



ELSEVIER

Computer Physics Communications 147 (2002) 64–70

Computer Physics
Communications

www.elsevier.com/locate/cpc

Go with the winners: a general Monte Carlo strategy

Peter Grassberger

John von Neumann, Institut für Computing, Forschungszentrum Jülich, D-52825 Germany

Abstract

We describe a general strategy for sampling configurations from a given distribution, *not* based on the standard Metropolis (Markov chain) strategy. It uses the fact that nontrivial problems in statistical physics are high dimensional and often close to Markovian. Therefore, configurations are built up in many, usually biased, steps. Due to the bias, each configuration carries its weight which changes at every step. If the bias is close to optimal, all weights are similar and importance sampling is perfect. If not, “population control” is applied by cloning/killing partial configurations with too high/low weight. This is done such that the final (weighted) distribution is unbiased. We apply this method (which is also closely related to diffusion type quantum Monte Carlo) to several problems of polymer statistics, reaction-diffusion models, sequence alignment, and percolation. © 2002 Elsevier Science B.V. All rights reserved.

PACS: 05.10.Ln; 36.20.Ey; 64.60.Ak

Keywords: Sequential Monte Carlo simulations with resampling; Pruned-enriched Rosenbluth method; Polymers; Percolation; Reaction-diffusion systems; Lattice animals; Sequence alignment

1. Introduction

Although Markov chain (Metropolis-type) Monte Carlo (MC) simulations dominate in statistical physics today, simulations not based on this strategy have been used from early times on. Well-known examples are evolutionary (in particular genetic) algorithms [1], diffusion type quantum MC [2], and several algorithms devised for the simulation of long chain molecules [3–7].

As these methods were developed independently in different communities, it was not generally recognized—or rather forgotten—that most of them are realizations of a common strategy, as pointed out by Aldous and Vazirani [8] who also coined the name “go with

the winners”. But essentially the same basic strategy was already discussed as a general purpose sampling method by Herman Kahn in 1956 [9] who called it “Russian Roulette and Splitting”, and attributed it to unpublished work by von Neumann and Ulam. For further applications of this strategy see [10–12]. The last two references also discuss applications in lattice spin systems and Bayesian inference, fields which will not be treated in the present review.

2. The basic strategy

2.1. Sequential importance sampling

As in any MC method, we draw configurations \mathbf{x} from some distribution $p(\mathbf{x})$. Writing the partition sum as

E-mail address: p.grassberger@fz-juelich.de (P. Grassberger).

$$Z = \sum_{\mathbf{x}} e^{-\beta H(\mathbf{x})} \approx M^{-1} \sum_{\alpha=1}^M e^{-\beta E(\mathbf{x}^\alpha)} / p(\mathbf{x}^\alpha), \quad (1)$$

we can interpret this as each configuration having its *weight* $W(\mathbf{x}) = e^{-\beta E(\mathbf{x})} / p(\mathbf{x})$. In importance sampling we try to choose $p(\mathbf{x}) \propto e^{-\beta E(\mathbf{x})}$, so that all weights become equal.

We now assume that we can break up the construction of a configuration into N single steps x_n , $n = 1, 2, \dots, N$. For a polymer, the n th step would, e.g., be the placement of the n th monomer. Then the weight W is obtained recursively as $W = W_N$, $W_0 = 1$,

$$W_n = W_{n-1} \frac{e^{-\beta(E(x_1, \dots, x_n) - E(x_1, \dots, x_{n-1}))}}{p_n(x_n | x_1 \dots x_{n-1})}. \quad (2)$$

In statistics this is called sequential importance sampling (SIS) [12].

In some cases, ‘natural’ values for the $p_n(x_n)$ are easy to guess. In the *Rosenbluth* method [3] for simulating a self-avoiding walk (SAW), e.g., one chooses uniformly among the free neighbors. But this is not optimal, a better choice is provided by *Markovian anticipation* [13]. In general, for choosing $p_n(x_n | x_1 \dots x_{n-1})$ one has to depend on heuristics, except in the case of diffusion quantum MC where perfect importance sampling ($W_N \equiv 1$) is possible if the ground state wave function is known [2,14]. Specific choices will be discussed together with applications.

