq n$, be a set of partially weighted clauses and let $\text{Aut}(G(\mathcal{S}))$ be the automorphism group of the colored graph constructed for \mathcal{S} . There is a one-to-one correspondence between $\text{Aut}(G(\mathcal{S}))$ and the automorphism group of \mathcal{S} .*

Given a set of partially weighted clauses \mathcal{S} with variables \mathbf{X} we have, by Theorem 3.2, that if we define a distribution Pr over random variables \mathbf{X} with features f_i and weights w_i , $1 \leq i \leq n$, then $\text{Aut}(G(\mathcal{S}))$ is an automorphism group for $(\{0, 1\}^{\mathbf{X}}, \text{Pr})$. Hence, we can use the method to find symmetries in a large class of graphical models. The complexity of computing generating sets of $\text{Aut}(G(\mathcal{S}))$ is in NP and not known to be in P or NP-complete. For graphs with bounded degree the problem is in P [25]. There are specialized algorithms for finding generating sets of automorphism groups of colored graphs such as SAUCY[8] and NAUTY[28] with remarkable performance. We will show that SAUCY computes *irredundant* sets of generators of automorphism groups for graphical models with millions of variables. The size of these generating sets is bounded by the number of graph vertices.

We briefly position the symmetry detection approach in the context of existing algorithms and concepts.

3.1 Lifted Message Passing

There are two different lifted message passing algorithms. Lifted First-Order Belief Propagation [41] operates on Markov logic networks whereas Counting Belief Propagation [22] operates on factor graphs. Both approaches leverage symmetries in the model to partition variables and features into equivalence classes. Each variable class (supernode/clusternode) contains those variable nodes that would send and receive the same messages were (loopy) belief propagation (BP) run on the original model. Each feature class (superfeature/clusterfactor) contains factor nodes that would send and receive the same BP messages.

The colored graph construction provides an alternative approach to partitioning the variables and features of a graphical model. We simply compute the *orbit partition* induced by the permutation group $\text{Aut}(G(\mathcal{S}))$ acting on the set of variables and features. For instance, the orbit partition of Example 3.1 is $\{\{a, b\}, \{c\}, \{f_1, f_2\}\}$. In general, orbit partitions have the following properties: For two variables v_1, v_2 in the same orbit we have that (a) v_1 and v_2 have identical marginal probabilities and (b) the variable nodes corresponding to v_1 and v_2 would send and receive the same messages were BP run on the original model; and for two features f_1 and f_2 in the same orbit we have that the factor nodes corresponding to f_1 and f_2 would send and receive the same BP messages.

3.2 Finite Partial Exchangeability

The notion of exchangeability was introduced by de Finetti [14]. Several theorems concerning finite (partial) exchangeability have been stated [10, 14]. Given

a finite sequence of n binary random variables \mathbf{X} , we say that \mathbf{X} is *exchangeable* with respect to the distribution \Pr if, for every $\mathbf{x} \in \{0, 1\}^n$ and *every* permutation \mathbf{g} acting on $\{0, 1\}^n$, we have that $\Pr(\mathbf{X} = \mathbf{x}) = \Pr(\mathbf{X} = \mathbf{x}^{\mathbf{g}})$. This is equivalent to saying that the symmetric group $\text{Sym}(n)$ is an automorphism group for $(\{0, 1\}^n, \Pr)$. Whenever we have finite exchangeability, there are $n + 1$ orbits each containing the variable assignments with Hamming weight $i, 0 \leq i \leq n$. Hence, every exchangeable probability distribution over n binary random variables is a unique mixture of draws from the $n + 1$ orbits. In some cases of partial exchangeability, namely when the orbits can be specified using a *statistic*, one can use this for a more compact representation of the distribution as a product of mixtures [10]. The symmetries that have to be present for such a re-parameterization to be feasible, however, are rare and constitute one end of the symmetry spectrum.

Therefore, a central question is how *arbitrary symmetries*, compactly represented with irredundant generators of permutation groups, can be utilized for efficient probabilistic inference algorithms that go beyond (a) single variable marginal inference via lifted message passing and (b) the limited applicability of finite partial exchangeability. In order to answer this question, we turn to the major contribution of the present work.

4 Orbital Markov Chains

Inspired by the previous observations, we introduce *orbital Markov chains*, a novel family of Markov chains. An orbital Markov chain is always derived from an existing Markov chain so as to leverage the symmetries in the underlying model. In the presence of symmetries orbital Markov chains are able to perform wide-ranging transitions reducing the time until convergence. In the absence of symmetries they are equivalent to the original Markov chains. Orbital Markov chains only require a generating set of a permutation group \mathfrak{G} acting on the chain’s state space as additional input. As we have seen, these sets of generators are computable with graph automorphism algorithms.

Let Ω be a finite set, let $\mathcal{M}' = (X'_0, X'_1, \dots)$ be a Markov chain with state space Ω , let π be a stationary distribution of \mathcal{M}' , and let \mathfrak{G} be an automorphism group for (Ω, π) . The *orbital Markov chain* $\mathcal{M} = (X_0, X_1, \dots)$ for \mathcal{M}' is a Markov chain which at each integer time $t + 1$ performs the following steps:

1. Let X'_{t+1} be the state of the original Markov chain \mathcal{M}' at time $t + 1$;
2. Sample X_{t+1} , the state of the orbital Markov chain \mathcal{M} at time $t + 1$, uniformly at random from $X'_{t+1}{}^{\mathfrak{G}}$, the orbit of X'_{t+1} .

The orbital Markov chain \mathcal{M} , therefore, runs at every time step $t \geq 1$ the original chain \mathcal{M}' first and samples the state of \mathcal{M} at time t uniformly at random from the orbit of the state of the original chain \mathcal{M}' at time t .

