NUMERICAL SIMULATIONS OF LATTICE QCD

Thesis by

Paul Ernest Stolorz

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

1987

(Submitted July 29, 1986)

Acknowledgements

I would like to thank my advisor, Geoffrey Fox, for the encouragement and patience he has shown me during the course of this work. My thanks also go to Rajan Gupta and Apoorva Patel for their help, support, and many useful discussions.

I am especially grateful to Steve Otto for his guidance, enthusiasm, and wisdom over the years, and also for his good humour and wit.

Finally, I would like to thank my parents for their astonishing forbearance and support, as well as those who have been closest to me during my stay at Caltech, especially Kathy, Mike, Steve, Betsy, Paul, John V., Jim, Dave, Mark M., Joel, Chi-Bin, Becky, John S., Lynn, Mary, Mark G., Joann, Nick, Lisa, Brian, Bette and James.

Abstract

Numerical methods are used to investigate some of the non-perturbative properties of lattice QCD. With the aid of Monte Carlo techniques based on the canonical ensemble, we calculate the QCD potential between a pair of heavy quarks, in the quenched approximation (no dynamical quarks). We find that the potential exhibits a linear dependence on distance at distances of the order of a fermi, which is consistent with the expected confining property of QCD. At smaller distances, we observe that the potential follows a 1/R type behaviour. We also compute the mass of the 0^{++} glueball for the SU(3) gauge group. We implement several statistical improvements in this calculation, in order to extract the mass reliably from the Monte Carlo simulations. We obtain a mass value of $\approx$1400 MeV for this glueball state (in the quenched approximation).

Finally, we use a numerical method, called the "demon" method, which is based upon the microcanonical ensemble, to measure the flow of lattice actions for the group SU(2) under renormalisation transformations generated by the Monte Carlo Renormalisation Group technique. We find that the demon method is ideally suited to the problem of tracking these renormalisation flows. Using the method, we are able to obtain an "improved" lattice action, which better describes physics near the continuum limit than the more straightforward naive actions.

Table of Contents

Introduction

Gauge field theories are by far the most successful attempts to date at explaining the behaviour of nature at high energies. The simplest quantised gauge therory, quantum electrodynamics (QED), which is based on the abelian group U(1), gives remarkably precise predictions for the behaviour of electrons and photons. We also have quantum chromodynamics (QCD), a gauge theory based on the more complex non-abelian group SU(3), which is believed to govern the behaviour of quarks inside hadrons. This theory has been successful, although in a less quantitative fashion, in explaining some of the features of high energy hadronic processes.

Unfortunately, most techniques of formulating gauge theories rely for their predictive power on perturbative expansions in some small dimensionless quantity, the coupling. In the case of QED, this is the fine structure constant α, which happily takes the truly small value 1/137 for the energy ranges of interest. In QCD we are not so lucky, however. The dimensionless coupling strength g is no longer constant, but changes, depending upon the energy scale at which one is looking at the process [2] (in truth, the coupling also changes for QED, but these changes are only significant at energies far larger than we can currently reach). This important discovery of a "running" coupling constant enables one to use perturbative techniques for QCD at high enough energy scales, when $g \ll 1$. Unfortunately, these energies correspond to distance scales much less than the size of a hadron. Therefore, in the case of QCD there is no hope of using perturbative methods to explain *all* the features of hadronic behaviour from the fundamental theory.

These problems can be circumvented by formulating the theories on a lattice [1]. When a gauge theory is formulated in this way, the observable quantities are all well-defined, and calculable, at *any* value of the coupling. Furthermore, the gauge symmetry can be preserved in a

natural way. The lattice, therefore, provides a powerful way of extracting non-perturbative information about a gauge theory.

Now, the lattice formulation turns out to be equivalent to a statistical system in thermal equilibrium, with the action of a lattice field configuration playing the part of the "energy" of a microstate of the system, and with the coupling g corresponding to the temperature of the entire ensemble. Observable quantities are obtained from ensemble averages in the limit of the lattice spacing going to zero. There has long been a substantial industry in statistical physics which studies the behaviour of large statistical systems using Monte Carlo methods. Many of these techniques are now being applied to the investigation of field theories in the lattice formulation, since they allow us to look at the previously inaccessible region of intermediate couplings.

Several important insights into QCD have been provided by this numerical work. The first [3], and crucial, one was the production of strong evidence that QCD is a confining theory, which means that infinite energy is required to pull quarks infinitely far apart. Confinement was indicated by the fact that the lattice formulation does not undergo a phase transition as one moves from the region of strong coupling where the lattice theory confines, to the region of weak coupling, which is where the continuum limit is taken and where perturbation theory is reliable. The result also holds for the simpler non-abelian group SU(2). The result is important because we know that the theory must at the same time account for the 2 facts that quarks are bound inside hadrons and that they seem to behave essentially as free particles at high energies. In contrast, lattice QED does undergo a phase transition, and we live in the deconfined phase (electrons may be pulled infinitely far apart).

Monte Carlo methods have also been used to attempt to calculate the spectrum of composite particles from the first principles of QCD, without requiring any extra parameters [4]. These have met with limited success, due partly to limited computational resources and partly to the fact that we do not yet understand how to incorporate all features of physical interest into the lattice. For instance, representing fermionic fields on the lattice is highly

problematical. Despite these difficulties, however, a certain amount of progress has also been made in understanding the important question of how the breakdown of chiral symmetry occurs in QCD [5]. There has also been substantial work done on the renormalisation group properties of QCD [6], in which one looks at how the various couplings parameterising the theory at one scale are related to the couplings parameterising it at another scale. This is an extemely important issue when one must deal with the approach to the continuum limit.

In Chapter 1 of this thesis, then, we present the basic principles of lattice gauge theories. Here, we also describe how Monte Carlo methods based upon the canonical ensemble are used to extract physical information from the lattice formulation. We also discuss the implementation of these ideas on a homogeneous parallel processor, the Hypercube, recently developed at Caltech.

In Chapter 2, we apply these methods to the computation of the QCD potential between a pair of static quarks. This quantity gives an important indication of confinement, namely a linear behaviour with distance at scales of the order of a fermi. In Chapter 3, we use the methods to compute, from first principles, the mass of the 0^{++} glueball, a state which is believed to consist essentially of interacting gluonic fields. In principle, our techniques could be used to calculate the mass of all the glueball states, but we have been limited by the computational demands of the problem. We have, therefore, concentrated here on improving the efficiency of the Monte Carlo methods for this problem.

Finally, in Chapter 4, we look at the Monte Carlo Renormalisation Group (MCRG), a renormalisation technique in which scale changes are performed directly on configurations generated by (canonical) Monte Carlo methods. We then describe a new approach to Monte Carlo simulations, which takes advantage of the fact that the canonical and microcanonical ensembles are equivalent in the thermodynamic limit. We discuss the various ways in which the microcanonical ensemble can be used to conduct Monte Carlo simulations. We then explain how one of these methods, the "demon" method, is ideally suited to the problem of tracking the behaviour of couplings in MCRG studies. In this way we are able to determine

the best way to approach the continuum limit in numerical simulations.

References

[1] K.G. Wilson, Phys. Rev. D10 (1974) 2445;

A.M. Polyakov, Phys. Lett. B59 (1975) 82;

F. Wegner, J. Math. Phys. 12 (1971) 2259.

[2] H.D. Politzer, Phys. Rev. Lett. 30 (1973) 1346;

D.J. Gross, F. Wilczek, Phys. Rev. Lett. 30 (1973) 1343.

[3] M. Creutz, Phys. Rev. Lett. 43 (1979) 553.

[4] H. Hamber, G. Parisi. Phys. Rev. Lett. 47 (1981) 1792;

H. Lipps et al., Phys. Lett. 126B (1983) 250;

D. Weingarten, Phys. Lett. 109B (1982) 57.

[5] J. Kogut et al., Phys. Rev. Lett. 48 (1982) 1140.

[6] K.G. Wilson, in: *Recent Developments in Gauge Theories* (Cargese, 1979), eds. G. 't Hooft et al. (Plenum, New York, 1980).

Chapter 1: Fundamentals of Lattice Gauge Theories

1. Gauge theories on a lattice

We will first describe the basic lattice formalism. Excellent reviews of the topic may be found in [1,2]. The fundamental objects in the lattice formulation of a gauge theory are the link variables U_{ij} that connect sites i and j of a regular Euclidian lattice (see Fig. 1.1). For a given gauge group G, these variables take the form of matrices representing elements of the group [3]. Thus,

$$U_{ij}=e^{igA_{\mu}^{\alpha}(x)\lambda_{\alpha}dx^{\mu}} \tag{1.1.1}$$

where g is the coupling constant, λ_{α} are the generators of the group (matrices), μ labels the direction in space-time of link ij, and where $A_{\mu}^{\alpha}(x)$ are the gauge fields (the α parameter is summed over, but not the μ parameter). The matrix U_{ij} determines the amount of rotation in an internal symmetry space as one moves from site i in the lattice to site j. In order for the return from site j to site i to result in a total effect of zero, it is natural to demand that

$$U_{ij}=U_{ji}^{-1} \tag{1.1.2}$$

Local gauge transformations are defined by applying arbitrary group matrices g at the sites i and j in the following way:

$$U_{ij}^{'}=g_i U_{ij} g_j^{-1} \tag{1.1.3}$$

Consider now the product of link matrices U_{ij} along some path in the lattice labelled by the sites 1 to N. Applying an independent gauge transformation at each site, the final gauge-transformed product will be given by

$$g_1 U_{12} U_{23} \ldots U_{N-1\,N} g_N^{-1} \tag{1.1.4}$$

since the g_i's and their inverses cancel at all the intermediate sites in the path. If we look at

only closed loops, so that site 1 and site N are the same site, the trace of the transformed product is

$$Trace\,(g_1 U_{12} U_{23} \cdots U_{N-1\,1} g_1^{-1}) = Trace\,(U_{12} U_{23} \cdots U_{N-1\,1}) \tag{1.1.5}$$

Hence, a locally gauge-invariant object is formed by taking the trace of a closed loop of links. This is called a Wilson loop. The ability to retain the gauge symmetry explicitly in this way is one of the main reasons for the importance of the lattice formulation, because the invariance of physical quantities under the internal rotations of fundamental fields is the distinctive feature of gauge theories. Other symmetries, such as those of space-time translation and rotation, are sacrificed by the formalism, but it is hoped that these symmetries can be restored in the continuum limit (and indeed, this restoration can be checked [4]).

