Ref.TH.3521-CERN

MONTE CARLO METHODS IN THEORETICAL
HIGH ENERGY PHYSICS *)

B. Lautrup **)
CERN -- Geneva

ABSTRACT

The basic theory and the recent uses of the Monte Carlo method are presented for a non-physicist audience.

*) To be published in Communications of the ACM.

**) On leave of absence from the Niels Bohr Institute, Copenhagen, Denmark.

1. - INTRODUCTION

Since 1979 there has been an explosive increase in the use of Monte Carlo methods in lattice field theory¹⁾. By now most of the easily obtainable results have been extracted and further developments have come to face the limits on computational speed set by current hardware. A lot of the hope for further improvement rests on the expected progress in hardware design for general as well as special purpose computers²⁾. In parallel with this development theorists are exerting themselves to find new approaches to the Monte Carlo method. The main problem is caused by fermion fields which are a necessary part of any realistic model of nature. It is the aim of this article to introduce the basic theory of the Monte Carlo method and to describe some of the more recent applications.

Theoretical high energy physics is almost synonymous with relativistic quantum field theory. Models of nature consist of a multitude of interacting fields, some corresponding to observable particles (photons, electrons), others corresponding to unobservable particles (gluons, quarks) out of which other observable particles (pions, protons) are composed. Each field has one or more components at every space-time point. A relativistic quantum field theory may be cast into the form of a statistical system by the trick of making time imaginary instead of real. A number of physical properties, such as the masses of low-lying composite particles may be determined from the correlation functions of this statistical system.

Each state of the system is a complete four-dimensional field configuration. The number of degrees of freedom is proportional to the number of points in space-time. Because space-time is continuous the number of field configurations is immense, similar in nature to the number of possible weather configurations of the atmosphere. In order to deal with field theory in a mathematically consistent way the space-time continuum is approximated by a finite lattice. But even a small hypercubical lattice with ten sites in each of the four directions has

10 000 allowed space-time positions and thus more than 10 000 degrees of freedom. Gauge fields, like electromagnetic radiation, are polarizable in each of the four directions, which makes for 40 000 gauge field components on the above lattice. Furthermore, each gauge field component is an element of a continuous group. In the most interesting case of the gluon field this group is $SU(3)$ which needs at least eight real parameters for specifying one element. Thus, we end up with a system of 320 000 degrees of freedom on the above lattice. Statistical averages are integrals over all degrees of freedom, and it is clearly no mean feat to calculate an integral over a space of 320 000 dimensions.

It would be inconceivable to evaluate statistical averages for these systems by direct numerical techniques: luckily statistical systems lend themselves naturally to stochastic sampling techniques where, instead of averaging over all states, the average is only calculated over a sample of "particularly important" states. This sample is selected by (pseudo)stochastic "Monte Carlo" methods, which allow for an estimate of the statistical error too. Although these methods are well known³⁾ a number of special techniques have been and are being developed for lattice gauge theories. In Section 2 the basic principles of importance sampling are exposed. In Section 3 the applications in high energy theory are discussed, in particular the boson calculations and in Section 4 the fermion problem is discussed.

Fermions (electrons, quarks, protons) are necessary in any realistic model. A fermion field component is not an ordinary real or complex number, but a Grassmann number⁴⁾, and it is not possible to represent such a field directly in the computer. It is, however, possible to integrate out the fermion fields by hand at the expense of leaving a highly complicated integrand for the remaining bosonic integrals, involving the inverse and determinant of a huge matrix. Its dimension is at least as big as the number of space-time points in the lattice. A large amount of time and effort has already gone into developing techniques for efficient evaluation of the inverse of this matrix and its determinant. Monte Carlo methods seem particularly promising.