First, let us analyze the complexity of the second step which differs from the original Markov chain. Given a state X_t and a permutation group \mathfrak{G} we need to sample an element from $X_t{}^{\mathfrak{G}}$, the orbit of X_t , uniformly at random. By the orbit-stabilizer theorem this is equivalent to sampling an element $\mathbf{g} \in \mathfrak{G}$ uniformly at random and computing $X_t{}^{\mathbf{g}}$. Sampling group elements nearly uniform at random is a well-researched problem [6] and computable in polynomial time in the size of the generating sets with product replacement algorithms [35]. These algorithms are implemented in several group algebra systems such as GAP[15] and exhibit remarkable performance. Once initialized, product replacement algorithms can generate pseudo-random elements by performing a small number of group multiplications. We could verify that the overhead of step 2 during the sampling process is indeed negligible.

Before we analyze the conditions under which orbital Markov chains are aperiodic, irreducible, and have the same stationary distribution as the original chain, we provide an example of an orbital Markov chain that is based on the standard Gibbs sampler which is commonly used to perform probabilistic inference.

Example 4.1. Let \mathbf{V} be a finite set of random variables with probability distribution π , and let \mathfrak{G} be an automorphism group for $(\times_{V \in \mathbf{V}} V, \pi)$. The orbital Markov chain for the *Gibbs sampler* is a Markov chain $\mathcal{M} = (X_0, X_1, \dots)$ which, being in state X_t , performs the following steps at time $t + 1$:

1. Select a variable $V \in \mathbf{V}$ uniformly at random;
2. Sample $X'_{t+1}(V)$, the value of V in the configuration X'_{t+1} , according to the conditional π -distribution of V given that all other variables take their values according to X_t ;
3. Let $X'_{t+1}(W) = X_t(W)$ for all variables $W \in \mathbf{V} \setminus \{V\}$; and
4. Sample X_{t+1} from $X'_{t+1}{}^{\mathfrak{G}}$, the orbit of X'_{t+1} , uniformly at random.

We call this Markov chain the *orbital Gibbs sampler*. In the absence of symmetries, that is, if \mathfrak{G} ’s only element is the identity permutation, the orbital Gibbs sampler is equivalent to the standard Gibbs sampler.

Let us now state a major result of this paper. It relates properties of the orbital Markov chain to those of the Markov chain it is derived from. A detailed proof can be found in the appendix.

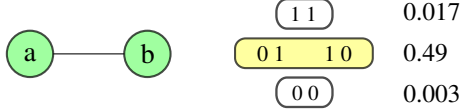


Figure 2: An undirected graphical model over two binary variables and one symmetric potential function with corresponding distribution shown on the right. There are three orbits each indicated by one of the rounded rectangles.

Theorem 4.2. *Let Ω be a finite set and let \mathcal{M}' be a Markov chain with state space Ω and transition matrix P' . Moreover, let π be a probability distribution on Ω , let \mathfrak{G} be an automorphism group for (Ω, π) , and let \mathcal{M} be the orbital Markov chain for \mathcal{M}' . Then,*

- (a) *if \mathcal{M}' is aperiodic then \mathcal{M} is also aperiodic;*
- (b) *if \mathcal{M}' is irreducible then \mathcal{M} is also irreducible;*
- (c) *if π is a reversible distribution for \mathcal{M}' and, for all $\mathfrak{g} \in \mathfrak{G}$ and all $x, y \in \Omega$ we have that $P'(x, y) = P'(x^{\mathfrak{g}}, y^{\mathfrak{g}})$, then π is also a reversible and, hence, a stationary distribution for \mathcal{M} .*

The condition in statement (c) requiring for all $\mathfrak{g} \in \mathfrak{G}$ and all $x, y \in \Omega$ that $P'(x, y) = P'(x^{\mathfrak{g}}, y^{\mathfrak{g}})$ conveys that the original Markov chain is compatible with the symmetries captured by the permutation group \mathfrak{G} . This rather weak assumption is met by all of the practical Markov chains we are aware of and, in particular, Metropolis chains and the standard Gibbs sampler.

Corollary 4.3. *Let \mathcal{M}' be the Markov chain of the Gibbs sampler with reversible distribution π . The orbital Gibbs sampler for \mathcal{M}' is aperiodic and has π as a reversible and, hence, a stationary distribution. Moreover, if \mathcal{M}' is irreducible then the orbital Gibbs sampler is also irreducible and it has π as its unique stationary distribution.*

We will show both analytically and empirically that, in the presence of symmetries, the orbital Gibbs sampler converges at least as fast or faster to the true distribution than state of the art sampling algorithms. First, however, we want to take a look at an example that illustrates the advantages of the orbital Gibbs sampler.

Example 4.4. Consider the undirected graphical model in Figure 2 with two binary random variables and a symmetric potential function. The probabilities of the states 01 and 10 are both 0.49. Due to the symmetry in the model, the states 10 and 01 are part of the same orbit. Now, let us assume a standard Gibbs sampler is in state 10. The probability for it to transition to one of the states 11 and 00 is only 0.02 and,

by definition of the standard Gibbs sampler, it cannot transition directly to the state 01. The chain is “stuck” in the state 10 until it is able to move to 11 or 00. Now, consider the orbital Gibbs sampler. Intuitively, while it is “waiting” to move to one of the low probability states, it samples the two high probability states horizontally uniformly at random from the orbit $\{01, 10\}$. In this particular case the orbital Gibbs sampler converges faster than the standard Gibbs sampler, a fact that we will also show analytically.

4.1 Mixing Time of Orbital Markov Chains

We will make our intuition about the faster convergence of the orbital Gibbs sampler more concrete. We accomplish this by showing that the more symmetry there is in the model the faster a coupling of the orbital Markov chain will coalesce and, therefore, the faster the chain will converge to its stationary distribution.

There are several methods available to prove rapid mixing of a finite Markov chain. The method we will use here is that of a *coupling*. A coupling for a Markov chain \mathcal{M} is a stochastic process (X_t, Y_t) on $\Omega \times \Omega$ such that (X_t) and (Y_t) considered marginally are *faithful copies* of \mathcal{M} . The *coupling lemma* expresses that the total variation distance of \mathcal{M} at time t is limited from above by the probability that the two chains have not *coalesced*, that is, have not met at time t (see for instance Aldous [1]). Coupling proofs on the joint space $\Omega \times \Omega$ are often rather involved and require complex combinatorial arguments. A possible simplification is provided by the *path coupling* method where a coupling is only required to hold on a subset of $\Omega \times \Omega$ (Bubley and Dyer [4]). The following theorem formalizes this idea.