Notice, in particular, that in constrast to the continuum case, the lattice formulation has the advantage of not requiring a gauge-fixing term in the action. This term is needed in the continuum formulations because the naive path integral contains equal contributions from many gauge-equivalent field configurations, which leads to infinities in the resulting expressions for Green's functions. The gauge-fixing term removes these divergences by counting a much smaller subset of configurations in each gauge-equivalent class. In the lattice case, the absence of such a term means that here also, many gauge-equivalent configurations contribute equally to the path integral. However, this does not present a problem, since for a finite lattice, and for the compact gauge groups that are usually considered, the total of these contributions is finite, and their effect can be divided out in the measure when looking at physical quantities.

We will restrict our attention to the gauge groups SU(N), since we are mainly looking at QCD, a theory based upon the SU(3) group. For SU(N) gauge groups, the inverse of a link matrix is its hermitian conjugate. The simplest lattice gauge theory for SU(N) is *defined* by the following partition function[3]:

$$Z = \int \prod_{links\ l} dU_l\, e^{-\beta S(U)} \tag{1.1.6}$$

where

$$S(U)=\sum_{p}[1-\frac{1}{2N}Tr(U_p+U_p^{\dagger})] \tag{1.1.7}$$

This particular form for the action function $S(U)$ is known as the Wilson action. The U_p represent the elementary squares ("plaquettes") of the lattice, and consist of the product of the 4 (directed) link matrices forming the sides of the plaquette (see Fig. 1.1). The inverse temperature β is $\frac{2N}{g^2}$, where g is the bare coupling constant. By performing a Taylor expansion about the plaquette centers, Z can be written (for small lattice spacing a) in the form

$$Z=\int d[A]e^{-1/4\int d^4x F_{\mu\nu}^{\alpha}F_{\alpha}^{\mu\nu}} \tag{1.1.8}$$

where

$$F_{\mu\nu}^{\alpha}=\partial_\mu A_\nu^{\alpha}-\partial_\nu A_\mu^{\alpha}+igf^{\alpha\beta\gamma}A_\mu^{\beta}A_\nu^{\gamma} \tag{1.1.9}$$

The $f^{\alpha\beta\gamma}$ are the structure constants of SU(N). We have now connected our formalism with the real world because this object is known to be the Euclidian form of the path integral defining the pure SU(N) gauge theory in the continuum! Provided that $1/a$ is much larger than any momentum present in the problem, there is hope the lattice will give sensible information about the theory described by the Lagrangian in Eq. (1.1.8).

The Wilson action is not the only form that reduces to the continuum limit of Eq. (1.1.8). It is, however, the simplest. It is believed that there is a large class of actions reducing to Eq. (1.1.8), which are "equivalent" to the Wilson action in the sense that physical observables are insensitive to the exact choice of action in the limit $a\rightarrow 0$. Of course, in the region where a is large, markedly different behaviours may be observed by the different actions.

Matter fields are incorporated in the action by defining the quark wave functions at lattice sites. However, because the quarks are fermions, certain difficulties arise when attempting to take the continuum limit. In fact, it has been shown that it is not possible to obtain a

lattice formulation of fermions without either introducing extra species of fermions that do not exist in the continuum theory, or explicitly breaking the chiral symmetry that exists in the massless fermion limit of the continuum theory [5].

The first approach was taken by Susskind [6], who devised a formulation in which chiral symmetry is retained in the limit of massless fermions at the expense of introducing extra fermion species. An alternative technique, due to Wilson [7], consists of introducing a term which explicitly breaks chiral symmetry, but allows the correct number of fermion fields to be described. Denoting by D the lattice Dirac operator for a given formalism, the full path integral is now

$$Z=\int\prod_l dU_l\prod_i d\bar{\psi}\prod_i d\psi e^{-\beta S(U)-\bar{\psi}(D+m)\psi} \tag{1.1.10}$$

where ψ and $\bar{\psi}$ are the fermion fields represented by anticommuting elements of a Grassman algebra. In practice, the best way to evaluate this integral is to integrate out the fermion fields to leave [8]

$$Z=\int\prod_l dU_l e^{-\beta S(U)}\,Det\,(D+m) \tag{1.1.11}$$

We are now presented with a further difficulty. The effective action in Eq. (1.1.11) is highly non-local. It turns out that this is a severe problem when Monte Carlo simulations are conducted (see Section 3), since the computer time required is an order of magnitude more than the time needed when the fermionic part of the action is not included. Many of the important aspects of QCD, such as the question of whether or not it confines, or the behaviour of states composed primarily of gluons, are believed to be influenced *mainly* by the self-interacting nature of the gluonic fields. This is well-described in the gauge part of the action S(U) above. Therefore, because of the limitations of computer power, the work in this thesis will deal with the properties of the pure gauge action without dynamical fermions.

2. The continuum limit

To obtain meaningful physical predictions, one must be able to calculate lattice observables as the lattice spacing a is allowed to approach zero. Let's see briefly what happens when this naive continuum limit is taken. It is easy to show [1] that physical quantities such as the masses of particle states occur in the lattice formulation as the rate of fall-off of 2-point correlation functions between zero-momentum operators O:

$$<O(T)O(0)> \rightarrow e^{-mT} \quad as \ T \rightarrow \infty \tag{1.2.1}$$

where the two operators are separated by the distance T in the time direction, and where the average <..> of some operator A is defined by

$$<A> = \frac{1}{Z} \int \prod_l dU_l A(U) e^{-\beta S(U)} \tag{1.2.2}$$

Now write T as ta, where t indicates the number of lattice spacings (a dimensionless quantity). If the mass is to remain finite when the lattice spacing is taken to zero, as befits a physical quantity, then the dimensionless quantity ma, which is the object calculable on the lattice, must approach zero also. But the regularised lattice formulation defines a statistical system (albeit one in 4 dimensions), and ma corresponds to the inverse correlation length of this system (measured in lattice units). Therefore the correlation length must become infinite as the limit $a \rightarrow 0$ is taken: in other words, the continuum limit corresponds to the critical point of a statistical system.

Looking at the lattice formulation as a statistical physics problem, it has been found to be easily soluble in the limit of strong coupling $g \rightarrow \infty$. Standard diagrammatic expansions have been developed to calculate correlation functions in this region. The trouble is that in the case of QCD, and indeed for all pure gauge SU(N) theories, the continuum limit $a \rightarrow 0$ occurs as the coupling g is made to approach zero: that is, the critical point is at $g=0$. Purely analytical approaches are therefore limited in the range of couplings that can be tackled.

It is here that Monte Carlo methods begin to make their contribution. Simulations have shown that the statistical system defined by the lattice QCD action does not have a phase transition as one moves from the region of strong to weak coupling [9]. This means that the lattice theory can at the same time describe the physics that occurs at 2 greatly differing energy scales.

3. The Monte Carlo method

In general, the observable quantities on a lattice take the form of averages of gauge-invariant loop operators over all field configurations C, with each configuration being given a weighting $e^{-\beta S(C)}$. Thus, we want to calculate

$$<O>=\frac{1}{Z}\int\prod_{l}dU_{l}O(C)e^{-\beta S(C)} \tag{1.3.1}$$

where

$$Z=\int\prod_{l}dU_{l}e^{-\beta S(C)} \tag{1.3.2}$$

As mentioned above, these integrals can be attacked analytically at low β, by using series expansion techniques developed to tackle similar problems in statistical physics. However, these methods fail at intermediate values of β. We therefore turn to numerical methods to approximate the integrals.

It would be possible, in principle, to generate a finite sequence of configurations evenly over the entire space of configurations, and then calculate the ensemble average of the quantity $O(C)e^{-\beta S(C)}$. However, this is extremely inefficient, since for almost every configuration, $e^{-\beta S(C)}$ is essentially zero and makes no contribution to the ensemble average. The solution is to use the well-known method of importance sampling [10], in which the factor $e^{-\beta S(C)}$ is included in the probability measure with which configurations are generated, so almost every configuration makes a non-zero contribution to the integral. It is always possible to include

this exponential factor in the probability measure because it is a real and positive quantity. Thus, one generates an ensemble of N configurations, each configuration C_i occurring in the ensemble with probability

$$p(C_i) \propto \prod dU e^{-\beta S(C_i)} \tag{1.3.3}$$

The estimate of the observable is then given by the ensemble average

$$<O> = \frac{1}{N} \sum_{i=1}^{N} O(C_i) \tag{1.3.4}$$

The ensemble of lattice configurations is generated by a Markov process. This consists of stochastically generating a sequence of configurations. Given a certain configuration C_i, one creates a new configuration C_j with a transition probability W which depends *only* upon C_i, and on no other configurations in the sequence. One chooses the form of the transition probability W so that the configurations in this sequence, called a Markov chain, will asymptotically be distributed with the probabililty distribution Eq. (1.3.3).

It can be shown [11] that a guaranteed way of producing a Markov chain which satisfies the property Eq. (1.3.3), independently of the starting configuration, is to choose a transition probability W which satisfies the condition of *detailed balance*, namely

$$\frac{W(C_i, C_j)}{W(C_j, C_i)} = \frac{e^{-\beta S(C_j)}}{e^{-\beta S(C_i)}} \tag{1.3.4}$$

This is not the only way of satisfying Eq. (1.3.3), but it is by far the most convenient in practice.