2.2. Population control

The main drawback of SIS is that the distribution of weights can become extremely wide. If long range correlations are weak (as, e.g., for SAWs), $\log W_N$ is roughly a sum of independent terms. This suggests the following strategy:

If at any step n the weight W_n is above a suitably chosen threshold W_n^+ , we make an additional copy of the configuration x_1, \dots, x_n , and give both copies the weight $W_n/2$. Both are then grown independently (with eventual later copyings) up to full length.¹ In this way high weights are suppressed and precious ‘good’ configurations are less likely to be lost entirely by bad

subsequent moves. In [4] a similar strategy (but not based on weights) was called ‘enrichment’.

On the other hand, if W_n falls below another threshold W_n^- , we draw a random number $r \in [0, 1]$. If $r < 1/2$ we kill the configuration and start a new one. If $r > 1/2$ we keep it and double its weight.

Obviously, for any choice of the thresholds, neither the cloning nor the pruning introduce any additional bias. Thus we can, in principle, use any choice for W_n^+ and W_n^- , and we can change them ad libitum during the simulation. Bad choices will, however, lead to inefficiency, just as do bad choices for $p_n(x)$.

Except at very low temperatures where special care is needed [15,16] we found the following strategy to be sufficient:

- For the first configuration(s) we do not clone at all and kill only if the weight is exactly zero.
- If we have already m previous configurations which had reached size $\geq n$, we estimate from them the partition sum

$$Z_n \approx \hat{Z}_n \equiv m^{-1} \sum_{\alpha=1}^m W_n(\mathbf{x}^\alpha).$$

We then set $W_n^\pm = C^\pm \hat{Z}_n$ with $C^+ \approx 1/C^- \approx O(1)$ – $O(10)$.

2.3. Depth first versus breadth first

As described above, the algorithm is most efficiently implemented in a *depth first* fashion, and as such was called PERM (pruned-enriched Rosenbluth method) in [7]. In a depth first approach [17], we follow one copy until its end before we take up the other copy. In *breadth first* search, on the other hand, we treat all copies in parallel and handle the n th steps of all copies before we go to $n + 1$.

Evolutionary algorithms [1] are usually implemented breadth first. One puts up a population of M replicas which are evolved simultaneously, and population control is exercised such that M stays constant during the evolution. The same is true for most implementations of the “go with the winners” strategy. This has several advantages:

- Breadth first approaches are well adapted for massively parallel computers. One simply puts one configuration on each processor.

¹ In some cases (e.g., at low temperatures, where Boltzmann factors are huge), it might be necessary to make several copies and to distribute the weight evenly among them.

- One has no problem with keeping the number of replicas constant.
- One can use more general population control strategies [12].

But the last two points seem minor in most applications we have studied. On the other hand, the main advantage of depth first is the elegance and efficiency of the codes. The easiest implementation is by means of recursion (for a pseudocode see [7]). Copies (or rather instructions to make copies) are then put on a stack which is maintained automatically if recursive function calls are used. Storage use is minimized (as only a single copy and its history is kept in memory), and communications are also less than in breadth first.

3. Multiple spanning percolation clusters

Let us consider percolation on a large but finite rectangular lattice in $2 \leq d < 6$. We single out one direction as ‘spanning’. In this direction boundaries are open, while periodic b.c. are used in the other direction(s). For a long time it was believed that there is at most one spanning cluster (which touches both open boundaries) in the limit of large lattices, keeping the aspect ratio fixed ($L_i = x_i L$, $L \rightarrow \infty$, $i = 1, \dots, d$).

