A Dynamic Programming Approach to Efficient Sampling from Boltzmann Distributions

Archis Ghate[a] and Robert L. Smith [b]

[a]Industrial Engineering, University of Washington, Box 352650, Seattle, Washington, 98195, USA, archis@u.washington.edu.

[b]Industrial and Operations Engineering, University of Michigan, Ann Arbor, Michigan, 48109, USA, rlsmith@umich.edu.

Markov chain methods for Boltzmann sampling work in phases with decreasing temperatures. The number of transitions in each phase crucially affects terminal state distribution. We employ dynamic programming to allocate iterations to phases to improve guarantees on sample quality. Numerical experiments on the Ising model are presented.

1. Introduction

Several algorithms in optimization, computer science and statistical mechanics are closely related to the problem of sampling from a Boltzmann distribution parameterized by the so-called temperature over a finite set. These include (i) Simulated Annealing (SA) [6], (ii) calculating the permanent of a nonnegative matrix [1], (iii) estimating the volume of a convex body [7], and (iv) computing partition functions of interacting particle systems such as the Ising model [8].

The finite set involved is typically exceedingly large rendering exact Boltzmann-sampling impractical. Hence one approach is to simulate "several" transitions of an ergodic Markov chain with the appropriate Boltzmann as its limiting distribution so that the final state-distribution well-approximates this Boltzmann. More precisely, to sample from a Boltzmann distribution at temperature $T_* > 0$, it is common to simulate a *sequence of ergodic Markov chains* whose limiting distributions are also Boltzmann at temperatures given by a "cooling schedule"—a strictly decreasing finite sequence of temperatures starting at a very high value, commonly ∞, and ending at T_*.

Initial motivation for implementing the sampling procedure in phases defined by a cooling schedule came from annealing processes in physics where a glass or metal is toughened by cooling it slowly, starting at a high temperature, to a low temperature equilibrium. The mathematical intuition behind gradual cooling is to implement phases such that the distribution of the state at the end of one phase is not very different from the limiting distribution of the next phase with the hope that this will reduce the total number of iterations required to well-approximate the target Boltzmann. This concept is termed a "warm start" and has been featured in recent work in this area [5].

A search for an "optimal" cooling strategy must consider several key questions including how to select the number of Markov transitions in each phase. Specifically, given a fixed number of total iterations (this may typically arise from knowledge of available computational power, and computation time), running too many iterations in early phases may be wasteful, whereas too few may leave us "far away" from the limiting distributions of subsequent phases ("cold start").

1.1. Contribution of this paper

A key contribution of this paper is that we model the above tradeoff as a resource allocation problem where the total number of iterations (resource) is to be allocated to a fixed number of phases (activities). In order to design an appro-

priate objective function for this resource allocation problem, we first note that a crucial performance measure for Markov chain based sampling algorithms is the probability that the final state is distributed exactly according to the limiting distribution. A well-known relation between this probability and the variation distance [9] between the Markov chain state-distribution and its limiting distribution is paramount in our context.

In particular, the Coupling Lemma (see [9] page 275) states that if we observe the state of the Markov chain when the variation distance is ϵ, the probability that the state is distributed exactly according to the limiting distribution is $(1-\epsilon)$. Consequently, the smaller the variation distance, the higher the probability of exact sampling. In our resource allocation formulation, we choose the number of Markov transitions in each phase so as to minimize an upper bound on the variation distance hence maximizing the corresponding lower bound on the probability of exact sampling from the target Boltzmann distribution.

More specifically, we note that the so-called χ^2 [4] distance provides an upper bound on the variation distance and derive an upper bound on this χ^2 distance as a function of the number of transitions in each phase. A benefit of this bound is that it leads to a convex formulation of the two-phase resource allocation problem, which can be solved analytically. This also helps the backward recursion in our dynamic programming procedure for the multi-phase case.