Actually, we choose configurations which differ by just one link value, applying a transition probability $W(U_{ij}, U'_{ij})$ to each link in turn. The locality of the gauge action then guarantees that this W depends only on the link values in the plaquettes containing the link ij. Again, we note in passing that the inclusion of dynamical fermions leads to a non-local factor in the action which destroys this useful computational feature of the pure gauge action.

The question now arises, "What form should the transition probability W take ?" There are basically 2 methods available. The earliest of these [12], referred to as the Metropolis scheme after one of the original authors of the method, selects a trial link variable $U_{ij}^{'}$ which is close (in group space) to the old link variable U_{ij}. This trial value is then accepted or rejected as the new link value, depending on the difference between the new and old actions. The exact acceptance probability is chosen so that Eq. (1.3.4) is satisfied. The main difficulty with this scheme is that, in order to achieve a sufficiently high acceptance ratio for trail link values, new link values must always be very close to the old ones in group space, which means that the lattice configurations are highly correlated. The lattice must therefore be "swept" by the algorithm many times before one can be confident of having adequately traversed the configuration space of the problem.

An alternative technique, the heat bath algorithm [13], consists of using a transition probability which does not depend at all on the initial link value U_{ij}, but only on the *environment* of links surrounding it (see Fig. 1.2), and on the new value $U_{ij}^{'}$ to be generated. The form of the transition probability is

$$dW(U_{ij} \rightarrow U_{ij}^{'}) = dU_{ij}^{'} e^{-S(U_{ij}^{'} E_{ij})} \tag{1.3.5}$$

where E_{ij} is the sum of the 3-link products of U's around the 6 plaquettes containing U_{ij}. Notice that since the group measure dU_{ij} is invariant, this transition probability trivially satisfies the detailed balance condition Eq. (1.3.4).

The advantage of the heat bath scheme over the Metropolis algorithm is that the $U_{ij}^{'}$ do not depend on the old U_{ij}, and may in fact be in any part of the group manifiold. Hence, the correlations between configurations are far less than in the Metropolis method. Also, for the particular group SU(2), an extremely elegant and efficient heat bath algorithm can be used [13]. For SU(3) the method is not quite so sleek, but here also a clever technique, called the pseudo heat bath [14], has been invented, which retains some of the advantages of the SU(2) heat bath. We have used the pseudo heat bath to generate lattice configurations efficiently, so

we will now describe the basic principles of the method.

4. Heat baths and pseudo heat baths

Let's look first at how the genuine heat bath algorithm works for SU(2) [13]. Consider the link ij shown in Fig. 1.2, where the sum of the partial plaquettes containing the link is denoted by E. We want to generate a new link value U from some old link value x with transition probability

$$dW(x,U)=\frac{1}{Z}e^{\frac{1}{2}\beta \mathrm{Re}\,\mathrm{Tr}(UE)}dU \tag{1.4.1}$$

We parameterise an SU(2) element by

$$\begin{bmatrix} a_o+ia_1 & -a_2+ia_3 \\ a_2+ia_3 & a_0-ia_1 \end{bmatrix} = a_0\mathbf{1}+i\,\mathbf{a}\cdot\boldsymbol{\sigma} \tag{1.4.2}$$

where σ_i are the Pauli matrices, and where the a_i are real and satisfy

$$\sum_{i=0}^{3} a_i^2=1 \tag{1.4.3}$$

The invariant group measure is then given by

$$dU=\frac{1}{2\pi^2}\delta(\sum_{i=0}^{3} a_i^2 - 1)\prod_{i=0}^{3} da_i \tag{1.4.4}$$

Th useful property of SU(2) is the fact the environment E, which is the sum of 6 matrix products, can be expressed as

$$E=kV \tag{1.4.5}$$

where

$$k=(detE)^{1/2} \tag{1.4.6}$$

and where V is an element of SU(2). Hence,

$$dW(x,UV^{\dagger})=\frac{1}{Z}e^{\frac{1}{2}\beta k\,\mathrm{Re}\,TrU}\,dU \tag{1.4.7}$$

since the group measure is invariant. Now when U is represented by the form Eq. (1.4.2), TrU is just $2a_0$. Hence, the right hand side of Eq. (1.4.7) effectively reduces to computing a_0 with essentially an exponential distribution (an extra factor of $\sqrt{1-a_0^2}$ actually arises because of the δ-function in the measure). One then generates the remaining a_i randomly on the resulting 3-sphere with length $\sqrt{1-a_0^2}$. The new link value is then given by the product $UV^{\dagger}$.

Unfortunately, in the case of SU(3), the elegant transformation which gives Eq. (1.4.7) is not available, because now the environment E cannot be expressed as the multiple of an SU(3) element. Hence, $\mathrm{Re}Tr(UE)$ must be written out in a cumbersome way in terms of the 8 parameters of SU(3) to construct a direct heat bath [15].

These problems can be circumvented, however, by using the SU(2) subgroups of SU(3), a fact that was realised by Cabibbo and Marinari [14]. It suffices to use the following two subgroups:

$$S_1=\begin{bmatrix}\alpha_1 & 0\\ 0 & 1\end{bmatrix}\quad S_2=\begin{bmatrix}1 & 0\\ 0 & \alpha_2\end{bmatrix} \tag{1.4.8}$$

where α_1 and α_2 are 2x2 matrices of the form Eq. (1.4.2) representing elements of SU(2). In fact, we also use the subgroup S_3 which consists of the 4 matrix elements of SU(2) residing at the 4 corners of the SU(3) matrix, with a 1 in the central location and all the rest of the entries 0.

The idea is that, given an old link value U, a new one U' is given by

$$U'=S_3S_2S_1U \tag{1.4.9}$$

The 3 subgroups are applied in turn in order to generate this new element. At each stage the cumulative product $\prod_{i}^{n-1} S_iU$ forms the old link value, to which the S_n subgroup is applied with

transition probability

$$dW(\prod^{n-1} S_i U, S_n \prod^{n-1} S_i U) = \frac{1}{Z_{\alpha_n}} e^{\frac{1}{3}\beta \mathrm{Re}\, Tr(S_n \prod^{n-1} S_i UE)} d\alpha_n \tag{1.4.10}$$

In order to be useful, this approach must have the following characteristics. Firstly, we must be able to apply the spirit of the SU(2) heat bath algorithm to generate an expression similar to Eq. (1.4.7). Secondly, we must use enough SU(2) subgroups so that there are no ideals; that is, no proper subgroups of SU(3) that are left invariant by the application of the SU(2) subgroups. This ensures that the configuration space is evenly scanned. In fact, S_1 and S_2 are enough to guarantee this. Thirdly, we must be able to show that the link values U' generated in this way are actually distributed according to the correct Boltzmann distribution. This is shown to be the case in [14].

To see how the algorithm works, consider the application of the S_1 subgroup to an old link value U. Observe that the exponent in Eq. (1.4.10) can be written as

$$\mathrm{Re}Tr(S_1 UE) = \mathrm{Re}Tr(\alpha_1 W_1) + \textit{terms independent of } \alpha_1 \tag{1.4.11}$$

where W_1 is the 2x2 submatrix forming the top left-hand block of the product UE. The SU(2) element α_1 is parameterised as usual by the form Eq. (1.4.2). Now W_1, a general 2x2 matrix, can be written as

$$W_1 = c_0 \mathbf{1} + i\, \mathbf{c}\cdot\boldsymbol{\sigma} \tag{1.4.12}$$

where the c_i are arbitrary complex numbers. It is then easy to show, using the properties of Pauli matrices, that

$$\mathrm{Re}Tr(\alpha_1 W_1) = \mathrm{Re}Tr(\alpha_1 F) \tag{1.4.13}$$

where

$$F = f_0 \mathbf{1} + i\, \mathbf{f}\cdot\boldsymbol{\sigma} \tag{1.4.14}$$

with $f_i = \mathrm{Re}(c_i)$. Since the f_i are all real, the matrix F is, in fact, a multiple of an SU(2) element. By extracting k, the square root of the determinant, to leave the SU(2) element G, we see that the transition probability can be written

$$dW(U,S_1U)=\frac{1}{Z_{\alpha_1}}e^{\frac{1}{3}\beta k\,Tr(\alpha_1 G)}d\alpha_1 \tag{1.4.15}$$

We have now essentially recovered the SU(2) heat bath. Upon multiplying α_1 on the right by $G^\dagger$, and using the invariance of the SU(2) measure, we get

$$dW(U,T_1U)=\frac{1}{Z_{\alpha_1}}e^{\frac{1}{3}\beta k\,Tr\,\alpha_1}d\alpha_1 \tag{1.4.16}$$

where

$$T_1=\begin{bmatrix}\alpha_1 G^\dagger & 0\\ 0 & 1\end{bmatrix} \tag{1.4.17}$$

The point now is that the right-hand side of Eq. (1.4.17) is easy to generate. It is analagous to the elegant expression Eq. (1.4.7) derived in the SU(2) case. Having obtained α_1 in this way, the new link value is gotten by forming the product

$$U^1=T_1U \tag{1.4.18}$$

The whole procedure is then repeated with the second subgroup being applied to the "old" link value U^1, and the resulting link U^2 is then used as the starting point for the application of the third subgroup.

The algorithm works because when an SU(2) subgroup is multiplied by the SU(3) product UE, the *relevant* part of the product UE (that is, the subgroup-dependent part contributing to the real part of the resulting trace) can be written as the multiple of an SU(2) element. However, the procedure is not *truly* a heat bath because the old link value U appears in the product UE. This cannot be avoided, because the top left 2x2 sub-matrix of the product S_1U

is not in general an SU(2) element, which is a crucial requirement if one wants to remove U from the environment W_1. The method is therefore not quite as elegant as the genuine SU(2) case, and it is also less efficient because the new link value is correlated with the old one via the presence of the old link value in the environment. Nevertheless, the spirit of the SU(2) improvement is retained, and the method is much more efficient than naive algorithms using the full SU(3) measure.