Since there is no spanning cluster for subcritical percolation and exactly one in supercritical, the only relevant case is critical percolation. There it is now known that the probabilities P_k to have exactly k spanning clusters are all non-zero in the limit $L \rightarrow \infty$. In $d = 2$ they are known exactly from conformal invariance, but for $d \geq 3$ no exact results are known. But there is a conjecture by Aizenman [18], stating that for a lattice of size $L \times \dots \times L \times (rL)$ (rL is the length in the spanning direction) $P_k \sim e^{-\alpha r}$ with

$$\alpha \propto k^{d/(d-1)} \quad \text{for } k \gg 1. \quad (3)$$

For $d = 2$ one has $\alpha \sim k^2$, in agreement with Eq. (3). A generalization in $d = 2$ consists in demanding that clusters are separated by at least q paths on the dual lattice [19]. In that case, and for periodic transverse b.c.,

$$\alpha = \frac{2\pi}{12} [((q+1)k)^2 - 1] \quad k \geq 2, \quad d = 2. \quad (4)$$

In order to test Eqs. (3), (4) for a wide range of values of k and r , one has to simulate events with




Fig. 1. Configuration of 5 spanning site percolation clusters on a lattice of size 500×900 . Any two clusters keep a distance of at least 2 lattice units. Lateral boundary conditions are periodic. The probability to find 5 such spanning clusters in a random disorder configuration is $\approx 10^{-92}$.

tiny probabilities, $\ln P_k \sim -10^2$ to -10^3 . It is thus not surprising that previous numerical studies have verified Eq. (4) only for small values of k , and have been unable to verify or disprove Eq. (3) [20,21].

To demonstrate that such rare events can be simulated with PERM, we show in Fig. 1 a lattice of size 500×900 with 5 spanning clusters which keep distances ≥ 2 . Eq. (4) predicts for it $P_k = \exp(-336\pi/5) \approx 10^{-92}$. This configuration was obtained by letting 5 clusters grow simultaneously, using a standard cluster growth algorithm [22], from the left border. Precautions were taken that they grew with the same speed towards the right, i.e. if one of them lagged behind, the growth of the others was stopped until the lagging cluster had caught up. If one of them died, or if two came closer than two lattice units, the entire configuration was discarded. If not, it was cloned if the weight W_n exceeded $3\hat{Z}_n$. Note that here the growth was made without bias, and therefore no pruning was necessary.

In this way we could check Eq. (4) with high precision, proving the correctness of our algorithm.

More interesting is the test of Eq. (3) for $d = 3$. Simulating up to 16 parallel clusters on lattices of sizes up to $128 \times 128 \times 2000$ (leading to probabilities as small as 10^{-300} !) gave perfect agreement with Eq. (3) [23].

4. Polymers

One of the main applications of the go-with-the-winners strategy is configurational statistics of long

polymer chains. For a breadth first algorithm which otherwise is very similar to PERM see [6].

4.1. Θ -polymers

PERM is particularly efficient near the so-called ‘theta-’ or coil-globule transition. According to the generally accepted scenario, the theta-point is tricritical with upper critical dimension $d_c = 3$ [24].

At T_θ , bias correction and Boltzmann factors nearly cancel in $d = 3$. Therefore, polymers have essentially random walk configurations with small (logarithmic) corrections. Therefore, a non-reversing random walk (U-turns are forbidden) for SIS is already sufficient to give good statistics with very few pruning and enrichment events. In [7] chains made of up to 1,000,000 steps could be sampled with high statistics within modest CPU time [7]. They were done in finite volumes (“dense limit”) and verified that the Θ -point indeed is a second order transition. The most precise verification of logarithmic corrections came from chains with $N = 10,000$ in infinite volume. The deviations from random walk behaviour turned out to be much stronger than the leading-log corrections predicted from the renormalization group [25], but agreement improves substantially when higher order corrections are included in the latter [26].

4.2. Critical unmixing

A related problem is the unmixing of semidilute polymer solutions. For any finite chain length N this is in the Ising universality class. But in addition to the Ising scaling laws, there are further universal scaling laws for parameters and amplitudes which, from the Ising point of view, would be non-universal. In particular, the critical temperature should approach T_θ when $N \rightarrow \infty$, $T_c - T_\theta \sim N^{-1/2}$, and the critical monomer concentration should tend to zero,

$$\phi_c \sim N^{-1/2}. \quad (5)$$

The exponents here are mean field, appropriate for $d = 3$. Indeed one should also expect logarithmic corrections [25]. Previous experiments had suggested an exponent 0.38 ± 0.01 in Eq. (5). This would be very hard to understand and has stirred a lot of theoretical activity (for a review see [27,29]). Simulations using PERM [27] showed that this is

wrong: The deviations from Eq. (5) can be understood most easily as logarithmic corrections.