2. Mathematical background and notation

Let X be a finite set with cardinality $|X| = N$. Let $f : X \rightarrow R$ be a real valued function on X. A Boltzmann (T) distribution π_T over X at temperature T is given by $\pi_T(x) = \frac{\exp(\frac{-f(x)}{T})}{Z} \ \ \forall x \in X$, where $Z = \sum_{y \in X} \exp(\frac{-f(y)}{T})$ is the normalization constant. Let $\mathcal{MC}$ be an irreducible, aperiodic and symmetric Markov chain on X. Then the limiting distribution of $\mathcal{MC}$ is uniform over X, denoted U_X.

A common way of designing a Markov chain that converges to π_T over X is to filter states generated by $\mathcal{MC}$ with the Metropolis acceptance criterion as follows. Start with an arbitrary initial state $x_0 \in X$ and let $x_k \in X$ be the state at the beginning of iteration k. Generate candidate $y \in X$ by a one-step transition of $\mathcal{MC}$ from x_k. The next iterate x_{k+1} is set equal to y if $f(y) \leq f(x_k)$. However, if $f(y) > f(x_k)$, the next iterate x_{k+1} is determined probabilistically as follows:

$$x_{k+1} = y \text{ w.p. } \exp\left(\frac{-(f(y)-f(x_k))}{T}\right) \text{ and}$$

$$x_{k+1} = x_k \text{ w.p. } 1-\exp\left(\frac{-(f(y)-f(x_k))}{T}\right).$$

The sequence of iterates $(x_0, x_1, \ldots, x_k \ldots)$ then forms an irreducible, aperiodic Markov chain $\mathcal{MC}_T$ that is reversible with respect to π_T. Thus π_T is the limiting distribution of $\mathcal{MC}_T$. The Markov chain $\mathcal{MC}$ is called the *candidate generator* for the Markov chain $\mathcal{MC}_T$. Observe that $\mathcal{MC}$ can also be viewed as $\mathcal{MC}_\infty$ since every candidate proposed by $\mathcal{MC}$ is accepted by the Metropolis filter when $T = \infty$.

Now consider the case when we employ a cooling schedule $\infty = T_0, \ldots, T_m = T_* > 0$ with $T_i > T_{i+1}$ and $m \geq 1$ to sample approximately from π_{T_*}. Specifically, starting with an arbitrary initial state in X, and initial temperature $T_0 = \infty$, one simulates chain $\mathcal{MC}_{T_i}$ for $K_i \geq 0$ steps in phase i. These chains are simulated sequentially, that is, the state delivered after K_i steps in phase i becomes the initial state in phase $i+1$. The limiting distribution of $\mathcal{MC}_{T_i}$ is π_{T_i}, and in particular, the state delivered at the end of phase m approximates π_{T_*} as desired.

It is well-known that the rate of convergence of a finite-state, irreducible, aperiodic, reversible Markov chain to its limiting distribution say ψ is governed by the Second Largest Eigenvalue Modulus (SLEM) λ^* of its transition matrix [11]. Let $\mathcal{Q}^k$ denote the state-distribution of a Markov chain after $k \geq 0$ steps, where $\mathcal{Q}^0$ stands for the initial distribution. The variation distance between $\mathcal{Q}^k$ and ψ is defined as $\Delta(k) = \frac{1}{2}\sum_{y \in X} |\mathcal{Q}^k(y) - \psi(y)|$. The variation distance is non-increasing in k (see [9] Theorem 11.4, page 280). The χ^2 distance [4] is given by $\chi^2(\mathcal{Q}^k, \psi) = \sum_{y \in X} \frac{(\mathcal{Q}^k(y)-\psi(y))^2}{\psi(y)}$ (the word distance is technically a misnomer here since χ^2 is

not symmetric). The following theorem will prove useful for us as in related work [10] on finite time performance bounds on Simulated Annealing.