5. Lattice gauge theories and parallel processing

Clearly, useful Monte Carlo simulations must be capable of producing results on a lattice whose grid size is somewhat smaller than the length scale we are interested in. At the same time, the volume of the lattice must be large enough that the finite, overall, physical size does not interfere with the physics at this length scale. However, since we are dealing with a 4-dimensional system, a limited increase in the number of lattice sites in each dimension greatly increases the number of degrees of freedom in the problem. For example, in our calculation lattices of size $12^3 \times 16$ were used, which required roughly 10^6 degrees of freedom. The ensuing computation places severe demands on both the memory and CPU time of conventional computers.

In order to make our calculations feasible, we have implemented them on an ensemble machine, the Hypercube, which is radically different from computers with the traditional von Neumann architecture. This machine, recently developed at Caltech, basically consists of 64 homogeneous processing elements, each of which is an independent computer with its own memory and roughly the power of a personal computer. The elements, or *nodes*, are wired together in the topology of a Boolean hypercube in 6 dimensions, so that each node is connected to 6 other nodes, its "nearest neighbours." A calculation proceeds by loading separate code into each of the nodes, and then having each node work independently on its own part of the problem. In this way the whole computer has roughly the power of 64 personal computers. We say roughly, because from time to time it is necessary for each of the nodes to

communicate information to its nearest neighbours. However, in the case of our problem the overhead for this is swamped by the amount of computation that needs to be done between communications, so the efficiency of the machine is extremely high. In fact, this class of machines has been found to appropriate for many problems of physical interest. Details of the machine's architecture, and the scientific applications that have been developed on it, can be found in [16].

As far as our particular problem goes, we mention briefly that it is ideally suited to the Hypercube. This is because the problem is a regular crystalline one, in which the degrees of freedom remain at fixed positions in the lattice. A sequential computer moves in a regular way through the lattice, performing calculations at each position in turn. On the Hypercube, each node can be given responsibility for a small subsection of the lattice. The code in every node moves in a regular way through its particular lattice subsection, and when the boundary of a subsection is reached, the information about the Wilson loops that include this boundary is passed onto the next node, as shown in Fig. 1.3.

Now, since each node is stepping regularly through its part of the lattice, at each point in the execution of the code one knows exactly which neighbouring node a given node must communicate with next. This is important because it means that the communication commands can be be included in the code that is loaded into each node. The whole calculation is kept synchronised by requiring each node to send and receive its Wilson loop information in consecutive steps. This software-driven communication scheme is very straightforward to implement, and results in our case in efficiencies of about 97%. Irregular problems, in contrast, require an interrupt-driven communication system because one does not always know beforehand which node must be spoken to next.

It turns out then, that the code required to perform lattice calculations on the Hypercube can be easily obtained by relatively minor modifications of well-modularised sequential code. The elegant way in which Wilson loops lend themselves to this process is described in detail in [17]. The nodes of the Hypercube accept code that has been compiled from the high-level

language C, so only modest investments of time are needed to decompose sensible sequential algorithms onto the concurrent processor. In this way we have been able to substantially boost the computer power available to us, enabling us to perform high statistics calculations of the important physical observables on large lattices.

References

[1] J. Kogut, Rev. Mod. Phys. 51 (1979) 659.

[2] J. Kogut, Rev. Mod. Phys. 55 (1983) 775.

[3] K.G. Wilson, Phys. Rev. D14 (1974) 2445.

[4] C.B. Lang, C. Rebbi, Phys. Lett. 115B (1982) 137.

[5] L.H. Karsten, J. Smit, Nucl. Phys. B183 (1981) 103.

[6] L. Susskind, Phys. Rev. D16 (1977) 3031.

[7] K.G. Wilson, in: *New Phenomena in Subnuclear Physics* (Erice, 1975), ed. A. Zichichi (Plenum Press, New York, 1977).

[8] see e.g., P. Ramond, *Field Theory, A Modern Primer* (Benjamin, 1981) 214.

[9] M. Creutz, Phys. Rev. Lett. 43 (1979) 553.

[10] J.M. Hammersley, D.C. Handscomb, *Monte Carlo Methods* (Chapman and Hall, London, 1979).

[11] M. Creutz, B. Freedman, Ann. Phys. 132 (1981) 427.

[12] N. Metropolis et al., J. Chem. Phys. 21 (1953) 1087.

[13] M. Creutz, Phys. Rev. D21 (1980) 2308.

[14] N. Cabibbo, E. Marinari, Phys. Lett. 119B (1982) 387.

[15] E. Pietarinen, Nucl. Phys. B190 FS3 (1981) 349.

[16] G.C. Fox, S.W. Otto, *Algorithms* for *Concurrent Processors*, Physics Today, May 1984; C.L.Seitz, *The Cosmic Cube*, CACM special issue on architecture, January 1985.

[17] S. Otto, Caltech. Ph.D. thesis (1983). Caltech preprint CALT-68-1047.

Figure Captions

(1.1) The basic plaquette "p" introduced in Eq. (1.1.7), shown for convenience on a 2-dimensional lattice. Shown are the 4 directed link matrices joining the sites i,j,k,l (the inverse matrix is used when a link is directed in the opposite direction).

(1.2) The heat bath environment "E" for a link ij shown on a 3-dimensional lattice (see Eq. (1.3.5) and Eq. (1.4.1)). For each of the 4 plaquettes shown, the product of the 3 dashed links is obtained. The sum of these products is the environment "E."

(1.3) Communications required to calculate the corner of a plaquette occurring at node boundaries on a 4-node Hypercube. The double arrows indicate the content and direction of the information that must be sent between nodes. Thus, node 1 sends link matrix A to node 2, which multiplies it by link matrix B to form the corner AB. Similarly, the other nodes form the products CD, EF and GH.