4.3. DNA melting

DNA in physiological conditions forms a double helix. Changing the pH value or increasing T can break the hydrogen bonds between the base pairs, and a phase transition to an open coil occurs. Experiments suggest it to be first order [35]. While a second order transition would be easy to explain [36,37], no previous model had been able to give a first order transition.

The model studied in [28] lives on a simple cubic lattice. A double strand of DNA with length N is described by a diblock copolymer of length $2N$, made of N monomers of type A and N monomers of type B . All monomers have excluded volume interactions, i.e. two monomers cannot occupy the same lattice site, with one exception: The k th A -monomer and the k th B -monomer, with k being counted from the center where both strands are joined together, can occupy the same site. If they do so, then they even gain an energy $-\epsilon$. This models the binding of complementary bases.

The surprising result of simulations of chains with N up to 4000 is that the transition is first order, but shows finite scaling behaviour as expected for a second order transition with cross-over exponent $\phi = 1$. To demonstrate this, we show in Fig. 2 energy histograms for different chain lengths. One sees two maxima, one at $n = 0$ and the other at $n \approx N/2$, whose




Fig. 2. Histograms of the number of contacts, for single strand length $N = 500, \dots, 3000$, at $\epsilon = \epsilon_c$. On the horizontal axis is plotted n/N as is appropriate for a first order transition.

distance scales proportionally to N . But in contrast to usual first order transitions the minimum in between does not deepen with increasing N . This is due to the absence of any analogon to a surface tension. The same conclusion is obtained from specific heat data and from end-to-end distances [28].

In [28] we also studied similar models with (partially) switched off excluded volume effects. They show that excluded volume is the main force making the transition first order, as also confirmed by subsequent analytic calculations [38].

4.4. Native configurations of toy proteins

Predicting the native (\approx ground) states of proteins is one of the most challenging problems in mathematical biology [39]. It is difficult because of the many local energy minima.

In view of this, there exists a large literature on finding ground states of artificially constructed heteropolymers. Most of these models are formulated on a (square or simple cubic) lattice and use only few monomer types. The best known example is the HP model of Dill [40] which has two types of amino acids: hydrophobic (H) and hydrophilic (polar, P) ones. With most algorithms, one can find ground states typically for random chains of lengths up to ~ 50 .

In [16] we used PERM to study several sequences, of the HP model and of similar models, which had been discussed previously by other authors. In all cases we found the known lowest energy states, but in several cases we found new ones. A particularly impressive example is a chain of length 80 with two types of monomers, constructed such that it should fold into a bundle of four ‘helices’ with an energy -94 [41]. Even with a specially designed algorithm, the authors of [41] were not able to recover this state. With PERM we not only found it easily, we also found several lower states, the lowest one having energy -98 and a completely different structure.

4.5. Miscellaneous

Applications of PERM to other polymer problems are treated in [13,15,30–34]. For problems with open coils, a bias strategy called *Markovian anticipation* in [13] worked very well. Integrating over the disorder, we recently could also map a biased random walker

in the presence of random traps onto a stretched collapsed polymer [42]. Without bias, the transition from the finite time *Rosenstock* to the asymptotic *Donsker–Varadhan* behaviour is in 3D akin to a cross-over in a first order phase transition. With bias, the delocalization (globule-stretch) transition is first order in $d \geq 2$ [42].

5. Lattice animals (randomly branched polymers)

Consider the set of all connected clusters of n sites on a regular lattice, with the origin being one of these sites, and with a weight defined on each cluster. Lattice animals are defined by giving the same weight to each cluster. This distinguishes them from percolation clusters where the weight depends on the ‘wetting’ probability p . In the limit $p \rightarrow 0$ this difference disappears, and the two statistics coincide. It is believed that lattice animals are a good model for randomly branched polymers [43]. While there existed no efficient algorithm for estimating the animal partition sum there exist very simple and efficient Leath-type [22] algorithms for percolation clusters.

