A faster method for sampling independent sets

Mark Huber*
Department of Statistics
Stanford University
Stanford, CA 94305-4065
mhuber@orie.cornell.edu

October 20, 1999

Abstract

The problem of finding a random independent set arises in several contexts. In statistical physics, it is known as the hard core gas model. These samples may also be used to approximate the number of independent sets in a graph or to find large independent sets. One sampling approach is to run a Markov chain "for a long time". One such chain for this problem was introduced by Dyer and Greenhill, and here we develop a bounding chain for the Dyer-Greenhill chain. This bounding chain allows us to develop experimental upper bounds on the mixing time of the chain, and to create perfect samples using this chain. An implementation of this algorithm shows that on the two dimensional lattice it is faster than a competing chain, the single-site heat bath chain.

1 The Problem

An independent set $\mathcal{I}$ of a graph is a collection of nodes such that no two nodes of the independent set are adjacent to one another. The distribution we seek to sample from is

$$\pi(\mathcal{I}) = \frac{\lambda^{|\mathcal{I}|}}{Z_\lambda}$$

where Z_λ is the normalization constant that makes π a probability distribution. For example, when $\lambda = 1$, Z_λ is the number of independent sets in the graph, and finding this number is a $\sharp P$-complete problem. The ability to sample quickly from this distribution immediately gives a method to approximate Z_λ (see [7]).

In statistical physics, π is the hard core gas model, where the graph is often a lattice, and the location of gas molecules correspond to nodes in the independent set. This model is one of the simplest to exhibit phase transition behavior [9]. Roughly put, when λ is small the independent sets are highly disorganized, but as λ grows large, long range effects begin to appear, and large structures become highly likely.

It has been shown [2] that when λ rises above a certain value is it NP-hard to obtain a sample from this distribution for general graphs. The techniques we describe here will only be useful below this threshold.

A Markov chain is a stochastic process where $X_{t+1} = \phi(X_t, R_t)$ where X_t is the state of the chain at step t, ϕ is a deterministic function, and R_t is some random bits. Under mild conditions (irreducibility and aperiodicity) the distribution of X_t will approach the stationary distribution of the chain as $t \to \infty$.

Dyer and Greenhill [3] created a Markov chain with π as its stationary distribution. Let $X(v) = 0$ for nodes v not in the independent set, and $X(v) = 1$ for nodes in the independent set.

Dyer and Greenhill hard core chain step
Set $X \leftarrow X_t$
Choose a vertex v uniformly at random from V
Choose U uniformly from $[0,1]$
Case 1: v has no neighbors colored 1 in X_t, then
  If $U \leq \frac{\lambda}{1+\lambda}$: **Set** $X(v) \leftarrow 1$
  Else: **Set** $X(v) \leftarrow 0$
Case 2: v has exactly 1 neighbor w colored 1
  If $U \leq p_{swap}\frac{\lambda}{1+\lambda}$: **Set** $X(v) \leftarrow 1$, $X(w) \leftarrow 0$
  Else: **Set** $X(v) \leftarrow 0$
Set $X_{t+1} \leftarrow X$

Here p_{swap} is a parameter in $[0,1]$ chosen by the user. If Δ is the maximum degree of the graph, this chain with $p_{swap} = 1/4$ is known to be rapidly mixing when $\lambda \leq 2/(\Delta - 2)$. However, we want to obtain samples for larger values of λ.

The Bounding Chain One method for doing so is bounding chains, which provide a basis for perfect sampling algorithms. The Dyer Greenhill chain lives on

*Supported by an NSF postdoc and NSF grants CCR-9307391, CCR-9700029, and DMS-9505155, and ONR through grants N00014-96-1-00500 and AASERT N0001197-1-0881

state space $\{0,1\}^V$. The bounding chain lives on state space $\{\{0\},\{1\},\{0,1\}\}^V$. It is designed so that if Y is a process on the bounding chain and X is a process on the original chain, $X_t(v) \in Y_t(v)\ \forall v$ implies that $X_{t+1}(v) \in Y_{t+1}(v)\ \forall v$. We start with $Y_0(v) = \{0,1\}$ for all v so that trivially $x_0(v) \in Y_0(v)$ for all x_0 and all v. Now if we ever come to a time that $|Y_t(v)| = 1$, then only one state is possibly contained the bounding chain and we say that *complete coalescence* has occurred (see [6, 5, 4] for a more detailed description of bounding chains and their relationship to perfect sampling).

In the bounding chain step, let $a(v)$ be the number of neighbors of v colored $\{0\}$, $b(v)$ the number of neighbors colored $\{1\}$, and $c(v)$ the number colored $\{0,1\}$. Finally, let $d(v)$ be the degree of v, so that $a(v) + b(v) + c(v) = d(v)$.