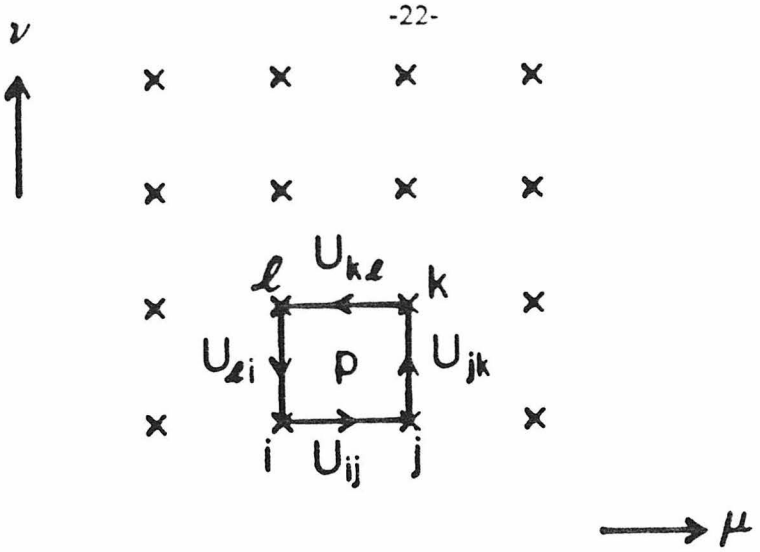

Fig. 1.1

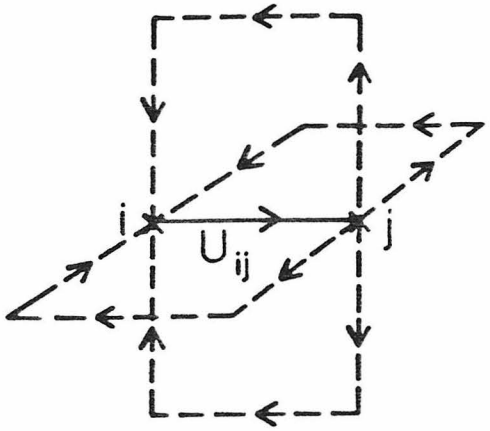

Fig. 1.2

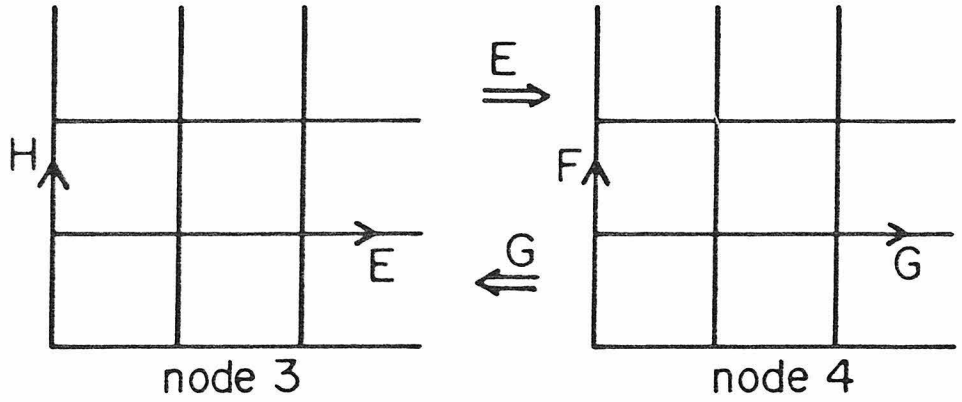

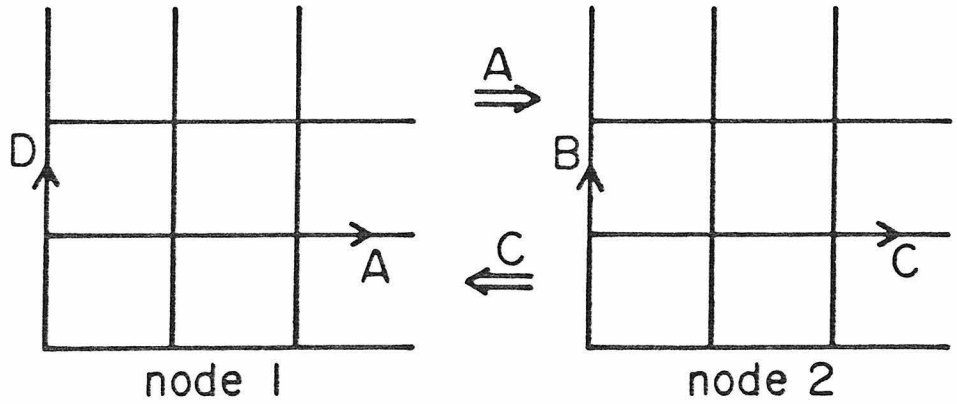

Fig. 1.3

Chapter 2: The SU(3) Heavy Quark Potential

Introduction

With the Monte Carlo machinery available to scrutinise the intermediate coupling regime of field theories, it is now possible to investigate some of the non-perturbative quantities of interest in this region. One of the more important of these is the form of the QCD potential between 2 static quarks in the limit of no dynamical fermions (the "quenched" approximation). Unlike the case of the electrodynamic field, the non-linear dynamics of the QCD gauge fields make the form of this potential a highly non-trivial object that cannot be obtained by traditional methods from the QCD Lagrangian. It is, however, an extremely important quantity to understand, because in the quenched approximation, the potential ought to display confining properties at large distances: i.e., infinite energy should be required to pull quarks infinitely far apart.

In this chapter we describe a high statistics calculation of the QCD static potential. We consider a pair of stationary quarks that are introduced into the lattice as an external current. We then measure the potential energy between these quarks due to the pure gauge part of the action (the quenched approximation), as a function of the distance between the quarks. We compute the potential on lattices of large size (by current standards), namely ones with $12^3 \times 16$ sites. In this way we are able to obtain measurements of the potential out to distances of the order of a hadron radius, 1 fermi or so. We have been able to obtain very high statistics by applying a variance reduction technique which greatly increases the efficiency of the Monte Carlo averaging procedure.

In Section 1 we discuss the lattice quantity from which the potential function can be extracted. This is essentially the coefficient of the exponential dependence of rectangular Wilson loops upon the time dimension, for loops that are much longer in the time than in the space dimension. In Section 2, we describe the statistical improvement which has enabled us to vastly increase the accuracy of our results. In Section 3, we outline the computation itself,

which was performed at several different lattice couplings. In Section 4 we discuss how the measurements made at different couplings are related to each other. We also describe how the self-energy effects of the static quarks impinge on this issue, and how they are accounted for to yield a physical potential function. In Section 5 we present our final results, and discuss how well they compare with phenomenology.

1. Static quark potential

To arrive at a sensible lattice quantity for the static quark potential, consider first the Euclidian path integral, in the continuum, for the gauge fields only:

$$Z[0]=\int d[A]e^{-1/4\int d^4x F^a_{\mu\nu}F^a_{\mu\nu}} \tag{2.1.1}$$

where $F^a_{\mu\nu}$ is given by Eq. (1.1.7). Now consider creating a pair of infinitely massive quarks a distance R apart, having them live for a time T>>R, and then letting them annihilate. This situation, shown in Fig. 2.1, can be described by a current J(x), which takes the value 1 on the contour C formed by the straight lines C_1,C_2,C_3 and C_4, and is 0 everywhere else. Denoting the action in Eq. (2.1.1) by $S_0(A)$, the path integral in the presence of J(x) is now

$$Z[J]=\int d[A]e^{-S_0(A)+ig\int_C A(l)dl} \tag{2.1.2}$$

In the language of statistical physics, the ratio of this quantity to Z[0] can be expressed as the expectation value of the path-ordered operator involving the external current in a system of "energy" $S_0(A)$:

$$\frac{Z[J]}{Z[0]} = <Pe^{ig\int_C A(l)dl}>_{S_0} \tag{2.1.3}$$

Following the discussion in Chapter 1, we formulate the situation on a lattice, introducing a link matrix U_i for every link in the contour C. We now have the correspondence

$$P\, e^{ig\int_C A(l)dl} \approx Tr(U_1 U_2 ... U_n) = W(C) \tag{2.1.4}$$

As required for a physical observable, this object is gauge invariant, because it is a closed loop (see the discussion in Chapter 1). Continuing in the statistical physics language, we can write the expectation value of this Wilson loop W(C) in terms of the free energy,

$$<W(C)>_{S_0} = \frac{Z[J]}{Z[0]} = e^{-(F[J]-F[0])} \tag{2.1.5}$$

Since the process is a static one, we expect

$$F[J]-F[0] \approx V(R)T \tag{2.1.6}$$

in the limit T>>R. We can therefore extract the potential V(R) as

$$V(R) = -\lim_{T\to\infty} \frac{\ln<W(C)>}{T} \tag{2.1.7}$$

Notice that a potential V which behaves linearly with distance,

$$V(R) = KR \tag{2.1.8}$$

results in the Wilson loop showing an exponential dependence on the area RxT of the loop, the classic confining signal first pointed out by Wilson [1]. The constant coefficient, K, is known as the string tension. It is the main physically accessible quantity that we will calculate.

This particular method of extracting the static potential on a lattice was first applied by Stack to the group SU(2), which is computationally less demanding, and then to SU(3) [2]. We have used the same principles in our calculation with the group SU(3) [3]. Other authors have also used the technique [4], and obtained results similar to our own.

2. Variance reduction - the PPR trick

It has been possible to drastically improve the statistical accuracy of the data points over those generated with naive Monte Carlo simulations by using a variance reduction procedure due to Parisi, Petronzio and Rapuano [5], and which we refer to in the following as the PPR trick. This improvement has been of crucial importance in allowing us to extract a statistically reliable signal for the larger Wilson loops, so we will now briefly describe how it works.

In their work, the above authors were attempting to measure the correlation between two parallel Wilson lines, i.e., endless loops that wrap around the lattice periodically. They observed that the product of these lines is most sensitive to the values of the link variables in the lines themselves. It therefore makes sense to average over links in the Wilson lines more rapidly than over the other links in the lattice. The PPR trick is therefore implemented by periodically freezing the lattice variables that do not occur in a Wilson line, and by then averaging each link in the loop itself over several link updates, before multiplying the link averages together to obtain the trace of the whole loop.

We have implemented a slight variation of this method in the measurement of rectangular Wilson loops. Consider again, for example, the loop shown in Fig. 2.1. Denote by $\bar{U}$ the integral over the link variable U with its surrounding links ("environment") held fixed.

$$\bar{U}=\int dU U e^{-\beta S(U)} \tag{2.2.1}$$

The integral for the whole Wilson loop W can be factorised into

$$<W>=<\prod_{C_1} U \prod_{C_2} \bar{U} \prod_{C_3} U^{-1} \prod_{C_4} \overline{U^{-1}}> \tag{2.2.2}$$

provided that the size of the loop is at least 2 lattice spacings in each direction (for the simple Wilson action). The easiest way to evaluate $\bar{U}$ is to estimate it by updating the link many times in a fixed environment (a "mini-update"):

$$\bar{U}=\frac{1}{N}\sum_{i=1}^{N} U_i \tag{2.2.3}$$

The procedure is free of systematic error, since each term in the resulting expansion of Eq. (2.2.2) satisfies the requirement of detailed balance. The method can be thought of as pausing at a certain point in a Markov chain of configurations, and "peeking ahead" to many different, but equally likely, Markov chains. Of course, in order to maintain detailed balance, it is necessary to return to the original lattice configuration after a set of mini-updates, before proceeding with the usual updating.

We have applied the trick to just the parts of the loop contours in the long (time) dimension. The reason for this is that it would not be possible to include the corner links in such a barring procedure. This is because each corner link occurs in the environment of another corner link, so the simultaneous updating of these links would lead to terms in the loop average that violate detailed balance. The proof that detailed balance is satisfied in general when using the PPR improvement, and the reason that it fails for overlapping environments, can be found in Appendix 1.

The real power of this trick stems from the fact that long straight lines of links occur in the observable. In a straight line of L links, applying M mini-updates to each of the links gives an average over L^M link configurations. Also, the computing time taken to do this is somewhat less than the time it would take to do LxM normal link updates, because the environment of each link remains constant, and need not be re-calculated for every mini-update. This environment calculation is the most time-consuming part of the entire computation. Of course, the configurations averaged over using this trick are very correlated by the fixed background, so one doesn't really get L^M independent measurements. The efficiency of the trick is also reduced by the inability to bar the corner links. Nevertheless, using 15 mini-updates per link, we have been able to effectively reduce the computation time by a factor of 100. This corresponds to a factor of 10 gain in the statistical errors of our data points. This variance reduction procedure has also been applied to studies with the SU(2) group [6].

3. The potential measurement

The calculation itself consisted of measuring the ensemble averages of all the planar Wilson loops up to size 6x9. The lattices used were of size $12^3 \times 16$, and the β values ranged from 5.8 to 7.6. The minimum value of 5.8 was chosen because previous work has shown [7] that in the fundamental-adjoint plane, there is a line of first order phase transitions which ends in a critical point extremely close to the fundamental axis near β_F of about 5.6. Other calculations show a peak in the specific heat near this point also [8]. These features are all artifacts of the lattice regularisation. They obscure the behaviour of physical quantities, so we avoid this region in our simulations. Since we are using the most naive Wilson action which contains only the fundamental coupling β_F, we stick to the regime above 5.7.