Theorem 4.5 (Dyer and Greenhill [12]). *Let δ be an integer valued metric defined on $\Omega \times \Omega$ taking values in $\{0, \dots, D\}$. Let $S \subseteq \Omega \times \Omega$ such that for all $(X_t, Y_t) \in \Omega \times \Omega$ there exists a path $X_t = Z_0, \dots, Z_r = Y_t$ between X_t and Y_t with $(Z_l, Z_{l+1}) \in S$ for $0 \leq l < r$ and*

$$\sum_{l=0}^{r-1} \delta(Z_l, Z_{l+1}) = \delta(X_t, Y_t).$$

Define a coupling $(X, Y) \rightarrow (X', Y')$ of the Markov chain \mathcal{M} on all pairs $(X, Y) \in S$. Suppose there exists $\beta \leq 1$ with $\mathbf{E}[\delta(X', Y')] \leq \beta \delta(X, Y)$ for all $(X, Y) \in S$. If $\beta < 1$ then the mixing time $\tau(\varepsilon)$ of \mathcal{M} satisfies

$$\tau(\varepsilon) \leq \frac{\ln(D\varepsilon^{-1})}{1 - \beta}.$$

If $\beta = 1$ and there exists an $\alpha > 0$ such that $\Pr[\delta(X_{t+1}, Y_{t+1}) \neq \delta(X_t, Y_t)] \geq \alpha$ for all t , then

$$\tau(\varepsilon) \leq \left\lceil \frac{eD^2}{\alpha} \right\rceil \lceil \ln(\varepsilon^{-1}) \rceil.$$

We selected the *insert/delete* Markov chain for independent sets of graphs for our analysis. Sampling independent sets is a classical problem motivated by numerous applications and with a considerable amount of recent research devoted to it [24, 13, 45, 42, 37]. The coupling proof for the orbital version of this Markov chain provides interesting insights into the construction of such a coupling and the influence of the graph symmetries on the mixing time. The proof strategy is in essence applicable to other sampling algorithms.

Let $G = (V, E)$ be a graph. A subset X of V is an *independent set* if $\{v, w\} \notin E$ for all $v, w \in X$. Let $\mathcal{I}(G)$ be the set of all independent sets in a given graph G and let λ be a positive real number. The partition function $Z = Z(\lambda)$ and the corresponding probability measure π_λ on $\mathcal{I}(G)$ are defined by

$$Z = Z(\lambda) = \sum_{X \in \mathcal{I}(G)} \lambda^{|X|} \quad \text{and} \quad \pi_\lambda(X) = \frac{\lambda^{|X|}}{Z}.$$

Approximating the partition function and sampling from $\mathcal{I}(G)$ can be accomplished using a rapidly mixing Markov chain with state space $\mathcal{I}(G)$ and stationary distribution π_λ . The simplest Markov chain for independent sets is the so-called *insert/delete* chain [13]. If X_t is the state at time t then the state at time $t+1$ is determined by the following procedure:

1. Select a vertex $v \in V$ uniformly at random;
2. If $v \in X_t$ then let $X_{t+1} = X_t \setminus \{v\}$ with probability $1/(1+\lambda)$;
3. If $v \notin X_t$ and v has no neighbors in X_t then let $X_{t+1} = X_t \cup \{v\}$ with probability $\lambda/(1+\lambda)$;
4. Otherwise let $X_{t+1} = X_t$.

Using a path coupling argument one can show that the *insert/delete* chain is rapidly mixing for $\lambda \leq 1/(\Delta-1)$ where Δ is the maximum degree of the graph [13]. We can turn the *insert/delete* Markov chain into the orbital *insert/delete* Markov chain $\mathcal{M}(\mathcal{I}(G))$ simply by adding the following fifth step:

5. Sample X_{t+1} uniformly at random from its orbit.

By Corollary 4.3 the orbital *insert/delete* chain for independent sets is aperiodic, irreducible, and has π_λ as its unique stationary distribution. We can now state the following theorem concerning the mixing time of this Markov chain. It relates the graph symmetries to the mixing time of the chain. The proof of the theorem is based on a path coupling that is constructed so as to make the two chains coalesce whenever their respective states are located in the same orbit. A detailed and instructive proof can be found in the appendix.

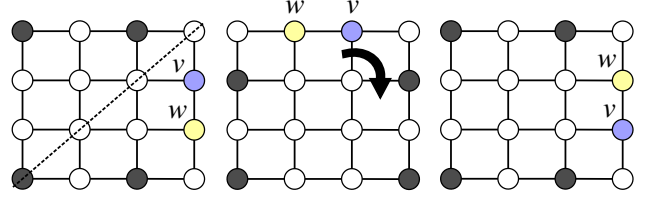


Figure 3: Two independent sets $X \cup \{v\}$ and $X \cup \{w\}$ of the 4×4 grid located in the same orbit. The first permutation is the reflection on the sketched diagonal the second a clockwise 90° rotation.

Theorem 4.6. *Let $G = (V, E)$ be a graph with maximum degree Δ , let λ be a positive real number, and let \mathfrak{G} be an automorphism group for $(\{0, 1\}^V, \pi_\lambda)$. Moreover, let $(X \cup \{v\}) \in \mathcal{I}(G)$, let $(X \cup \{w\}) \in \mathcal{I}(G)$, let $\{v, w\} \in E$, and let $\rho = \Pr[(X \cup \{v\}) \notin (X \cup \{w\})^\mathfrak{G}]$. The orbital insert/delete chain $\mathcal{M}(\mathcal{I}(G))$ is rapidly mixing if either $\rho \leq 0.5$ or $\lambda \leq 1/((2\rho-1)\Delta-1)$.*

The theorem establishes the important link between the graph automorphisms and the mixing time of the orbital *insert/delete* chain. The more symmetries the graph exhibits the larger the orbits and the sooner the chains coalesce. Figure 3 depicts the 4×4 grid with two independent sets $X \cup \{v\}$ and $X \cup \{w\}$ with $\{v, w\} \in E$ and $(X \cup \{v\}) \in (X \cup \{w\})^\mathfrak{G}$. Since $\rho < 1$ for $n \times n$ grids, $n \geq 4$, we can prove (a) rapid mixing of the orbital *insert/delete* chain for larger λ values and (b) more rapid mixing for identical λ values.