Bounding chain step for DG chain
Set $Y \leftarrow Y_t$
Choose vertex v uniformly at random from V
Choose U uniformly from $[0,1]$
If $U > \frac{\lambda}{1+\lambda}$ **Set** $Y(v) \leftarrow \{0\}$
Else
Case 1: $a(v) = d(v), b(v) = 0, c(v) = 0$
Set $Y(v) \leftarrow \{1\}$
Case 2: $a(v) = d(v) - 1, b(v) = 1, c(v) = 0$
If $U \leq p_{swap}\frac{\lambda}{1+\lambda}$ and $w \sim v$ with $Y(w) = \{1\}$
Set $Y(v) \leftarrow \{1\}$, $Y(w) \leftarrow \{0\}$
Else: Set $Y(v) \leftarrow \{0\}$
Case 3: $b(v) > 1$: **Set** $Y(v) \leftarrow \{0\}$
Case 4: $a(v) = d(v) - 1, b(v) = 0, c(v) = 1$
If $U \leq p_{swap}\frac{\lambda}{1+\lambda}$ and $w \sim v$ with $Y(w) = \{0,1\}$
Set $Y(v) \leftarrow \{1\}$, $Y(w) = \{0\}$
Else: Set $Y(v) \leftarrow \{0,1\}$
Case 5: $b(v) = 1, c(v) \geq 1$
if $U \leq p_{swap}\frac{\lambda}{1+\lambda}$ and $w \sim v$ with $Y(w) = \{1\}$
Set $Y(v) \leftarrow \{0,1\}$, $Y(w) \leftarrow \{0,1\}$
Case 6: $a(v) = d(v) - c(v), b(v) = 0, c(v) > 1$
Set $Y(v) \leftarrow \{0,1\}$

Determining complete coalescence via the bounding chain gives us a means of getting experimental upper bounds on the mixing time via a classic theorem of Doeblin [1] and a means for generating samples exactly from π using coupling from the past (CFTP) [8] (see [6] for more details). The expected running time of CFTP is proportional to the expected time needed for complete coalescence to be detected.

Another Markov chain used to obtaining samples from π is the single site heat bath chain. Haggstrom and Nelander used antimonotonicity to give a means for determining complete coalescence of this chain [4]. From an extension of [10], we know that this approach will always detect complete coupling in $O(n \ln n)$ time when $\lambda \leq 2/(\Delta - 2)$. Similarly, our approach to bounding the Dyer Greenhill chain can also be shown to only require $O(n \ln n)$ time when $\lambda \leq 2/(\Delta - 2)$.

For $\lambda > 1$, the plot below shows the expected time the two algorithms take to reach complete coalescence when $p_{swap} = 1/4$ on a square lattice with 100 points on a side. As is usual with Monte Carlo Markov chain methods, running time is measured in steps of the Markov chain needed. The factor of speed gain seems to grow roughly exponentially in lambda; when $\lambda = 1$, our method is about twice as fast, but by the time $\lambda = 3.5$, it is roughly 3.6 times as fast.

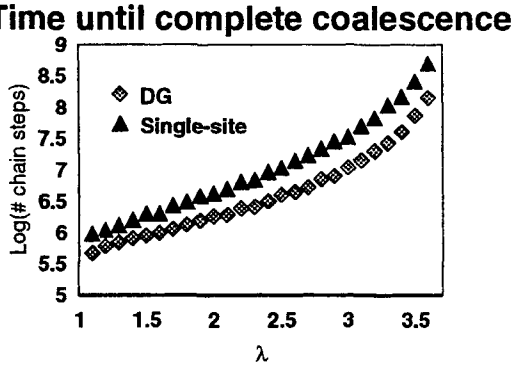

References

[1] W. Doeblin. Exposé de la théorie des chains simples constantes de Markov à un nombre fini d'états. *Rev. Math. de l'Union Interbalkanique*, 2:77–105, 1933.

[2] Martin Dyer, Alan Frieze, and Mark Jerrum. On counting independent sets in sparse graphs. Technical Report ECS-LFCS-98-391, Univ. of Edinburgh, 1998.

[3] Martin Dyer and Catherine Greenhill. On Markov chains for independent sets. *Preprint*, 1997.

[4] Olle Häggström and Karin Nelander. Exact sampling from anti-monotone systems, 1997. Preprint.

[5] Mark L. Huber. Exact sampling and approximate counting techniques. In *Proceedings of the 30th Symposium on the Theory of Computing*, pages 31–40, 1998.

[6] Mark L. Huber. *Perfect Sampling with Bounding Chains*. PhD thesis, Cornell University, 1999.

[7] M. Jerrum, L. Valiant, and V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986.

[8] James G. Propp and David B. Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms*, 9(1 and 2):223–252, 1996.

[9] Barry Simon. *The Statistical Mechanics of Lattice Gasses*, volume 1. Princeton University Press, 1993.

[10] Eric Vigoda. Fast convergence of the Glauber dynamics for sampling independent sets: Part II. Technical Report TR-99-003, International Computer Science Institute, January 1998.