Finally it should be mentioned that a small finite lattice is at best a rather poor approximation to a space-time continuum. A lattice field theory behaves in some respects like a four-dimensional crystal and exhibits many features, such as phase transitions, known from ordinary crystals. In order to extract meaningful results it is necessary to extrapolate the results from finite lattices to the continuum limit. This rather subtle extrapolation is guided by the renormalization group^{5),6)} which governs the behaviour of field theories under transformations of the length scale. It is an interesting and surprising fact that even relatively small lattices seem to agree with the predictions of the renormalization group; this is the deeper reason for the upsurge of interest in these systems.

The Monte Carlo method has seen many applications in physics besides these more recent ones. In experimental physics Monte Carlo methods are used to generate events according to theoretical models. This will not be discussed here. In perturbative field theory Monte Carlo methods have been quite widely used to evaluate the integrals that arise from higher order Feynman diagrams. The most spectacular example is here the calculation of the 8th order anomalous magnetic moment of the electron, which involves a large number of ten-dimensional integrals⁷⁾. The methods used are variations on the theme presented in the following section and will not per se be discussed further here^{8),9)}.

2. - IMPORTANCE SAMPLING

In spite of the diversity of the applications of the Monte Carlo method, the basic issue boils down to the evaluation of one or more sums of the type

$$A = \sum_{s=1}^S F(s) \quad ; \quad (2.1)$$

even an integral may be thought of as an extreme limit of such a sum. For simplicity we shall limit ourselves to discussing this case. The "configuration space" S is a finite but extremely large collection of "states" $s = 1, 2, \dots, S$ and the function $F(s)$ is assumed to be real.

Since the sum is so large that direct evaluation is hopeless the best one can do is to generate a much smaller sample of states $s_1, s_2, \dots, s_t, \dots, s_T$ and try to estimate the value of the sum from the values $F(s_t)$. The index is called t because the sample states usually are generated at successive instants of time. In principle we must be able to continue the sample to $T \gg S$ but in practice we always have $T \ll S$, as we otherwise could have evaluated the sum by direct means. Since a realistic sample explores only a tiny part of the phase space certain smoothness assumptions are usually necessary about the behaviour of $F(s)$ in order that the sample may yield sufficient information to estimate A reliably.

Let $T(s)$ be the number of times the state s occurs in the sample. For $T \ll S$ every state presumably only occurs once or not at all, so $T(s)$ is 0 or 1. For $T \gg S$ most states occur many times. We shall assume that the sample is ergodic, which means that for sufficiently large T all states will occur in the sample, so that $T(s) > 0$. Then we can rewrite the original sum exactly in the form

$$A = \sum_{t=1}^T \frac{F(s_t)}{T(s_t)}$$

because each occurrence of any state s should only contribute a fraction $1/T(s)$ to the sum. We shall also assume that the sampling algorithm is constructed such that the relative frequency of occurrence of any state $T(s)/T$ converges towards a known distribution $\rho(s)$ for $T \rightarrow \infty$. Then the above formula indicates that the most natural guess we can make for the value of the sum when using a small sample is

$$\hat{A} = \frac{1}{T} \sum_{t=1}^T \frac{F(s_t)}{\rho(s_t)} \quad (2.2)$$

For $T \rightarrow \infty$ this estimator will converge towards the exact sum. Note that ergodicity implies that $\rho(s) > 0$ for all s . Without further knowledge about the nature of the sampling algorithm it is not possible to say anything about the error in this estimate. It is, however, intuitively clear that the error will be smallest if we choose a limiting distribution which is big in those regions of phase space where $|F(s)|$ is large. This is the principle of importance sampling.

The Monte Carlo method selects the particular sample for which the estimate is made, stochastically among all possible samples of the same length. This allows us not only to calculate the expected value of the estimator $\langle \hat{A} \rangle$ but also the expected error $\sigma = \sqrt{\langle (\hat{A} - A)^2 \rangle}$. It even permits us to estimate this error from the chosen sample itself. Although such an error estimate is useful, one should always bear in mind, that it is no better than the sample from which it is constructed. If the sample has missed an important part of configuration space, not only will the estimate for the sum be wrong, but the error estimate will not be representative of the actual deviation between the estimate and the original sum.