The next corollary follows from Theorem 4.6 and the simple fact that $X' \in X^\mathfrak{G}$ for all $X \subseteq V, X' \subseteq V$ with $|X'| = |X|$ whenever \mathfrak{G} is the symmetric group on V .

Corollary 4.7. *Let $G = (V, E)$ be a graph, let λ be a positive real number, and let \mathfrak{G} be an automorphism group for $(\{0, 1\}^V, \pi_\lambda)$. If \mathfrak{G} is the symmetric group $\text{Sym}(V)$ then $\mathcal{M}(\mathcal{I}(G))$ is rapidly mixing with $\tau(\varepsilon) \leq |V| \ln(|V|\varepsilon^{-1})$.*

By analyzing the coupling proof of Theorem 4.6 and, in particular, the moves leading to states (X', Y') with $|X'| = |Y'|$ and the probability that X' and Y' are located in the same orbit in these cases, it is possible to provide more refined bounds. Moreover, to capture the full power of orbital Markov chains, a coupling argument should not merely consider pairs of states with Hamming distance 1. Indeed, the strength of the orbital chains is that, in the presence of symmetries in the graph topology, there is a non-zero probability that states with large Hamming distance (up to $|V|$) are located in the same orbit. The method presented here is also applicable to Markov chains known to mix rapidly for larger λ values than the *insert/delete* chain such as the *insert/delete/drag* chain [13].

social network model [41]					
people	20	50	100	250	500
vertices	1740	10200	40700	251750	1003500
edges	2120	10350	50600	314000	1253000
time [s]	0.04	0.15	0.81	22.5	261.3
features	860	5150	20300	125750	501500
orbs w/o	7	7	7	7	7
orbs w/	238	1244	6237	30192	78303
$k \times k$ grid model					
k	20	50	100	250	500
vertices	800	5000	20000	125000	500000
edges	1160	7400	29800	187000	749000
time [s]	0.02	0.03	0.2	0.6	2.5

Table 1: Number of vertices and edges of the colored graphs, the runtime of SAUCY, and the number of (super-)features of the social network model without and with 10% evidence.

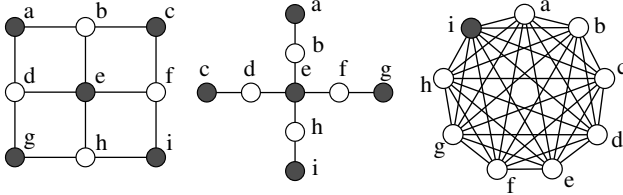


Figure 4: From left to right: the 3-grid, the 3-connected cliques, and the 3-complete graph models.

5 Experiments

Two graphical models were used to evaluate the symmetry detection approach. The “Friends & Smokers” Markov logic network where for a random 10% of all people it is known (a) whether they smoke or not and (b) who 10 of their friends are [41]. Moreover, we used the $k \times k$ grid model, an established and well-motivated lattice model with numerous applications [37]. All experiments were conducted on a PC with an AMD Athlon dual core 5400B 1.0 GHz processor and 3 GB RAM. Table 1 lists the results for varying model sizes. SAUCY’s runtime scales roughly quadratic with the number of vertices and it performs better for the $k \times k$ grid models. This might be due to the larger sets of generators for the permutation groups of the social network model. Table 1 also lists the number of features of the ground social network model (features), the number of feature orbits without (orbs w/o) and with (orbs w/) 10% evidence.

We proceeded to compare the performance of the orbital Markov chains with state-of-the-art algorithms for sampling independent sets. We used GAP[15], a system for computational discrete algebra, and the ORB package[31]¹ to implement the sampling algo-

¹<http://www.gap-system.org/Packages/orb.html>

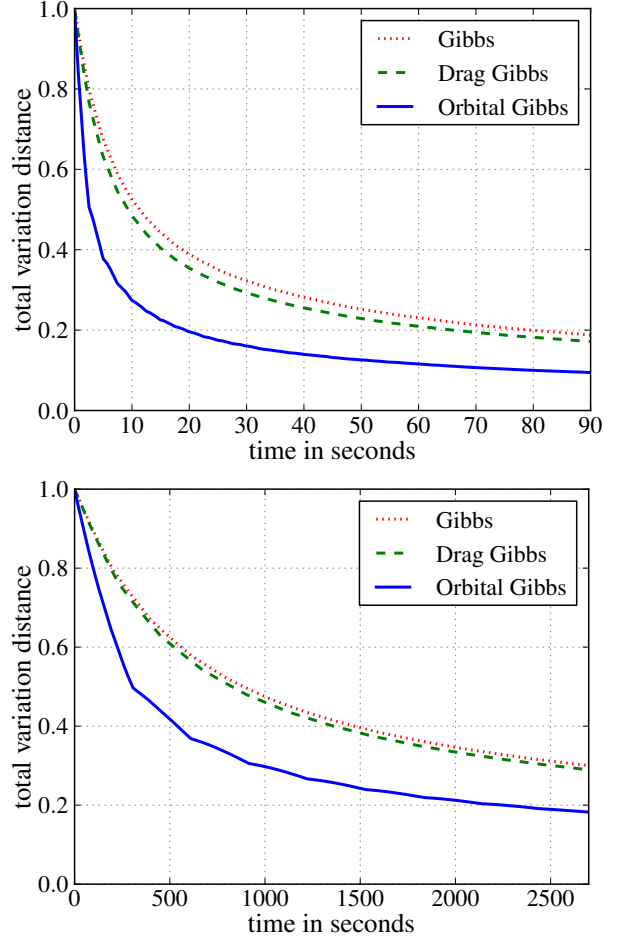


Figure 5: The results of the three Gibbs samplers for the 5-grid model (top) and the 6-grid model (bottom).

gorithms. The experiments can easily be replicated by installing GAP and the ORB package and by running the GAP files available at a dedicated code repository². For the evaluation of the sampling algorithms we selected three different graph topologies exhibiting varying degrees of symmetry:

The k -grid model is the 2-dimensional $k \times k$ grid. An instance of the model for $k = 3$ is depicted in Figure 4 (left). Here, the generating set of the permutation group \mathfrak{G} computed by SAUCY is $\{(a\ c)(d\ f)(g\ i), (a\ i)(b\ f)(d\ h)\}$ and $|\mathfrak{G}| = 8$. The permutation group \mathfrak{G} partitions the set $\{0, 1\}^9$ in 102 orbits with each orbit having a cardinality in $\{1, 2, 4, 8\}$.