Our PERM strategy [14,44] consists in starting off to generate subcritical percolation clusters by a (breadth first) Leath method, re-weighting them as animals while they are still growing, and in making clones of ‘good’ ones. Since we work at $p < p_c$, we




Fig. 3. A typical lattice animal with 8000 sites on the square lattice.

do not need pruning. The threshold W_+ for cloning is chosen such that it depends both on the present animal weight and on the anticipated success for further growth.

In this way we obtained good statistics for animals of several thousand sites, independent of the dimension of the lattice. A typical 2D animal with 8000 sites is shown in Fig. 3. We also simulated animal collapse (when each nearest neighbor pair contributes $-\epsilon$ to the energy), and animals near an adsorbing surface [44].

6. Error estimates and reliability tests

Statistical errors can be estimated as usual by dividing a long run into several bunches, computing averages over each bunch, and studying the fluctuations between them. For PERM the situation is indeed rather easy, since each *tour* (set of all configurations generated by cloning from one common ancestor) is independent of any other.

To check for excessive fluctuations in weights W of entire tours, we make a histogram on a logarithmic scale, $P(\log(W))$, and compare it with the weighted histogram $WP(\log(W))$. If the latter has its maximum for values of $\log(W)$ where the former is already large (i.e. where the sampling is already sufficient), we are presumably on the safe side. However, if $WP(\log(W))$ has its maximum at or near the upper end of the sampled range, we should be skeptical.

In Fig. 4 we illustrate this with two figures taken from [32]. While the left panel gave rise to correct results, the right one did not.




Fig. 4. Full lines are histograms of logarithms of tour weights, normalized as tours per bin. Broken lines show the corresponding weighted distributions, normalized so as to have the same maximal heights. Weights W are only fixed up to a β -dependent multiplicative constant. While the left panel suggests a reliable simulation, the right one was indeed wrong (from Ref. [32]).

7. Conclusion

We have seen that MC simulations *not* following the Metropolis scheme can be very efficient. We have illustrated this with a wide range of problems. Conspicuously, the Ising model was not among them. It simply would be very hard to beat, say, the Swendsen–Wang algorithm. In principle, the go-with-the-winners strategy has as wide a range of applications as the Metropolis scheme. Its only requirement is that instances (configurations, histories, ...) are built up in small steps, and that the growth of their weights during the early steps of this build-up is not too misleading.

The method is not new. It has its roots in algorithms which have been regularly used for several decades. Some of them, like genetic algorithms, are familiar to most scientists, but it is in general not well appreciated that they can be made into a general purpose tool. And it seems even less appreciated how closely related are methods developed for quantum MC simulations, polymer simulations, and optimization methods. I firmly believe that this close relationship can be made use of in many more applications to come.

Among these are significance tests for sequence alignment, where one needs large samples of random pairs of sequences in order to check whether an observed alignment is significant. Instead of really drawing random pairs, one can use PERM to draw biased pairs which are more similar than random ones, enhancing thereby the interesting high-score region [45].

Another application is to epidemic models where one can follow the fate of epidemics which have a

very low chance of survival since, e.g., they started in a very hostile environment which they first have to adopt to their needs. Here simulations with PERM [46] allowed to verify with very high statistics the claim of [47] that no power laws result, in contrast to previous suggestions. A final application to a toy ‘population’ model [48] is discussed in [14].