2.1 Independent Sampling

The most efficient stochastic sampling is obtained when a random generator for single states can be constructed with a probability distribution $\rho(s)$. Each state in the sample is generated independently of the rest of the sample, so that the probability of obtaining a particular sample is $\rho(s_1), \rho(s_2), \dots, \rho(s_T)$. The expected value of the estimator (2.2) when averaged over all samples of length T is $\langle \hat{A} \rangle = A$ for all T . Thus the estimator is unbiased. The expected error is

$$\sigma^2 = \frac{1}{T} (B - A^2) \quad (2.3)$$

where

$$B = \sum_{s=1}^S \frac{|F(s)|^2}{\rho(s)}$$

This quantity may itself be estimated from the sample

$$\hat{B} = \frac{1}{T} \sum_{t=1}^T \left| \frac{F(\lambda_t)}{\rho(\lambda_t)} \right|^2$$

The unbiased estimator for the error is

$$\hat{\sigma}^2 = \frac{1}{T-1} (\hat{B} - \hat{A}^2)$$

The error shows the convergence $1/\sqrt{T}$ typical of all Monte Carlo methods. Apart from making T as large as possible, the only influence one can have on the magnitude of the error lies in the choice of $\rho(s)$. Minimizing (2.3) with respect to $\rho(s)$ shows that the optimal distribution is proportional to $|F(s)|$ in agreement with the intuitive expectations. Even if we could generate a distribution of this shape we would, however, not be able to normalize it because that would entail the evaluation of a sum of the same kind as the original one. But even a rough approximation to the optimal distribution may yield quite good results. This is the basis for adaptive integration techniques for integrals in few dimensions^{8),9)}.

2.2 Markovian Sampling

When the desired "target" distribution $\rho(s)$ is very complicated it may be impossible in practice to construct a random generator which will generate each state independently. It may instead be feasible to generate the sample incrementally such that each state in the sample depends stochastically on its immediate predecessor. In this case the sample becomes a Markov chain.

Markov chains have very nice properties. Let $p(s, s')$ be the one-step probability for the next state being s' , when the present one is s . The basic problem of Markovian sampling is to tailor the one-step probability to yield exactly the right limiting distribution for the chain. A necessary condition is that the one-step probability must be in balance with the target distribution

$$\sum_{s=1}^S \rho(s) p(s, s') = \rho(s') \quad (2.4)$$

such that a single step will not change it. This also turns out to be nearly a sufficient condition. The only other requirement is that the chain must be ergodic: it must be able to reach every state in S from any initial state. In practical cases ergodicity is often self-evident.

The probability of generating a particular sample is $\rho_1(s_1)$, $p(s_1, s_2)$, ..., $p(s_{T-1}, s_T)$, where $\rho_1(s)$ is the probability distribution for the initial state. Summing over all chains while keeping $s_T = s$ we obtain the probability distribution $\rho_T(s)$ of the T^{th} state. Under the above-mentioned conditions of balance and ergodicity this distribution approaches the target distribution with an exponential rate. There exists a $\lambda > 0$ such that $|\rho_t(s) - \rho(s)| < e^{-\lambda t}$ for all t . The relaxation time for the Markov chain is $1/\lambda$ and it is only meaningful to use samples that are longer than the relaxation time. The estimator (2.2) is biased by the choice of initial state, but approaches A with $1/\lambda T$ corrections. Likewise the estimator for the error $\hat{\sigma}$ approaches the expected error σ with relative $1/\lambda T$ corrections.