The k -connected cliques model is a graph with $k + 1$ distinct cliques each of size $k - 1$ and each connected with one edge to the same vertex. Statistical relational formalisms such as Markov logic networks often lead to similar graph topologies. An in-

²<http://code.google.com/p/lifted-mcmc/>

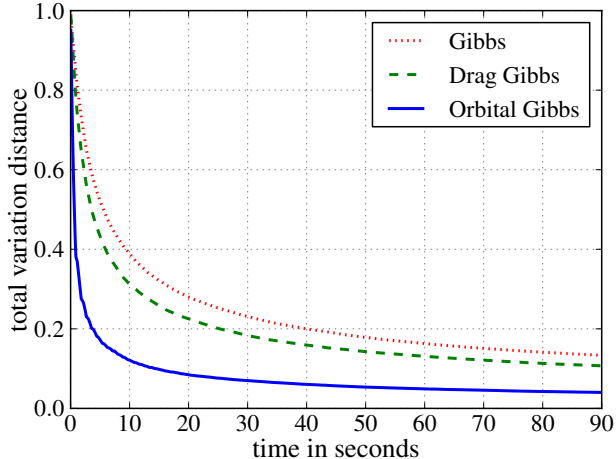


Figure 6: The results of the three Gibbs samplers for the 5-connected cliques model.

stance for $k = 3$ is depicted in Figure 4 (center). Here, the generating set of \mathfrak{G} computed by SAUCY is $\{(a g)(b f), (a c)(b d), (a i)(b h)\}$ and $|\mathfrak{G}| = 24$. The permutation group \mathfrak{G} partitions the set $\{0, 1\}^9$ in 70 orbits with cardinalities in $\{1, 4, 6, 12, 24\}$.

The k -complete graph model is a complete graph with k^2 vertices. Figure 4 (right) depicts an instance for $k = 3$. Here, the generating set of \mathfrak{G} computed by SAUCY is $\{(b c), (b d), (b e), (b f), (b g), (b h), (b i), (a b)\}$ and $|\mathfrak{G}| = 9! = 362880$. The permutation group \mathfrak{G} partitions the set $\{0, 1\}^9$ in 10 orbits with each orbit having a cardinality in $\{1, 9, 36, 84, 126\}$.

SAUCY needed at most 5 ms to compute the sets of generators for the permutation groups of the three models for $k = 6$. We generated samples of the probability measure π_λ on $\mathcal{I}(G)$ for $\lambda = 1$ and the three different graph topologies by running (a) the *insert/delete* chain, (b) the *insert/delete/drag* chain [13], and (c) the orbital *insert/delete* chain. Each chain was started in the state corresponding to the empty set and no burn-in period was used. The orbital *insert/delete* chain did not require more RAM and needed 50 microseconds per sample which amounts to an overhead of about 25% relative to the 40 microseconds of the *insert/delete* chain. The 25% overhead remained constant and independent of the size of the graphs. Since the sampling algorithms create large files with all accumulated samples, I/O overhead is included in these times. For each of the three topologies and each of the three Gibbs samplers, we computed the total variation distance between the distribution approximated using *all* accumulated samples and the true distribution π_1 . Figure 5 plots the total variation distance over elapsed time for the k -grid model for $k = 5$ and $k = 6$. The

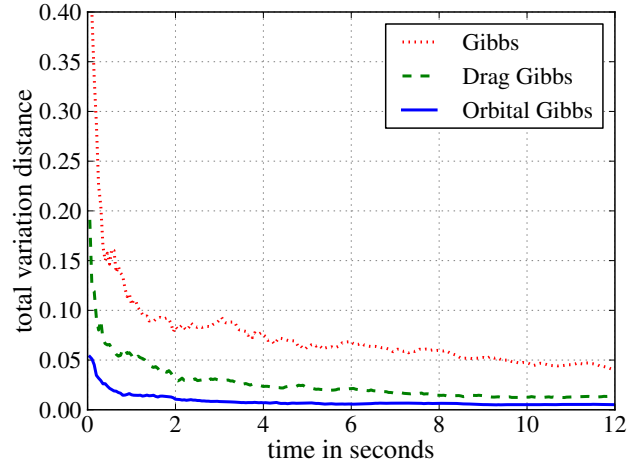


Figure 7: The results of the three Gibbs samplers for the 5-complete graph model.

orbital *insert/delete* chain (Orbital Gibbs) converges the fastest. The *insert/delete/drag* chain (Drag Gibbs) converges faster than the *insert/delete* chain (Gibbs) but since there is a small computational overhead of the *insert/delete/drag* chain the difference is less pronounced for $k = 6$. The same results are observable for the other graph topologies (see Figures 6 and 7) where the orbital Markov chain outperforms the others. In summary, the larger the cardinalities of the orbits induced by the symmetries the faster converges the orbital Gibbs sampler relative to the other chains.

6 Discussion

The mindful reader might have recognized a similarity to *lumping* of Markov chains which amounts to partitioning the state space of the chain [5]. Computing the coarsest lumping quotient of a Markov chain with a bisimulation procedure is linear in the *number of non-zero probability transitions* of the chain and, hence, in most cases exponential in the number of random variables. Since merely counting equivalence classes in the Pólya theory setting is a #P-complete problem [18] there are clear computational limitations to this approach. Orbital Markov chains, on the other hand, combine the advantages of a compact representation of symmetries as generating sets of permutation groups with highly efficient product replacement algorithms and, therefore, provide the advantages of *lumping* while avoiding the intractable explicit computation of the partition of the state space.