We must also be careful to take account of finite size effects. The string tension K determines the energy scale. The inverse of this quantity is therefore a measure of the correlation length, and in our case this reaches half the size of the shortest lattice dimension at β=6.4. Also, if the time dimension is short enough, we are not really simulating a 4-dimensional world extending infinitely in all directions, but rather a 3-dimensional one held at a finite temperature T, given by the inverse of the time dimension. In fact, this system has been found to have a phase transition beyond which the confining behaviour breaks down [9]. This deconfining transition occurs when the smallest dimension is below $\frac{1}{66\Lambda_L}$, which corresponds in our case to $\beta \approx 6.55$. Hence our results above $\beta \approx 6.4$ must be treated with caution, as they are somewhat prone to systematic size effects.

Measurements of the loops were made on every 10^{th} configuration, the latter being generated by the pseudo heat-bath algorithm described in Chapter 1. The PPR trick was applied with 15 mini-updates per link to obtain high statistics averages. The results for a typical set of loops calculated on a lattice at the coupling β=6.0 are shown in Fig 2.2. Notice that for each value of R, the logarithm of the loop values follows, with the exception of the small T points, a straight line as a function of T. This strongly exponential behvaviour of the Wilson loops shows that the system we are measuring does indeed lie almost entirely in the ground state,

and therefore our results can be expected to give reliable information about a static quark/anti-quark system.

4. Scaling and the quark self-energy

Now that we know the details of the exponential behaviour of large Wilson loops at different lattice couplings, we must relate the measurements at different couplings to obtain phys-ically useful results. If we assume we are sufficiently near the continuum limit that lattice spacing effects are negligible, and if we also assume that the effects of the finite size of the lattice are small, then there is only one dimensional quantity remaining in the theory, namely the correlation length ξ. Any physical quantity must be a dimensionless number multiplied by this length. If we consider describing the theory on lattices of different couplings g, there must be a well-defined dependence of the coupling g on the lattice spacing a in order to keep the phys-ical correlation length constant. This well-defined behaviour $g(a)$ in the region $a<<\xi<<L$ is known as *scaling*. To see if one is in the scaling region, one finds the ratios of various observable quantities calculated on a lattice at a certain coupling. If these ratios remain unchanged when a different coupling is used, one must be in the scaling region, since the effects of the lattice spacing have clearly been washed out.

In general, the functional form $g(a)$ depends on the specific lattice formulation that is chosen. However, for small g there is a unique functional form that is independent of the lattice formulation (and, in fact, independent of any sort of regularisation scheme). It is calculable in perturbation theory (i.e., for $g<<1$). Up to 2-loop order, it is given by

$$\rho(\beta) = a\,\Lambda_L = \left(\frac{8\pi^2\beta}{33}\right)^{51/121} e^{-4\pi^2\beta/33} \qquad (2.4.1)$$

where Λ_L is a dimensional constant that must be set by appeal to experiment. Up to 2-loop order, all regularisation schemes, including lattice regularisation, differ only in the value of Λ_L. When higher loops are taken into account, the form Eq. (2.4.1) changes, depending on the regularisation method.

Fortuitously, it turns out that this expression for $g(a)$ seems to remain applicable (to a large extent) even in the region we are studying, which is $g \approx 1$. What this means is that physical observables, such as masses, computed on lattices at different couplings, remain roughly constant if Eq. (2.4.1) is assumed to relate the lattices. We can also justify this form *a postiori* from the potential calculation itself in the following way. We obtain a very smooth potential function involving just 2 free parameters from calculations on lattices at several different couplings, assuming that Eq. (2.4.1) holds. Both these parameters have natural physical interpretations. If Eq. (2.4.1) had not been at all appropriate, we would presumably have obtained a much more complicated "potential" function requiring many more arbitrary parameters to obtain a good fit to the data. This scaling behaviour associated specifically with the form predicted by perturbation theory is known as *asymptotic scaling*.

It could be objected, of course, that perturbation theory has now reared its head in the calculation and destroyed its non-perturbative character. However, this is not really the case. We are using the perturbative scaling function strictly as the best "guess" for the functional dependence $g(a)$, but the area dependence of Wilson loops remains a fundamentally non-perturbative phenomenon. The reason for the precocious onset of the 2-loop result at $g \approx 1$ is not known (and indeed, there do seem to be some deviations from this asymptotic scaling in the Monte Carlo results). This corresponds to correlation lengths of roughly 2 lattice spacings. Nor is it understood why the correlation lengths can be so close to the lattice size. Indeed, these two fortuitous but mysterious events are the reasons that Monte Carlo simulations are possible at all on present-day computers.

We are now in a position to look at the asymptotic scaling properties of the potential. Denote by r and t the number of lattice spacings in the spatial and time directions respectively representing the dimensional distances R and T. If we are in a region of β where the scaling function Eq. (2.4.1) holds, then the dimensionless function $V(R)/\Lambda_L$ must be a function solely of the only dimensionless variable available, $R\Lambda_L$, by simple dimensional analysis:

$$\frac{1}{\Lambda_L} V(R = ra, \Lambda_L) = \frac{V}{\Lambda_L}(ra\,\Lambda_L) \tag{2.4.2}$$

Therefore, in a graph of the variable $x = ra\Lambda = r\rho(\beta)$ versus $\frac{1}{\rho(\beta)t}\ln W(ra, ta)$, all the points should fall along the same curve. We can actually obtain results in more convenient units by dealing in terms of the length $\xi = 0.011/\Lambda_L$, and plotting the function ξV against the dimensionless variable $y = R/\xi$. Fig. 2.3 shows these plots obtained at each of the lattice couplings.

From Fig. 2.3, it is clear that scaling eludes us. The reason is that the quantity V(R) of Eq. (2.1.7) is not physically measurable! It contains the self-energy of the static quark lines. This self-energy gives a contribution to the Wilson loops which depends on the length of the quark lines. It does not scale and must be removed.

There are two approaches to this problem. One is to observe that the self-energy contribution depends only upon the lattice coupling β, and in particular is independent of the quark separation R. This means that the force between the quarks, i.e., the gradient of the potential function, will not contain the self-energy term (this must, of course, be true since the force between quarks is a genuinely measurable quantity). One could therefore eliminate the self-energy by taking the finite differences of potentials at different couplings and then numerically integrating the resulting force function.

Alternatively, since the self-energy is a constant for a given coupling, one can just move the y-intercept of the potential function obtained on each lattice so that the superimposed curves produce a single function that best fits the data points. In previous calculations [2], the former method was used. However, in practice the second, simpler method, was found to achieve essentially the same results, so it was adopted for the current calculation. The curves obtained at different couplings overlap sufficiently to make the fitting of the points to a single function a non-trivial test of the scaling of the results.

5. Results

In order to detect quantitatively any non-scaling behaviour, we introduced a parameter f to the 2-loop expression for Λ_L. Thus, we used a family of Λ parameters given by

$$\rho_f(\beta)=a\,\Lambda_L(f)=\left(\frac{8\pi^2\beta}{33}\right)^{51/121}e^{-4\pi^2\beta f/33} \tag{2.5.1}$$

The 2-loop perturbation result corresponds to $f=1.0$. For each value of f, a least sqares fit was performed to the following 3-parameter functional form:

$$\xi V=-\alpha/y+Ay+B \tag{2.5.2}$$

The reasons for choosing this particular functional form will be discussed shortly. It was found that when the $R=a$ points were included in these fits, unacceptably high χ^2 values were obtained for all values of the parameter f. This is to be expected, since clearly the lattice spacing effects will be dominant at these distances. The $R=a$ points were therefore discarded from the analysis. With the remaining points, the optimal value of f (i.e., the one that gave the minimum χ^2) was found to be $f_0=1.05\pm.08$. The uncertainty was obtained by finding the value it took to raise χ^2 by one unit. Clearly the perturbative scaling function Eq. (2.4.1) is a perfectly acceptable fit to the data. In Fig. 2.4 we show all our data points plotted assuming the perturbative scaling form. For the alpha parameter in Eq. (2.5.2), our best fit gives the value

$$\alpha=0.25\pm0.02 \tag{2.5.3}$$

Since the parameter A is related to the string tension by $K=A/\xi^2$, we re-express it as the more conventional ratio $\Lambda_L/\sqrt{K}$. We obtain

$$\Lambda_L/\sqrt{K}=(9.4\pm0.3)\times10^{-3} \tag{2.5.4}$$

Now, why is it that we expect to see the potential follow the functional form Eq.(2.5.2)? Let's look first at the piece that grows linearly with distance. It has long been known that the

hadron spectrum, roughly speaking, obeys the Regge-Chew-Frautschi relation [10]

$$J = \alpha_0 + \alpha' M_J^2 \tag{2.5.5}$$

where J is the angular momentum of a particle, and M_J is its mass. The coefficient α' has a value of $0.85 \pm 0.05 GeV^{-2}$. This linear trajectory can be obtained from a simple model in which the hadrons are assumed to consist of massless quarks rotating about the end of a string held under a constant string tension k (energy per unit length) [11]. The string tension is related to α' by

$$k = \frac{1}{2\pi\alpha'} \tag{2.5.6}$$

Thus, phenomenology gives an estimate for the string tension of

$$k = 0.19 \pm 0.01 \ GeV^2 \tag{2.5.7}$$

Our calculation shows a definite linear behaviour of the potential at large distances, and from the slope of this line we are able to extract a string tension K in terms of the renormalisation group invariant Λ_L (Eq. (2.5.4)). In order to find out what this corresponds to in physical units, we need some independent way of determining the value of Λ_L, i.e., setting the absolute scale of the problem.

To do this we have used the results of a Monte Carlo calculation which computes a different observable, namely the masses of hadrons in the quenched approximation [12]. A reasonable fit to the meson spectrum is made by these authors when the scale is set by

$$a(\beta=6.0) = 0.1 Fermi = (2GeV)^{-1} \tag{2.5.8}$$

Using this absolute scale, our estimate for the physical string tension becomes

$$K = 0.23 \pm 0.01 \ GeV^2 \tag{2.5.9}$$

The uncertainty quoted here is purely statistical. The setting of the scale actually introduces a

somewhat larger systematic error, since spectrum calculations are in their infancy and have only been made in the quenched approximation. It is unlikely that the absolute scale is known to better than 10–20%. Given this uncertainty, it is remarkable that our string tension result Eq. (2.5.9) agrees to within 15% with the value obtained from phenomenology, Eq. (2.5.7).