The relaxation time depends mainly on the choice of one-step probability, whereas the error depends mainly on the choice of target distribution. Thus, there are two levels of adaptation possible. Firstly, among all the one-step probabilities in balance with the same target distribution, one may choose the one that yields fastest convergence. Since independent sampling (i.e., $p(s, s') = \rho(s')$) has zero relaxation time this is of course the best, but may not be obtainable. Secondly, the limiting distribution should be chosen to resemble $|F(s)|$ as much as possible. In statistical mechanics calculations it is natural to choose the target distribution to be the Boltzmann factor of the system, but it should be emphasized that this is neither the only possible choice, nor the optimal choice for all averages. A priori there seems to be no coupling between the relaxation time and the target distribution, because we might choose to sample independently. Due to other requirements, for example simplicity, a coupling is normally present, and the relaxation time will depend on the external parameters characterizing the system. The statistical error may also diverge at critical points causing critical slowdown of the Markov chain's convergence near such points.

The balance condition is usually implemented by means of detailed balance

$$\rho(s) p(s', s) = \rho(s') p(s, s') \quad (2.5)$$

which implies (2.4). Detailed balance may in turn be implemented in various ways, of which we shall only discuss the most useful one, the Metropolis algorithm¹⁰⁾. The basic feature of this algorithm is a Markovian random state generator q with symmetric one-step probability; $q(s, s') = q(s', s)$. Its target distribution is uniform in S . This generator is used to make a suggestion for the next state s' . If $s' \neq s$ and $\rho(s') > \rho(s)$ this suggestion is accepted. If $\rho(s') < \rho(s)$ a random number x is generated uniformly in the unit interval

and the suggestion is accepted if $p(s') > x p(s)$. Otherwise the suggestion is rejected and the present state s is used once more. Thus the probability is

$$p(s, s') = \begin{cases} q(s, s') & p(s') \geq p(s) \quad s' \neq s \\ q(s, s') \frac{p(s')}{p(s)} & p(s') < p(s) \\ q(s, s) + r(s) & s' = s \end{cases}$$

where $r(s)$ is the probability of rejection. It is easy to demonstrate that this algorithm leads to detailed balance independently of the choice of $q(s, s')$.

It is only via the choice of the suggesting algorithm q that we may influence the relaxation time. It is often possible to couple q to the observed performance of the sample and, for example, aim for 50% rejections. The Metropolis algorithm has the clear advantage that it is not necessary to normalize the target distribution in order to obtain convergence of the sample towards the correctly normalized distribution. In cases where only ratios of sums over the same space are desired, as is the case for statistical averages, this is a major simplification. On the other hand, it is then impossible to calculate the partition function which is the normalization constant for the Boltzmann factor. If it is desired to calculate the partition function one must choose a target distribution which is analytically normalizable. As long as it resembles the Boltzmann factor sufficiently the estimator will converge towards the desired sum.

3. - BOSON FIELDS

As mentioned in the Introduction a lattice gauge theory with bosonic fields is a collection of group elements distributed over the links between the sites of the lattice. Although the first studies of such systems used the heat bath algorithm¹¹⁾, most later studies have used the Metropolis algorithm. A single step in the Markov chain (an update) affects normally a single link variable. Due to the local nature of gauge field interactions, the probability density will be a product over simple geometric objects (plaquettes). The one-step probability only depends on the immediate vicinity of the link in question and has a relatively simple form. The Metropolis algorithm also has the important advantage over the heat bath technique, that it will automatically generate the correct Haar measure for the gauge group, provided the single steps in the group are done by group multiplication. Ergodicity is ensured either by choosing each link at random

or - more commonly - by sweeping through all the links of the lattice in a systematic way. In the latter case the updating procedure is explicitly time dependent and the analysis of the previous section does not directly apply. If all the link updates of a sweep are collected in a grand Markov step, time independence will ensue. But then detailed balance will be lost because the reverse transition involves the reverse sweep. Balance - and thereby convergence - is nevertheless assured.

The numerous Monte Carlo studies of bosonic lattice gauge theories¹²⁾ which have so far been carried out, have revealed that these theories possess a rich phase structure. Although it has no direct significance for the continuum physics, the phase structure nevertheless helps to reveal which non-perturbative mechanisms are at play. The confinement of quarks and gluons preventing these particles from being observed individually has been amply confirmed by lattice calculations. The continuum physics results which have been extracted using the renormalization group extrapolation technique comprise the tension in the glue string that connects the quarks and the masses of purely gluonic particles, glueballs. For a discussion of these results we refer to the literature¹²⁾.