One can apply orbital Markov chains to other graphical models that exhibit symmetries such as the Ising model. Since Markov chains in general and Gibbs samplers in particular are components in numerous algo-

rithms (cf. [43, 20, 38, 19, 23, 3]), we expect orbital Markov chains to improve the algorithms' performance when applied to problems that exhibit symmetries. For instance, sampling algorithms for statistical relational languages are obvious candidates for improvement. Future work will include the integration of orbital Markov chains with algorithms for marginal as well as maximum a-posteriori inference. We will also apply the symmetry detection approach to make existing inference algorithms more efficient by, for instance, using symmetry breaking constraints in combinatorial optimization approaches to maximum a-posteriori inference in Markov logic networks (cf. [39, 32, 33]).

While we have shown that permutation groups are computable with graph automorphism algorithms for a large class of models it is also possible to *assume* certain symmetries in the model in the same way (conditional) independencies are assumed in the design stage of a probabilistic graphical model. Orbital Markov chains could easily incorporate these symmetries in form of permutation groups.

Acknowledgments

Many thanks to Guy Van den Broeck, Kristian Kersting, Martin Mladenov, and Babak Ahmadi for insightful discussions concerning lifted inference, to Jürgen Müller for helpful remarks on the product replacement algorithm, and to all those who have contributed to the GAP system and the ORB package.

References

- [1] D. Aldous. Random walks on finite groups and rapidly mixing Markov chains. In *Seminaire de Probabilités XVII*, pages 243–297. 1983.
- [2] F. A. Aloul, K. A. Sakallah, and I. L. Markov. Efficient symmetry breaking for boolean satisfiability. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 271–276, 2003.
- [3] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, Mar. 2003.
- [4] R. Buble and M. Dyer. Path coupling: A technique for proving rapid mixing in markov chains. In *Proceedings of the 38th Symposium on Foundations of Computer Science*, pages 223–231, 1997.
- [5] P. Buchholz. Exact and ordinary lumpability in finite markov chains. *Journal of Applied Probability*, 31(1):pp. 59–75, 1994.
- [6] F. Celler, C. R. Leedham-Green, S. H. Murray, A. C. Niemeyer, and E. O'brien. Generating random elements of a finite group. *Communications in Algebra*, 23(13):4931–4948, 1995.
- [7] J. Crawford. A theoretical analysis of reasoning by symmetry in first-order logic. In *Proceedings of the Workshop on Tractable Reasoning*, 1992.
- [8] P. T. Darga, K. A. Sakallah, and I. L. Markov. Faster symmetry discovery using sparsity of symmetries. In *Proceedings of the 45th annual Design Automation Conference*, pages 149–154, 2008.
- [9] R. de Salvo Braz, E. Amir, and D. Roth. Lifted first-order probabilistic inference. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, pages 1319–1325, 2005.
- [10] P. Diaconis and D. Freedman. De finetti's generalizations of exchangeability. In *Studies in Inductive Logic and Probability*, volume II. 1980.
- [11] C. Drescher, O. Tifrea, and T. Walsh. Symmetry-breaking answer set solving. *AI Commun.*, 24(2):177–194, 2011.
- [12] M. Dyer and C. Greenhill. A more rapidly mixing markov chain for graph colorings. *Random Struct. Algorithms*, 13(3-4):285–317, 1998.
- [13] M. Dyer and C. Greenhill. On markov chains for independent sets. *Journal of Algorithms*, 35(1):17–49, 2000.
- [14] B. Finetti. *Probability, induction and statistics: the art of guessing*. Probability and mathematical statistics. Wiley, 1972.
- [15] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.4.12*, 2008.
- [16] V. Gogate and P. Domingos. Probabilistic theorem proving. In *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence*, pages 256–265, 2011.
- [17] V. Gogate, A. Jha, and D. Venugopal. Advances in lifted importance sampling. In *Proceedings of the 26th Conference on Artificial Intelligence*, 2012.
- [18] L. A. Goldberg. Computation in permutation groups: counting and randomly sampling orbits. In *Surveys in Combinatorics*, pages 109–143. Cambridge University Press, 2001.
- [19] C. Gomes, J. Hoffmann, A. Sabharwal, and B. Selman. From sampling to model counting. In *Proceedings of the 20th International Joint*