The reason for the presence of the $1/R$ term in Eq. (2.5.2) is somewhat more obscure. At very short distance scales, $\sqrt{K}R \ll 1$, one expects to have a $1/R$ potential due to single gluon exchange between quarks, with the coefficient being the running coupling calculable by renormalisation group improved perturbation theory [2,13]. On the other hand, at large distances, $\sqrt{K}R \gg 1$, low frequency fluctuations of the string give rise to a $1/R$ potential also [14]. The trouble is that our calculation was carried out in the intermediate region where neither approach is really valid, so it is difficult to interpret our results. Various analyses have been attempted to resolve the issue, but they have been inconclusive [15]. Perhaps a small clue that the perturbative origin is more reasonable is given by a new calculation in which the 3-quark potential is computed. The $1/R$ coefficient appears to be roughly half the 2-quark value, as predicted by perturbation theory [16]. However, these results are very preliminary, and should be treated with caution.

References

[1] K.G. Wilson, Phys. Rev. D14 (1974) 2445.

[2] J.D. Stack, Phys. Rev. D27 (1983) 412;

J.D. Stack, Phys. Rev. D29 (1984) 1213.

[3] E. Brooks III et al., Phys. Rev. Lett. 52 (1984) 2324;

S.W. Otto, J.D. Stack, Phys. Rev. Lett. 52 (1984) 2328.

[4] D.D. Barkai, K. Moriarty, C. Rebbi, Phys. Rev. D30 (1984) 1296.

[5] G. Parisi, R. Petronzio, F. Rapuano, Phys. Lett. 128B (1983) 418.

[6] F. Karsch, C.B. Lang, Phys. Lett. 138B (1984) 176.

[7] G. Bhanot, M. Creutz, Phys. Rev. D24 (1981) 3212:

G. Bhanot, Phys. Lett. 108B (1982) 337.

[8] K.C. Bowler et al., Nucl. Phys. B240 (1984) 213.

[9] F. Karsch, R. Petronzio, Phys. Lett. 139B (1984) 403.

[10] G.F. Chew, S.C. Frautschi, Phys. Rev. 123 (1961) 1478;

T. Regge, Nuovo Cimento 14 (1959) 951.

[11] Y. Nambu, Phys. Rev. D10 (1974) 4262.

[12] H. Lipps et al., Phys. Lett. 126B (1983) 250.

[13] E. Kovacs, Phys. Rev. D25 (1982) 871.

[14] M. Luscher, K. Symanzik, P. Weisz, Nucl. Phys. B173 (1980) 365;

J.D. Stack, M. Stone, Phys. Lett. 100B (1981) 476.

[15] P.M. Fishbane, P. Kaus, S. Meshkov, Phys. Rev. D33 (1986) 852;

M. Flensburg, C. Peterson, University of Lund preprint LU TP 85-14.

[16] J.W. Flower, S.W. Otto, private communication.

Figure Captions

(2.1) The Wilson loop used to calculate the string tension (see Eq. (2.1.4)). The distance R is the separation between a pair of quarks that live for time T. The Wilson loop W(C) (or W(R,T)) is the trace of the product of all the directed link matrices along the contour C formed by the straight lines C_1,C_2,C_3 and C_4.

(2.2) The lattice observable $-\ln\langle W(R,T)\rangle$ of Fig. 2.1, plotted as a function of the number of lattice spacings in the time direction T/a, measured at the coupling β=6.4. The dimensionless potential V/Λ_L for a given separation R is obtained from the slope of the appropriate line on this graph (see Eq. (2.4.2)).

(2.3) The dimensionless potential ξV plotted as a function of quark separation R, for several different values of β. The distance scale is set by Eq. (2.5.8).

(2.4) The dimensionless potential $V/\sqrt{K}V$ plotted as a function of the scaling variable $\sqrt{K}R$ after self-energy corrections have been made, for all values of the coupling β. Eq. (2.5.4) has been used to relate $\sqrt{K}$ to Λ_L.

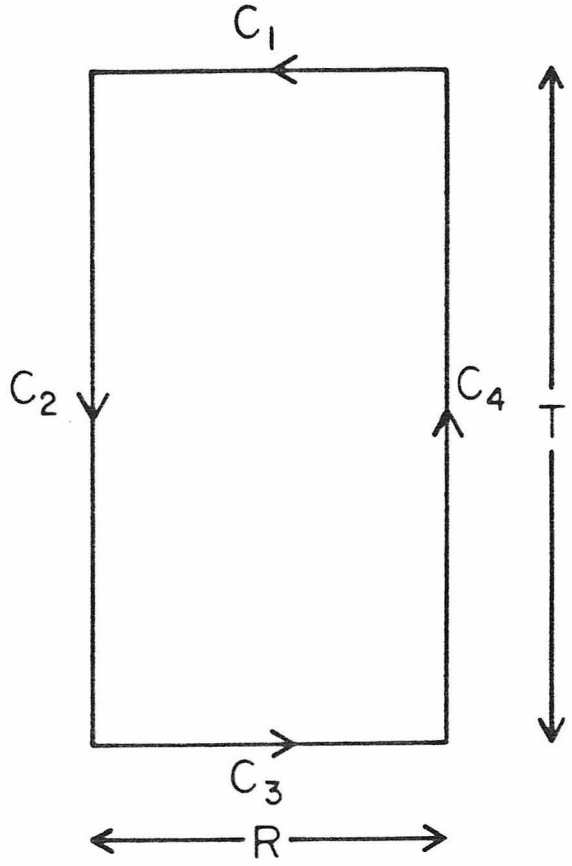

Fig. 2.1

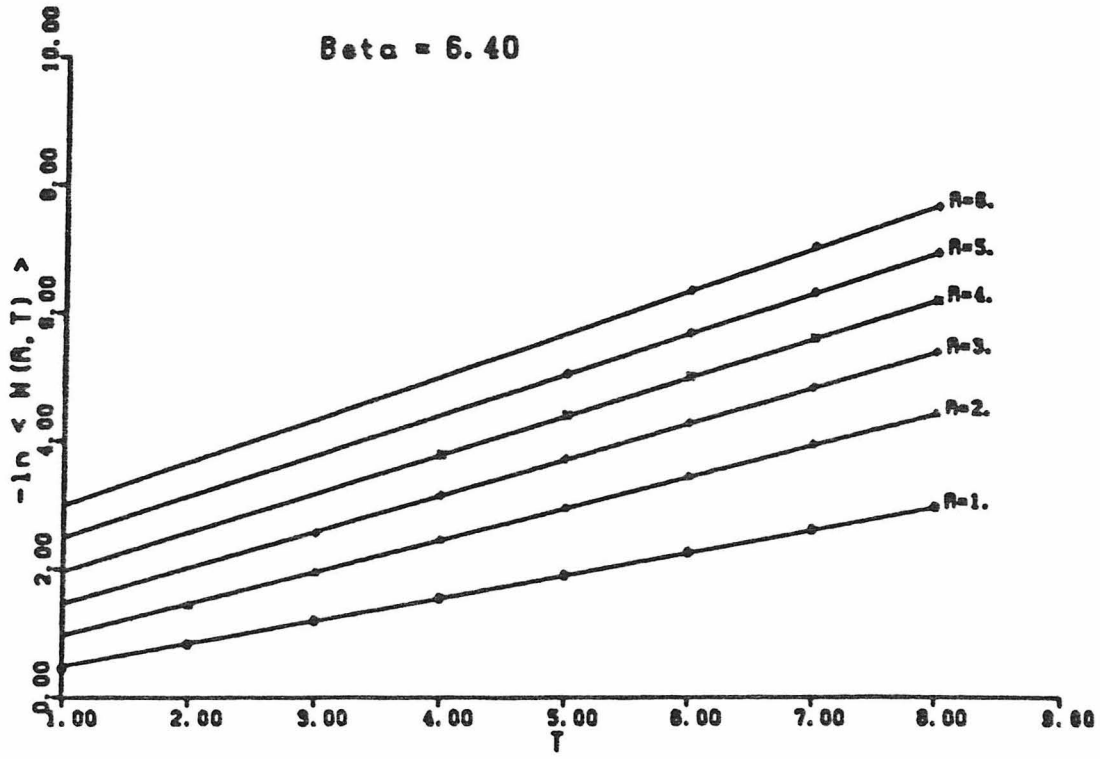
Beta = 6.40
-ln < W(R,T) >
0.00
2.00
4.00
6.00
8.00
10.00
1.00
2.00
3.00
4.00
5.00
6.00
7.00
8.00
9.00
T
R=6.
R=5.
R=4.
R=3.
R=2.
R=1.