Theorem 2.1. *[4] The variation and the χ^2 distances for a finite-state, irreducible, aperiodic, reversible Markov chain with limiting distribution ψ and SLEM λ^* satisfy*

$$2\Delta(k) \leq \sqrt{\chi^2(\mathcal{Q}^k, \psi)} \leq (\lambda^*)^k \sqrt{\chi^2(\mathcal{Q}^0, \psi)}. \quad (1)$$

We will repeatedly apply this result to Markov chains $\mathcal{MC}_{T_i}$ using λ_{T_i}, and Q_{T_i} to denote their SLEM and state distributions respectively for $i = 0, \ldots, m$ (with the exception that we use ∞ instead of T_0 in the subscript). We denote the χ^2 distance between $Q^k_{T_i}$ and π_{T_i} by $\chi^2_k(i)$ for integers $k \geq 0$ and $i = 0, \ldots, m$ for brevity.

3. Problem formulation and solution

All through the sequel we use K to denote a fixed given number of iterations that we will allocate to different phases of sampling. We first consider a two-phase problem where $m = 1$, $\infty = T_0, T_1 = T_*$. Thus, the question now is how to split K iterations between Markov chains $\mathcal{MC}_\infty$ and $\mathcal{MC}_{T_*}$.

3.1. A convex two-phase formulation

Phase Zero

In this phase, we run Markov chain $\mathcal{MC}_\infty$ for K_0 steps starting at $x_0 \in X$. Let $\delta_{x_0}(x)$ be the Dirac-delta function on X that takes value 1 if $x = x_0$ and value 0 otherwise. Observe that $\chi^2_0(0)$ equals

$$\frac{(\delta_{x_0}(x_0) - U_X(x_0))^2}{U_X(x_0)} + \sum_{y \in X \setminus x_0} \frac{(\delta_{x_0}(y) - U_X(y))^2}{U_X(y)}$$
$$= \sum_{y \in X \setminus x_0} \frac{1}{N} + \frac{(1 - \frac{1}{N})^2}{\frac{1}{N}} = (N-1) \leq N.$$

Then by the second inequality in Equation (1) in Theorem 2.1

$$\sqrt{\chi^2_{K_0}(0)} \leq \lambda_\infty^{K_0} \sqrt{\chi^2_0(0)} \leq \lambda_\infty^{K_0} \sqrt{N}. \quad (2)$$

Phase one

In this phase, we run Markov chain $\mathcal{MC}_{T_1}$ (recall $T_1 = T_*$) for K_1 steps. Let $|.|_2$ denote the Euclidean norm in R^N. We employ a standard algebraic manipulation to bound the χ^2 distance using a scaled triangle inequality that involves the Euclidean norm [4,10]. Specifically, $\sqrt{\chi^2_0(1)}$ can be written as

$$= \sqrt{\sum_{y \in X} \frac{(Q^{K_0}_\infty(y) - \pi_{T_1}(y))^2}{\pi_{T_1}(y)}}$$
$$= \left| \left(\frac{(Q^{K_0}_\infty(\cdot) - U_X(\cdot) + U_X(\cdot) - \pi_{T_1}(\cdot))}{\sqrt{\pi_{T_1}(\cdot)}} \right) \right|_2,$$

and the right hand side is bounded above as

$$\leq \left| \left(\frac{(Q^{K_0}_\infty(\cdot) - U_X(\cdot))}{\sqrt{U_X(\cdot)}} \frac{\sqrt{U_X(\cdot)}}{\sqrt{\pi_{T_1}(\cdot)}} \right) \right|_2$$
$$+ \sqrt{\sum_{y \in X} \frac{(U_X(y) - \pi_{T_1}(y))^2}{\pi_{T_1}(y)}}$$
$$\leq \left| \left(\frac{(Q^{K_0}_\infty - U_X)}{\sqrt{U_X}} \right) \right|_2 \max_{y \in X} \left(\sqrt{\frac{U_X(y)}{\pi_{T_1}(y)}} \right)$$
$$+ \max_{y \in X} \left(\frac{U_X(y)}{\pi_{T_1}(y)} - 1 \right).$$