- Conference on Artificial Intelligence*, pages 2293–2299, 2007.
- [20] O. Häggström. *Finite Markov Chains and Algorithmic Applications*. London Mathematical Society Student Texts. Cambridge University Press, 2002.
- [21] M. Jaeger. Lower complexity bounds for lifted inference. *CoRR*, abs/1204.3255, 2012.
- [22] K. Kersting, B. Ahmadi, and S. Natarajan. Counting belief propagation. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, pages 277–284, 2009.
- [23] P. Liang, M. I. Jordan, and B. Taskar. A permutation-augmented sampler for dp mixture models. In *Proceedings of the 24th International Conference on Machine Learning*, pages 545–552, 2007.
- [24] M. Luby and E. Vigoda. Fast convergence of the glauber dynamics for sampling independent sets. *Random Struct. Algorithms*, 15(3-4):229–241, 1999.
- [25] E. M. Luks. Isomorphism of graphs of bounded valence can be tested in polynomial time. *J. Computer and System Sciences*, 25(1):42–65, 1982.
- [26] F. Margot. Exploiting orbits in symmetric ilp. *Math. Program.*, 98(1-3):3–21, 2003.
- [27] F. Margot. Symmetry in integer linear programming. In *50 Years of Integer Programming 1958–2008*, pages 647–686. Springer-Verlag, 2010.
- [28] B. D. McKay. Practical graph isomorphism. *Congressus Numerantium*, 30:45–87, 1981.
- [29] B. Milch and S. J. Russell. General-purpose mcmc inference over relational structures. In *Proceedings of the 22nd Conference in Uncertainty in Artificial Intelligence*, pages 349–358, 2006.
- [30] M. Mladenov, B. Ahmadi, and K. Kersting. Lifted linear programming. In *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics*, 2012.
- [31] J. Müller, M. Neunhöffer, and R. Wilson. Enumerating big orbits and an application: b acting on the cosets of fi_{23} . *Journal of Algebra*, 314(1):75–96, 2007.
- [32] M. Niepert. A delayed column generation strategy for exact k -bounded map inference in markov logic networks. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, pages 384–391, 2010.
- [33] M. Niepert, J. Noessner, and H. Stuckenschmidt. Log-linear description logics. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, pages 2153–2158, 2011.
- [34] J. Ostrowski, J. Linderöth, F. Rossi, and S. Smriglio. Orbital branching. *Math. Program.*, 126(1):147–178, 2011.
- [35] I. Pak. The product replacement algorithm is polynomial. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 476–485, 2000.
- [36] D. Poole. First-order probabilistic inference. In *Proceedings of the 18th Joint Conference on Artificial Intelligence*, pages 985–991, 2003.
- [37] R. Restrepo, J. Shin, P. Tetali, E. Vigoda, and L. Yang. Improved mixing condition on the grid for counting and sampling independent sets. In *Proceedings of the 52nd Symposium on Foundations of Computer Science*, pages 140–149, 2011.
- [38] M. Richardson and P. Domingos. Markov logic networks. *Machine Learning*, 62(1-2), 2006.
- [39] S. Riedel. Improving the accuracy and efficiency of map inference for markov logic. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 468–475, 2008.
- [40] P. Sen, A. Deshpande, and L. Getoor. Bisimulation-based approximate lifted inference. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, 2009.
- [41] P. Singla and P. Domingos. Lifted first-order belief propagation. In *Proceedings of the 23rd Conference on Artificial Intelligence*, pages 1094–1099, 2008.
- [42] A. Sly. Computational transition at the uniqueness threshold. In *Proceedings of the 51st Symposium on Foundations of Computer Science*, pages 287–296, 2010.
- [43] M. Stephens. Dealing with label switching in mixture models. *Journal of the Royal Statistical Society Series B*, 62(4):795–809, 2000.
- [44] G. Van den Broeck. On the completeness of first-order knowledge compilation for lifted probabilistic inference. In *Advances in Neural Information Processing Systems*, pages 1386–1394, 2011.
- [45] D. Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the 38th ACM Symposium on Theory of computing*, pages 140–149, 2006.

A Proof of Theorem 4.2

We first prove (a). Since \mathcal{M}' is aperiodic we have, for each state $x \in \Omega$ and every time step $t \geq 0$, a non-zero probability for the Markov chain \mathcal{M}' to remain in state x at time $t + 1$. At each time $t + 1$, the orbital Markov chain transitions uniformly at random to one of the states in the orbit of the original chain's state at time $t + 1$. Since every state is an element of its own orbit, we have, for every state $x \in \Omega$ and every time step $t \geq 0$, a non-zero probability for the Markov chain \mathcal{M} to remain in state x at time $t + 1$. Hence, \mathcal{M} is aperiodic. The proof of statement (b) is accomplished in an analogous fashion and omitted.

Let $P(x, y)$ and $P'(x, y)$ be the probabilities of \mathcal{M} and \mathcal{M}' , respectively, to transition from state x to state y . Since π is a reversible distribution for \mathcal{M}' we have that $\pi(x)P'(x, y) = \pi(y)P'(y, x)$ for all states $x, y \in \Omega$. For every state $x \in \Omega$ let $x^{\mathfrak{G}}$ be the orbit of x . Let $\mathfrak{G}_x := \{\mathfrak{g} \in \mathfrak{G} \mid x^{\mathfrak{g}} = x\}$ be the stabilizer subgroup of x with respect to \mathfrak{G} . We have that

$$\begin{aligned} \sum_{\mathfrak{g} \in \mathfrak{G}} P'(x, y^{\mathfrak{g}}) &= \sum_{y' \in y^{\mathfrak{G}}} |\mathfrak{G}_{y'}| P'(x, y') \\ &= |\mathfrak{G}_y| \sum_{y' \in y^{\mathfrak{G}}} P'(x, y') \\ &= (|\mathfrak{G}|/|y^{\mathfrak{G}}|) \sum_{y' \in y^{\mathfrak{G}}} P'(x, y') \end{aligned} \quad (1)$$

where the last two equalities follow from the orbit-stabilizer theorem. We will now prove that $\pi(x)P(x, y) = \pi(y)P(y, x)$ for all states $x, y \in \Omega$. By definition of the orbital Markov chain we have that $\pi(x)P(x, y) = \pi(x)(1/|y^{\mathfrak{G}}|) \sum_{y' \in y^{\mathfrak{G}}} P'(x, y')$ and, by equation (1), $\pi(x)(1/|y^{\mathfrak{G}}|) \sum_{y' \in y^{\mathfrak{G}}} P'(x, y')$

$$\begin{aligned} &= \pi(x)(1/|y^{\mathfrak{G}}|)(|y^{\mathfrak{G}}|/|\mathfrak{G}|) \sum_{\mathfrak{g} \in \mathfrak{G}} P'(x, y^{\mathfrak{g}}) \\ &= \pi(x)(1/|\mathfrak{G}|) \sum_{\mathfrak{g} \in \mathfrak{G}} P'(x, y^{\mathfrak{g}}) \\ &= (1/|\mathfrak{G}|) \sum_{\mathfrak{g} \in \mathfrak{G}} \pi(x)P'(x, y^{\mathfrak{g}}). \end{aligned}$$