References

- [1] H.P. Schwefel, *Evolution and Optimum Seeking*, Wiley, New York, 1995.
- [2] W. von der Linden, *Phys. Rep.* 220 (1992) 53.
- [3] M.N. Rosenbluth et al., *J. Chem. Phys.* 23 (1955) 356.
- [4] F.T. Wall et al., *J. Chem. Phys.* 30 (1959) 634, 637.
- [5] S. Redner et al., *J. Phys. A* 14 (1981) 2679.
- [6] T. Garell, H. Orland, *J. Phys. A* 23 (1990) L621.
- [7] P. Grassberger, *Phys. Rev. E* 56 (1997) 3682.
- [8] D. Aldous, U. Vazirani, “Go with the winners” algorithms, in: *Proc. 35th IEEE Sympos. on Foundations of Computer Science*, 1994.
- [9] H. Kahn, *Use of Different Monte Carlo Sampling Techniques*, in: H.A. Meyer (Ed.), *Symposium on the Monte Carlo Method*, Wiley, New York, 1956.
- [10] P. Grassberger et al., PERM: a Monte Carlo strategy for simulating polymers and other things, in: P. Grassberger et al. (Ed.), *Monte Carlo Approach to Biopolymers and Protein Folding*, World Scientific, Singapore, 1998.
- [11] Y. Iba, *cond-mat/0008226*, 2000.
- [12] J.S. Liu, *Monte Carlo Strategies in Scientific Computing*, Springer Series in Statistics, Springer, New York, 2001.
- [13] H. Frauenkron et al., *Phys. Rev. E* 59 (1999) R16.
- [14] P. Grassberger et al., *cond-mat/0010265*, 2000.
- [15] U. Bastolla et al., *J. Stat. Phys.* 89 (1997) 1061.
- [16] H. Frauenkron et al., *Phys. Rev. Lett.* 80 (1998) 3149; *Proteins* 32 (1998) 52.
- [17] R. Tarjan, *SIAM J. Comput.* 1 (1972) 146.
- [18] M. Aizenman, *Nucl. Phys. B* 485 (1997) 551.
- [19] M. Aizenman et al., *Phys. Rev. Lett.* 8 (2000).
- [20] P. Sen, *Int. J. Mod. Phys. C* 8 (1997) 229; *Int. J. Mod. Phys. C* 10 (1999) 747.
- [21] L. Shchur, *cond-mat/9906013 v2*, 1999.
- [22] P. Leath, *Phys. Rev. B* 14 (1976) 5046.
- [23] P. Grassberger, R. Ziff, to be published.
- [24] P.G. de Gennes, *Scaling Concepts in Polymer Physics*, Cornell Univ. Press, Ithaca, 1979.
- [25] B. Duplantier, *J. Phys. (France)* 43 (1982) 991; *J. Chem. Phys.* 86 (1987) 4233.
- [26] J. Hager et al., *Phys. Rev. E* 60 (1999) 2071.
- [27] H. Frauenkron et al., *J. Chem. Phys.* 107 (1997) 9599.
- [28] M.S. Causo et al., *Phys. Rev. E* 62 (2000) 3958.
- [29] B. Widom, *Physica A* 194 (1993) 532.
- [30] G.T. Barkema et al., *J. Stat. Phys.* 90 (1998) 1311.
- [31] S. Caracciolo et al., *J. Phys. A* 32 (1999) 2931.
- [32] P. Grassberger, *J. Chem. Phys.* 111 (1999) 440.
- [33] U. Bastolla et al., *J. Molec. Liquids* 84 (2000) 111; *cond-mat/0004169*, 2000.
- [34] T. Prellberg et al., *Phys. Rev. E* 62 (2000) 3780.
- [35] R.M. Wartell et al., *Phys. Rep.* 126 (1985) 67.
- [36] D. Poland et al., *J. Chem. Phys.* 45 (1966) 1456, 1464.
- [37] M.E. Fisher, *J. Chem. Phys.* 45 (1966) 1469.
- [38] Y. Kafri et al., *cond-mat/0007141*, 2000.
- [39] T.E. Creighton (Ed.), *Protein Folding*, Freeman, New York, 1992.
- [40] K.A. Dill, *Biochemistry* 24 (1985) 1501.
- [41] E. O’Toole et al., *J. Chem. Phys.* 97 (1992) 8644.
- [42] V. Mehra et al., *cond-mat/0107525*, 2001.
- [43] T.C. Lubensky et al., *Phys. Rev. Lett.* 41 (1978) 829.
- [44] P. Grassberger, W. Nadler, to be published.
- [45] R. Bundschuh, P. Grassberger, to be published.
- [46] P. Grassberger, to be published.
- [47] P. Grassberger et al., *Phys. Rev. E* 55 (1997) 2488.
- [48] S. Redner et al., *Amer. J. Phys.* 67 (1999) 1277.