Let $D \equiv D(N) \geq \left(\max_{i \in X} f(i) - \min_{j \in X} f(j) \right)$, an upper bound on the depth of f on X. Note that such an upper bound is typically available from the physics of the problem as for example in the Ising model in Section 4. We focus on the ratio $\frac{U_X(y)}{\pi_{T_1}(y)}$ for any fixed $y \in X$. We have

$$\frac{U_X(y)}{\pi_{T_1}(y)} = \frac{1/N}{\frac{\exp(\frac{-f(y)}{T_1})}{\sum_{z \in X} \exp(\frac{-f(z)}{T_1})}}$$
$$= \frac{1}{N} \sum_{z \in X} \exp(\frac{f(y) - f(z)}{T_1}) \leq \exp(D/T_1).$$

Using this in the above equation, $\sqrt{\chi^2_0(1)}$ can be bounded as follows:

$$\sqrt{\chi^2_0(1)} \leq \lambda_\infty^{K_0} \sqrt{N} \sqrt{\exp(D/T_1)} + (\exp(D/T_1)).$$

Then by the second inequality in Equation (1) in Theorem 2.1, $\sqrt{\chi^2_{K_1}(1)}$ is at most

$$\lambda_{T_1}^{K_1} \left(\lambda_\infty^{K_0} \sqrt{N} \sqrt{\exp(D/T_1)} + (\exp(D/T_1)) \right). \quad (3)$$

Using $\Psi(K_0, K_1)$ to denote expression (3), and recalling that we wish to minimize this upper bound, we formulate the two-phase problem as

$$\begin{aligned} \min_{K_0,K_1} \quad & \Psi(K_0, K_1) \\ \text{such that} \quad & K_0 + K_1 \leq K \\ & K_0 \geq 0, \quad K_1 \geq 0 \\ & K_0, K_1 \text{ integers.} \end{aligned} \tag{4}$$

Since expression (3) is decreasing in both K_0 and K_1, the inequality constraint $K_0 + K_1 \leq K$ will be binding in an optimal solution. Therefore, we can eliminate K_1 as $K_1 = K - K_0$ and rewrite (4) as

$$\min_{K \geq K_0 \geq 0} \quad \Phi(K_0), \quad K_0 \text{ integer}$$

where $\Phi(K_0) \equiv \Psi(K_0, K - K_0)$.

Lemma 3.1. *The continuous relaxation of the above two-phase optimization problem is convex.*

Proof. $\frac{d^2\Phi(K_0)}{dK_0^2}$ is equal to

$$\lambda_{T_1}^{K-K_0} \sqrt{\exp\left(\frac{D}{T_1}\right)} \Big[\lambda_\infty^{K_0} \log^2\left(\frac{\lambda_\infty}{\lambda_{T_1}}\right) \sqrt{N} + \sqrt{\exp\left(\frac{D}{T_1}\right)} \log^2 \lambda_{T_1} \Big],$$

which is clearly strictly positive. Thus, the objective function is convex. The feasible region $K \geq K_0 \geq 0$ is also convex. Hence the continuous relaxation is a convex optimization problem. □

The derivative of $\Phi(K_0)$ is zero at

$$\bar{K}_0 = \frac{1}{\log \lambda_\infty} \log \left(\frac{\sqrt{\exp(D/T_1)} \log \lambda_{T_1}}{\sqrt{N} \log\left(\frac{\lambda_\infty}{\lambda_{T_1}}\right)} \right).$$

Let $\tilde{K}_0 = \lceil \bar{K}_0 \rceil$ and $\hat{K}_0 = \lfloor \bar{K}_0 \rfloor$, where $\lceil x \rceil$ is the smallest integer no smaller than a real number x and $\lfloor x \rfloor$ is the biggest integer no larger than a real number x. Clearly, $\lceil x \rceil = \lfloor x \rfloor = x$ when x is an integer. Then, the optimal solution K_0^* is

$$\begin{aligned} K_0^* &= \text{argmin}\{\Phi(\tilde{K}_0), \Phi(\hat{K}_0)\} \text{ if } 0 \leq \bar{K}_0 \leq K \\ K_0^* &= 0 \text{ if } \bar{K}_0 < 0 \\ K_0^* &= K \text{ if } K < \bar{K}_0. \end{aligned}$$