4. - FERMION FIELDS

When fermions are included - as they must be - the difficulties are in some sense "squared". The fermionic degrees of freedom cannot be directly represented in a computer and must be removed before numeric evaluation is attempted. Luckily they occur in a very simple way which permits the explicit integration over these degrees of freedom. The remaining integral is only over bosonic variables, but the Boltzmann probability density has become non-local by being multiplied with the determinant of a huge matrix. The interesting physics quantities directly concerning fermions are calculated from the inverse of this matrix.

We are thus faced with the two related problems of evaluating the determinant and the inverse of a very large matrix. One simplifying feature is that this matrix is sparse: most of its elements are zero. The non-vanishing matrix elements only connect nearest neighbour sites of the lattice. A number of methods have been and are being tried out in order to solve this problem. At this moment the situation is still quite confused and no consensus has been reached concerning the best method. We shall now briefly review the most promising techniques. For a discussion of the physics results see Ref. 13).

4.1 The determinant

The simplest approach is to ignore it. This stop-gap technique has been dignified with the name "quenching". It is not clear to what extent quenching destroys the physics we are looking for.

Direct evaluation of the determinant is forbiddingly expensive involving at least $N^2 \log N$ operations where N is the dimension of the matrix.

The most promising method is incremental evaluation. Proceeding from a known simple state each bosonic variable is updated as described in the previous section and the associated change in the determinant is monitored using the formula

$$\det\{M + \delta M\} = \det M \det\{1 + M^{-1} \delta M\}$$

Since δM is even more sparse than M itself the last factor may be evaluated from the knowledge of only a few matrix elements of M^{-1} ¹⁴⁾.

4.2 The inverse

The calculation of the inverse cannot be ignored. Direct evaluation is again forbiddingly expensive. In practice one does not need to know the full inverse but only a few selected elements of it. One should therefore concentrate on methods of inversion that do not require storage of the complete inverse but only part of it.

The fermion matrix is of the form $M = 1 - R$ where R is "small". Standard methods of relaxation [Jacobi, Gauss-Seidel ¹⁵⁾] iterate the equation

$$M^{-1} = 1 + R M^{-1}$$

generating effectively the Neumann series

$$M^{-1} = 1 + R + R^2 + \dots$$

The iterations can be carried out columnwise so that only one column of the inverse need to be stored at any time. The problem with this method is that the process has to be cut off at some point and that all matrix elements are treated at equal footing even if they are very small and unimportant.

These problems seem to be overcome by a stochastic method originally due to van Neumann and Ulam ¹⁶⁾. The Neumann series may be written explicitly in the form

$$M^{-1}_{ij} = \delta_{ij} + \sum_{k=1}^{\infty} \sum_{k_1 \dots k_{k-1}} R_{ik} R_{k_1 k_2} \dots R_{k_{k-1} j}$$

This may be interpreted as a sum over paths $\pi = (i, k_1, k_2, \dots, k_{L-1}, j)$ of varying lengths L in the space of matrix indices. This sum may then be attacked by the standard methods of importance sampling discussed in Section 2. What is needed is therefore a random generator for paths of varying length. The simplest way to generate a path is by a random walk with step probability r_{ij} and probability s of stopping in each step. Thus the probability density of paths becomes

$$\rho(\pi) = r_{i,k_1} r_{k_1,k_2} \dots r_{k_{L-1},j} (1-s)^L$$

Notice that the paths are sampled independently by this method. In order to select the most important paths the natural choice of random walk probability is $r_{ij} \sim |R_{ij}|$. Since R only connects nearest neighbours the random walk proceeds through the lattice in small steps from site to site. There are many ways to improve this method, for example concerning the way the paths are stopped, or the best adaptation to the desired distribution, but we shall refrain from doing this here¹⁷⁾. The beauty of this method lies in the systematic use of stochastic methods in the bosonic as well as the fermionic part of the calculations, and that the running time is not proportional to the size of the lattice.