Since P' is reversible and $\pi(x) = \pi(x^{\mathfrak{g}})$ for all $\mathfrak{g} \in \mathfrak{G}$ we have $(1/|\mathfrak{G}|) \sum_{\mathfrak{g} \in \mathfrak{G}} \pi(x)P'(x, y^{\mathfrak{g}}) = (1/|\mathfrak{G}|) \sum_{\mathfrak{g} \in \mathfrak{G}} \pi(y^{\mathfrak{g}})P'(y^{\mathfrak{g}}, x) = \pi(y)(1/|\mathfrak{G}|) \sum_{\mathfrak{g} \in \mathfrak{G}} P'(y^{\mathfrak{g}}, x)$. Now, since $P'(x, y) = P'(x^{\mathfrak{g}}, y^{\mathfrak{g}})$ for all $x, y \in \Omega$ and all $\mathfrak{g} \in \mathfrak{G}$ by assumption, we have that $\pi(y)(1/|\mathfrak{G}|) \sum_{\mathfrak{g} \in \mathfrak{G}} P'(y^{\mathfrak{g}}, x) = \pi(y)(1/|\mathfrak{G}|) \sum_{\mathfrak{g} \in \mathfrak{G}} P'(y, x^{-\mathfrak{g}}) = \pi(y)(1/|\mathfrak{G}|) \sum_{\mathfrak{g} \in \mathfrak{G}} P'(y, x^{\mathfrak{g}})$ and, again by equation (1), $\pi(y)(1/|\mathfrak{G}|) \sum_{\mathfrak{g} \in \mathfrak{G}} P'(y, x^{\mathfrak{g}})$

$$\begin{aligned} &= \pi(y)(1/|\mathfrak{G}|)(|\mathfrak{G}|/|x^{\mathfrak{G}}|) \sum_{x' \in x^{\mathfrak{G}}} P'(y, x') \\ &= \pi(y)(1/|x^{\mathfrak{G}}|) \sum_{x' \in x^{\mathfrak{G}}} P'(y, x') = \pi(y)P(y, x). \quad \square \end{aligned}$$

B Proof of Theorem 4.6

Let $H : \Omega \times \Omega \rightarrow \mathbb{N}$ be the Hamming distance between any two elements in Ω . We provide a path coupling argument on the set of pairs having Hamming distance 1. Let X and Y be two independent sets which differ only at one vertex v with degree d . We assume, without loss of generality, that $v \in X \setminus Y$. Choose a vertex w uniformly at random. We distinguish five cases:

- (i) if $w = v$ then sample one $\mathfrak{g} \in \mathfrak{G}$ uniformly at random and let $(X', Y') = (X^{\mathfrak{g}}, X^{\mathfrak{g}})$ with probability $\lambda/(1+\lambda)$, otherwise let $(X', Y') = (Y^{\mathfrak{g}}, Y^{\mathfrak{g}})$; Hence, $H(X', Y') = 0$ with probability 1.
- (ii) if $w \neq v$ and $w \in X$ then sample one $\mathfrak{g} \in \mathfrak{G}$ uniformly at random and let $(X', Y') = ((X \setminus \{w\})^{\mathfrak{g}}, (Y \setminus \{w\})^{\mathfrak{g}})$ with probability $1/(1+\lambda)$, otherwise let $(X', Y') = (X^{\mathfrak{g}}, Y^{\mathfrak{g}})$. In both cases, we have that $H(X', Y') = 1$.
- (iii) if $w \neq v$, $w \notin X$ and w has no neighbor in X then sample one $\mathfrak{g} \in \mathfrak{G}$ uniformly at random and let $(X', Y') = ((X \cup \{w\})^{\mathfrak{g}}, (Y \cup \{w\})^{\mathfrak{g}})$ with probability $\lambda/(1+\lambda)$, otherwise let $(X', Y') = (X^{\mathfrak{g}}, Y^{\mathfrak{g}})$; In both cases, we have that $H(X', Y') = 1$.
- (iv) if $w \neq v$, $w \notin X$ and w has a neighbor in X but not in Y , then sample one $\mathfrak{g} \in \mathfrak{G}$ uniformly at random. Let $\mathfrak{G}' := \{\mathfrak{g}' \in \mathfrak{G} \mid X^{\mathfrak{g}'} = (Y \cup \{w\})^{\mathfrak{g}'}\}$. If $\mathfrak{G}' \neq \emptyset$ then sample one $\mathfrak{g}' \in \mathfrak{G}'$ uniformly at random and let $(X', Y') = (X^{\mathfrak{g}'}, (Y \cup \{w\})^{\mathfrak{g}'})$ with probability $\lambda/(1+\lambda)$. In this case we have $H(X', Y') = 0$. If $\mathfrak{G}' = \emptyset$ then let $(X', Y') = (X^{\mathfrak{g}}, (Y \cup \{w\})^{\mathfrak{g}})$ with probability $\lambda/(1+\lambda)$. Here we have $H(X', Y') = 2$. Otherwise let $(X', Y') = (X^{\mathfrak{g}}, Y^{\mathfrak{g}})$. Here, we have $H(X', Y') = 1$.
- (v) in all other cases sample one $\mathfrak{g} \in \mathfrak{G}$ uniformly at random and let $(X', Y') = (X^{\mathfrak{g}}, Y^{\mathfrak{g}})$. Here we have with probability 1 that $H(X', Y') = 1$.

In summary, we have that

$$\mathbf{E}[H(X', Y') - 1] \leq -\frac{1}{n} + \varrho(2\rho - 1)\frac{\lambda}{(1+\lambda)}$$

where $\varrho = \Pr[w \neq v, w \notin X, \text{ and } w \text{ has a neighbor in } X \text{ but not in } Y]$ and $\rho = \Pr[X \notin (Y \cup \{w\})^{\mathfrak{G}} \mid w \neq v, w \notin X, \text{ and } w \text{ has a neighbor in } X \text{ but not in } Y]$. If $\rho \leq 0.5$ we have that $\mathbf{E}[H(X', Y') - 1] \leq -\frac{1}{n}$ otherwise we have that $\mathbf{E}[H(X', Y') - 1] \leq$

$$-\frac{1}{n} + \frac{d}{n}(2\rho - 1)\frac{\lambda}{(1+\lambda)} \leq -\frac{1}{n} + \frac{\Delta}{n}(2\rho - 1)\frac{\lambda}{(1+\lambda)}.$$

Hence, $\mathcal{M}(\mathcal{I}(G))$ mixes rapidly if either $\rho \leq 0.5$ or $\lambda((2\rho - 1)\Delta - 1) < 1$. For $\lambda((2\rho - 1)\Delta - 1) = 1$ one can verify that there exists an $\alpha > 0$ such that $\Pr[H(X_{t+1}, Y_{t+1}) \neq H(X_t, Y_t)] \geq \alpha$ for all t . \square