3.2. Multi-phase dynamic programming formulation

We now provide a dynamic programming formulation for the multi-phase problem. Toward that end, for $i = 0, 1, \ldots, m$, we defined the state of the dynamic program as (ϵ_i, n), which indicates that the $\sqrt{\chi^2}$ value for the state distribution of $\mathcal{MC}_{T_i}$ and π_{T_i} is at most ϵ_i at the beginning of phase i and we have n out of the total K transitions remaining. Note that K_i is the decision variable in phase i. We first define the dynamic programming state transitions by noting that

$$\begin{aligned} \sqrt{\chi_0^2(i+1)} &= \sqrt{\sum_{y \in X} \frac{(Q_{T_{i+1}}^0(y) - \pi_{T_{i+1}}(y))^2}{\pi_{T_{i+1}}(y)}} \\ &= \sqrt{\sum_{y \in X} \frac{(Q_{T_i}^{K_i}(y) - \pi_{T_{i+1}}(y))^2}{\pi_{T_{i+1}}(y)}} \end{aligned}$$

Then using algebraic manipulation similar to phase one, the above right hand side is bounded as

$$\begin{aligned} &\leq \left| \left(\frac{(Q_{T_i}^{K_i}(\cdot) - \pi_{T_i}(\cdot))}{\sqrt{\pi_{T_i}(\cdot)}} \frac{\sqrt{\pi_{T_i}(\cdot)}}{\sqrt{\pi_{T_{i+1}}(\cdot)}} \right) \right|_2 \\ &+ \sqrt{\sum_{y \in X} \frac{(\pi_{T_i}(y) - \pi_{T_{i+1}}(y))^2}{\pi_{T_{i+1}}(y)}}. \end{aligned}$$

Thus letting $G_i = \left(\exp\left(D \left(\frac{T_i - T_{i+1}}{T_i T_{i+1}} \right) \right) \right)$ for brevity, $\sqrt{\chi_0^2(i+1)}$ is bounded above by $\sqrt{\chi_{K_i}^2(i)}\sqrt{G_i} + G_i$. Our definition of the dynamic programming state and the second inequality in Equation (1) imply that $\sqrt{\chi_{K_i}^2(i)} \leq \lambda_{T_i}^{K_i} \epsilon_i$. Hence $\sqrt{\chi_0^2(i+1)} \leq \lambda_{T_i}^{K_i} \epsilon_i \sqrt{G_i} + G_i$. Therefore, $\epsilon_{i+1} = \lambda_{T_i}^{K_i} \epsilon_i \sqrt{G_i} + G_i$. Also notice that if we have n out of K transitions remaining at the beginning of phase i and we choose to make K_i transitions in that phase, then only $n - K_i$ transitions remain. As a result, the dynamic programming state at the beginning of phase $i+1$ is $(\lambda_{T_i}^{K_i} \epsilon_i \sqrt{G_i} + G_i, n - K_i)$. The Optimal Value Function $V_i(\epsilon_i, n)$ is defined as the minimum value of an upper bound on $\sqrt{\chi_{K_m}^2(m)}$ obtained by the above approach, given that the dynamic program-

ming state at the beginning of phase i is (ϵ_i, n). Our aim is to find $V_0(\sqrt{N}, K)$ where $V_i(\epsilon_i, n)$ is

$$= \min\{V_{i+1}(\lambda_{T_i}^{K_i}\epsilon_i\sqrt{G_i} + G_i, n - K_i)\}, \quad (5)$$

for $i = 0, 1, \ldots, m-1$, all feasible ϵ_i and non-negative integers n that are at most K, where the minimum is taken over all integers K_i such that $0 \leq K_i \leq n$. The boundary condition is

$$V_m(\epsilon, n) = \lambda_{T_*}^n \epsilon, \ \ \forall\, \epsilon \in R_+, \ \ 0 \leq n \leq K. \quad (6)$$