5. - CONCLUSIONS

The conventional Monte Carlo calculations as described above are now entering into the "industrial" era in which thousands of hours of computer time are being used to obtain a few physics results. Alternative Monte Carlo approaches are being investigated¹⁸⁾ but have not yet been shown to be competitive with conventional Monte Carlo in "production runs".

The initial excitement over the possibility that small lattices could be used to extract meaningful continuum physics results has now been replaced by a wary scepticism. Although the results from the purely bosonic calculations seem to be beautifully in agreement with the renormalization group expectations the fermionic calculations have barely progressed to the point where similar comparisons can be made. The nuclear particle masses that have been extracted in different Monte Carlo "experiments" seem not to agree very well with each other. Maybe much larger lattices are needed in the fermionic case because of the non-local nature of the interactions that these fermions induce among the boson fields. Larger lattices require larger memory and faster numerical computation. One may hope that the next generation of supercomputers will be adequate for settling this very important question.

ACKNOWLEDGEMENTS

The author would like to thank Dr. A. Kennedy for valuable comments and criticisms as well as the ITP, Santa Barbara where part of this article was written.

REFERENCES

- 1) For an introduction to the physics and the numerical methods, see for example C. DeTar, "Introduction to Lattice Gauge Theories and Numerical Methods", Lectures presented at the Arctic School of Physics (1982), Helsinki preprint HU-TFT-82-28 (1982).
- 2) See the article by Pearson, Richardson and Toussaint, this issue.
- 3) A review of the Monte Carlo method going back to the old testament has been given by
J.H. Halton, SIAM review 12 (1970) 1;
A more recent review is found in K. Binder, "Phase Transitions and Critical Phenomena", Eds. C. Domb and M.S. Green, Academic Press N.Y., Vol. 5B P. 1 (1976) and in F. James, Reps. on Progress in Physics 43 (1980) 73.
- 4) Grassmann numbers anticommute. The standard presentation of their properties is found in
F.A. Berezin, "The Method of Second Quantization", Pure and Applied Physics 24, Academic Press (N.Y., 1966).
- 5) A general presentation of the method of the renormalization group has been given by
Bambi Hu, Phys. Rep. 91 (1982) 234.
- 6) The typical application to lattice gauge theory may be found in
M. Creutz, Phys. Rev. D21 (1980) 2308.
- 7) T. Kinoshita and W.B. Lindquist, Phys. Rev. Lett. 47 (1981) 1573.
- 8) See, for example
G.P. Lepage, J. Comp. Phys. 27 (1978) 192.
- 9) B. Lautrup, Proceedings of the Second Colloquium in Advanced Computer Methods in Theoretical Physics, Marseille, 1971.
- 10) N. Metropolis, A.W. Rosenbluth, A.H. Teller and E. Teller, J. Chem. Phys. 21 (1953) 1087.
- 11) See Ref. 2).
- 12) M. Creutz, C. Rebbi and L. Jacobs, Phys. Rep. (to be published).
- 13) See Ref. 12) and
G. Martinelli, "Recent Results from Lattice QCD with Fermions", Frascati preprint (June 1982), LFN-82/42(P).
- 14) D.J. Scalapino and R.L. Sugar, Phys. Rev. Lett. 46 (1981) 519.
- 15) D. Weingarten and D. Petcher, Phys. Lett. 99B (1981) 333.

- 16) J. Kuti, Phys. Rev. Lett. 49 (1982) 183.
- 17) T. DeGrand, J. Dreitlein and D.J. Toms, Phys. Lett. 120B (1982) 391.
- 18) One can here mention the Langevin method of
G. Parisi, Nucl. Phys. B205 [FS5] (1982) 337 in which the statistical averages
are evaluated from "real time" averages of solutions to differential
equations with white noise sources.