Let $0 \leq n \leq K$ and ϵ_{m-1} be feasible. Then $V_{m-1}(\epsilon_{m-1}, n)$ is the minimum value of

$$V_m(\lambda_{T_{m-1}}^{K_{m-1}}\epsilon_{m-1}\sqrt{G_{m-1}} + G_{m-1}, n - K_{m-1}) \quad (7)$$
$$\text{such that } \ 0 \leq K_{m-1} \leq n, \ K_{m-1} \ \text{ integer.}$$

Observe that substituting boundary condition (6) the above objective function changes to

$$\lambda_{T_*}^{n-K_{m-1}} \left(\lambda_{T_{m-1}}^{K_{m-1}}\epsilon_{m-1}\sqrt{G_{m-1}} + G_{m-1}\right).$$

Thus, problem (7) is similar to the two-phase problem that can be analytically solved as in Section 3.1 to start the backward recursion. The computational overhead mainly depends on the number of phases and is not directly affected by the problem size or complexity of function evaluations.

In practice only upper bounds on SLEM λ_∞ and λ_{T_i} are typically known and have to be used in place of the actual SLEM values in all the formulations above. For example, canonical path and multi-commodity flow techniques from [11] can be used to bound λ_∞. Moreover, the following inequality, which is a straightforward consequence of a well-known result from [2] (Proposition 21.3 page 210) is often useful to bound SLEM λ_{T_i}: $\lambda_{T_i} \leq 1 - \left(\frac{\exp(-2D/T_i)}{2}\right)(1 - \lambda_\infty)$.

4. Application to the Ising Model

The Ising model is a canonical model of interacting particle systems in statistical mechanics. We consider the one-dimensional Ising model with periodic boundary conditions here. In particular, we have η "particles" arranged on a circle in the order $1, 2 \ldots, \eta$ so that each particle has two neighbors. The set of neighbors of particle i is denoted $N(i)$. The "spin" of particle i is denoted s_i, where $s_i \in \{+1, -1\}$. As a result, this system has $N = 2^\eta$ possible configurations. The "energy" $f(x)$ of configuration x is given by $-\frac{1}{2}\sum_{i=1}^{\eta}\sum_{j\in N(i)} s_i(x)s_j(x)$, where $s_i(x) \in \{+1, -1\}$ is the spin of particle i in configuration x and the factor $1/2$ is included to compensate for double counting. The depth D of this energy function is at most 2η.

In order to sample from the Boltzmann distribution over the set of N configurations of the above Ising model, we first define a candidate generator Markov chain $\mathcal{MC}_\infty$. The common natural choice in this case is the "single-site update" chain, which chooses one of the n particles randomly and flips its spin value. In order to make this Markov chain aperiodic, we use a holding probability of $1/2$ at each configuration. That is, with probability $1/2$, we do not change the configuration and with probability $1/2$ we change the configuration using the single-site update. Note that this chain is equivalent to the "lazy" nearest neighbor random walk on the η-dimensional binary cube $\{0, 1\}^\eta$, whose SLEM λ_∞ is $(1 - 1/\eta)$ [3]. The candidates proposed by this Markov chain can then be filtered using the Metropolis criterion at temperature T described in Section 1 to design Markov chain $\mathcal{MC}_T$ whose limiting distribution is Boltzmann (T).

Table 1 shows results of two-phase resource allocation problems for the Ising model. We only list K_0^* since $K_1^* = K - K_0^*$. As another example, we consider geometric cooling schedules of the form $\infty = T_0, T_1 = 5\eta^2, T_2 = T_* = 5\eta$ and use the dynamic programming formulation presented in Section 3.2 to find the optimal number of steps K_0^*, K_1^*, K_2^* when the total number of steps is $K = 2\eta^2$. Such geometric cooling schedules are often used in finite time investigations of Simulated Annealing [5]. Numerical results for different values of η are presented in Table 2.

Observe that our approach consistently achieved high performance guarantees whereas there was significant variability when all iterations were allotted to the last phase. Thus the

Table 1
Results for the two-phase Ising problem. The fifth column lists a lower bound on the probability of sampling exactly guaranteed by the allocation K_0^*, K_1^*. The sixth column lists this lower bound when all K iterations are allocated to the last phase. The seventh column lists the approximate percentage improvement, i.e., $\frac{100(p_*-p)}{p}$.

η	T_*	K	K_0^*	p_*	p	%
10	75	150	40	0.9651	0.7823	23
	50	200	43	0.973	0.7844	24
15	100	300	88	0.9817	0.5851	67
	50	500	97	0.9818	0.2093	369
20	200	400	147	0.9873	0.3433	187
	150	500	150	0.9948	0.6374	56
25	150	800	235	0.9973	0.1123	788
	75	2000	251	0.9999	0.8964	11
30	60	5000	370	0.9999	0.6588	51
	50	7000	380	0.9999	0.2419	313

Table 2
Results for the multi-phase Ising problem for $\infty = T_0, T_1 = 5\eta^2, T_2 = T_* = 5\eta$ and $K = 2\eta^2$. The last three columns are defined in Table 1.

η	K_0^*	K_1^*	K_2^*	p_*	p	%
10	35	11	154	0.9584	0.7844	22
15	77	18	355	0.9935	0.8741	14
20	137	25	638	0.9989	0.9255	8
25	214	32	1004	0.9998	0.9555	5
30	309	39	1452	0.9999	0.9733	3

benefit of resource allocation was most significant when the naive procedure performed poorly. The small performance benefits toward the end of Table 2 are not because resource allocation performed poorly improving a bad solution slightly but rather because the naive approach worked well in these cases. As a final note to indicate the relatively minor overhead in solving our resource allocation problems, the dynamic program arising from the last row of Table 1 was solved using GNU C++ in 0.019 seconds on a MacBook laptop running Mac OS X 10.4.11 with a 1.83 GHz Intel Core Duo processor and 1GB memory.

Acknowledgements

Research supported in part by the NSF under DMI-0244291. The first author appreciates summer support from the University of Washington.

REFERENCES

1. Bezakova, I., Stefankovic, D., Vazirani, V. V., and Vigoda, E., Accelerating simulated annealing for the permanent and combinatorial counting problems, *Proc. 17th ACM-SIAM symp. on discrete algorithms*, 900-907, 2006.
2. Behrends, E., Introduction to Markov chains, Vieweg-Verlag, Braunschweig, 2000.
3. Diaconis, P., and Stroock, D., Geometric bounds for eigenvalues of Markov chains, *Ann. App. Prob.*, 1(1),36-61, 1991.
4. Fill, J. A., Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, *Ann. App. Prob.*, 1(1),62-87, 1991.
5. Kalai A., and Vempala, S., Convex optimization by simulated annealing, Math. Operations Research, 31 (2), 253-266, (2006).
6. Kirkpatrick, S., Gelatt C. D. Jr., and Vecchi, M. P., Optimization by Simulated Annealing, *Science*, 220, 671-680, (1983).
7. Lovasz, L., and Vempala, S., Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithm, *Journal of Computer and System Sciences*, 72 (2), 392-417, 2006.
8. Martinelli, F., Relaxation times of Markov chains in statistical mechanics and combinatorial structures, Encyclopedia of Math. Sci., 110, Springer, 2003.
9. Mitzenmacher, M., and Upfal, E., Probability and Computing: Randomized Algorithms and Probabilistic Analysis, Cambridge University Press, 2005.
10. Nolte, A., and Schrader, R., A note on the finite time behavior of simulated annealing, *Math. of Operations Research*, 25 (3), 476-484, 2000.
11. Sinclair, A., Improved bounds for mixing rates of Markov chains and multicommodity flow, *Combinatorics, Probability & Computing*, 1,351-370, 1992.