Fig. 2.2

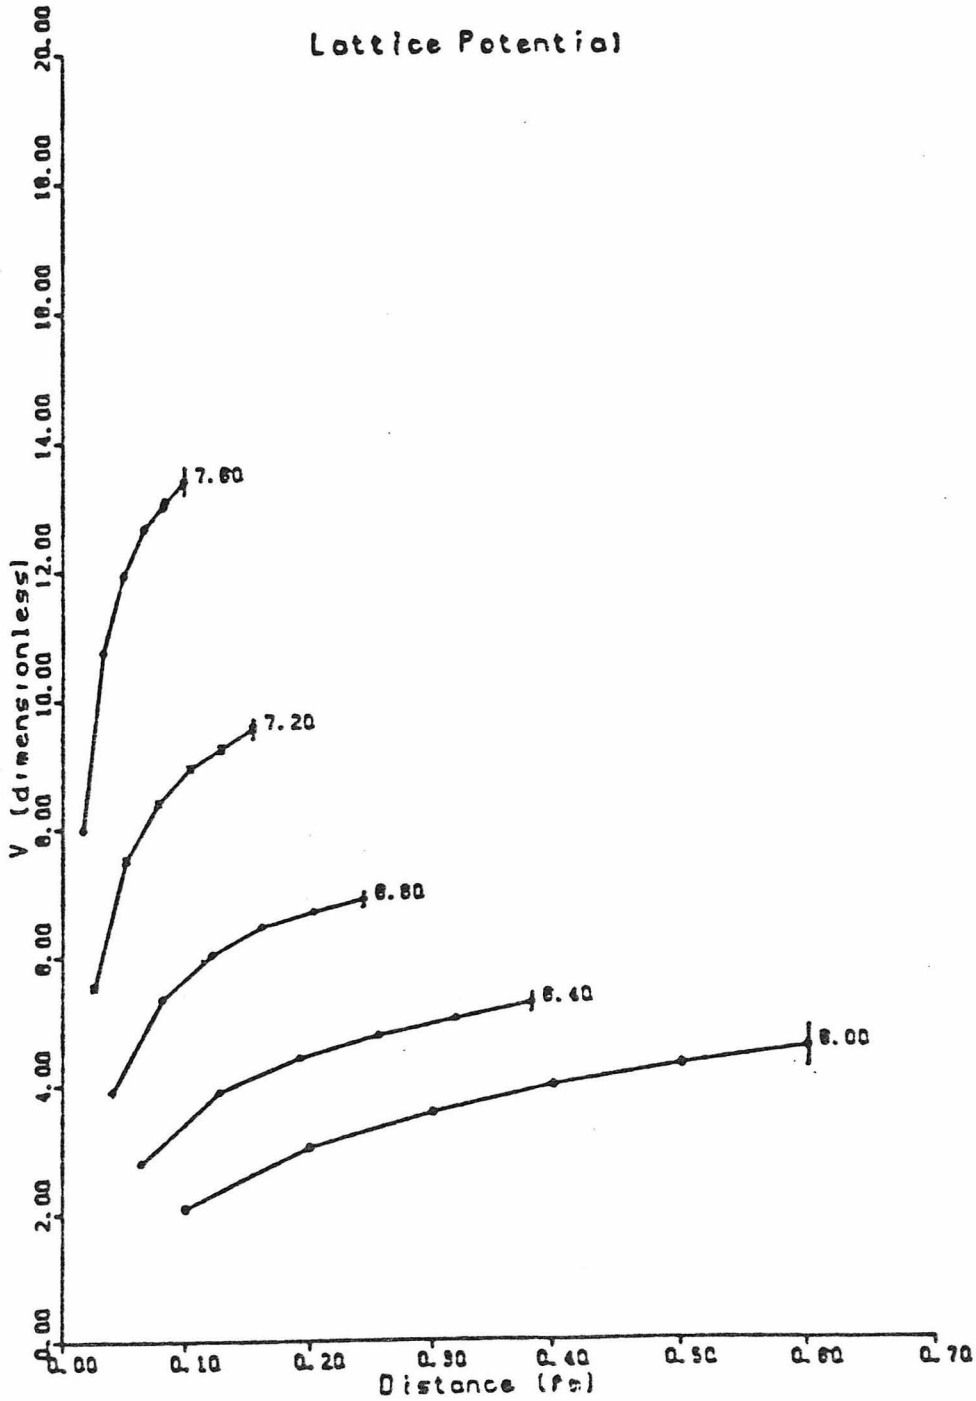
Lattice Potential
V (dimensionless)
Distance (fm)
7.60
7.20
6.80
6.40
6.00
0.00
2.00
4.00
6.00
8.00
10.00
12.00
14.00
16.00
18.00
20.00
0.00
0.10
0.20
0.30
0.40
0.50
0.60
0.70

Fig. 2.3

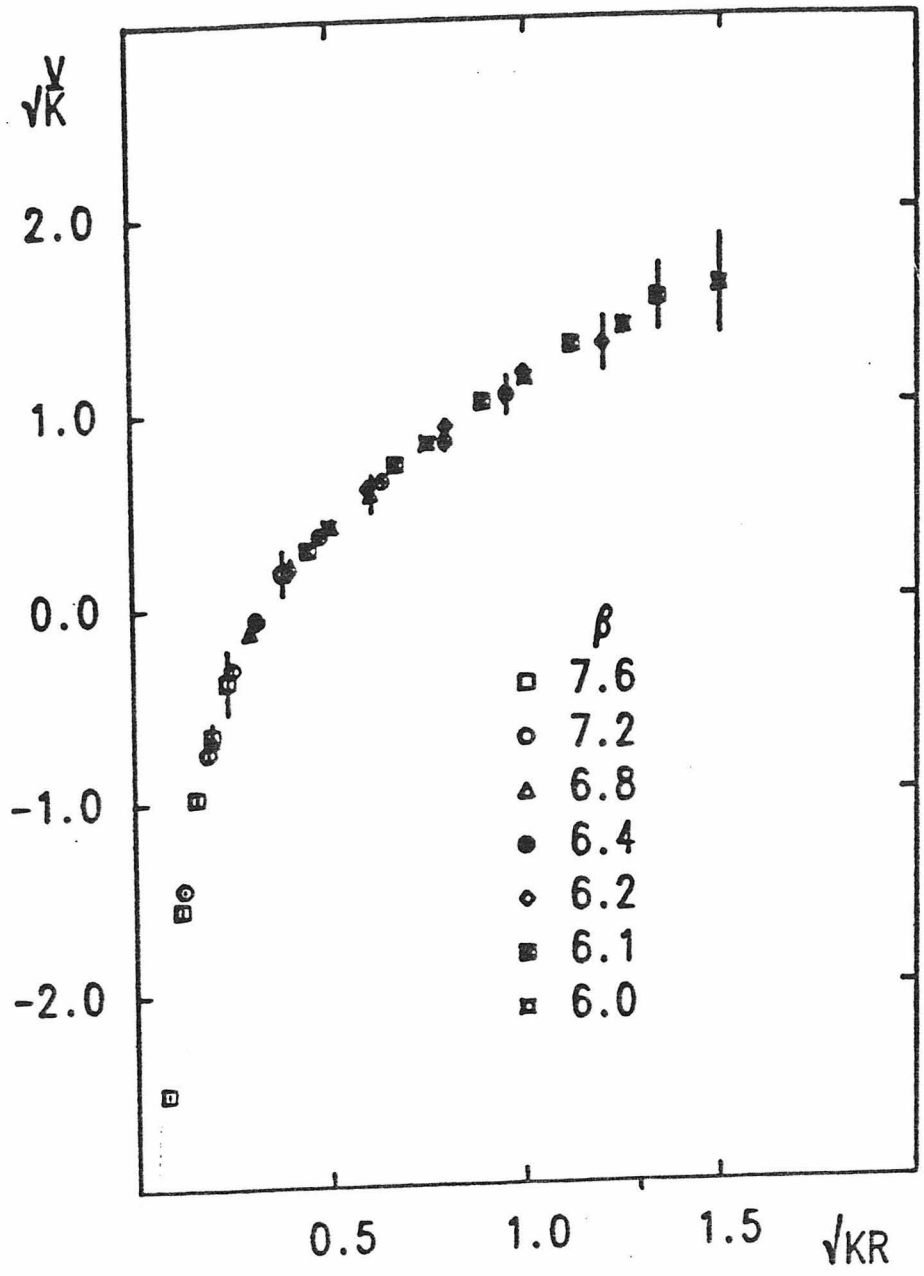

Fig. 2.4

Chapter 3: The Glueball Mass on Large Lattices

Introduction

Since the QCD Lagrangian contains a non-linear term involving only the gauge degrees of freedom, one expects the spectrum of particles from this theory to include states that are formed purely (or mainly) of gluonic fields interacting with each other. In contrast, QED contains no such self-interacting term, which means that the QED gauge field (the photon) can occur only in conjunction with fermions. These states of pure glue occurring in QCD are referred to as glueballs. Since, with the lattice, we now have a concrete non-perturbative formalism which is described in terms of the gauge fields, one might hope to be able to investigate some of the properties of these glueball states on the lattice.

In this chapter, we describe a calculation of the mass of the 0^{++} glueball in the quenched approximation of SU(3), where there are no dynamical quark loops. The glueball states in general are labelled J^{PC}, which specifies their angular momentum, parity and charge conjugation quantum numbers. As we have discussed in Chapter 1, this approximation is expected to explain many features of the glueball spectrum, although presumably the inclusion of fermions will alter the quantitative results somewhat.

It turns out that the observable quantity on the lattice from which one extracts the glueball mass is extremely sensitive to statistical noise. Therefore, the main thrust of our work has been to devise various ways of extracting the useful signal more efficiently. For this reason, we have concentrated on looking at the properties of a single glueball state. Eventually, of course, one wants to compute the masses of all the glueball states of interest, since on a sufficiently fine-grained lattice the ratios of the masses of all these states will be determined uniquely by the lattice, without requiring any external parameters.

In Section 1, we describe how glueball masses are computed on the lattice. In the following 3 sections, we discuss the different methods that we have employed to extract these masses efficiently from the Monte Carlo simulations. In Section 5, we present our results for

the mass of the 0^{++} glueball, and discuss the benefits and limitations of the technique.

1. Glueball masses on the lattice

To see how a glueball state would manifest itself on the lattice, consider the 2-point correlation function of a gauge-invariant operator, O, over a time t. It is easy to show [1] that the conventional Lagrangian path integral can be reformulated in terms of a transfer matrix T, in which a state defined at some zero time evolves into a new state over a time t by application of the operator $T = e^{-Ht}$, where H is the Hamiltonian of the system. The equivalence is demonstrated by splitting up the time interval t in the transfer matrix formalism into infinitesimal segments Δt, and then inserting a complete set of states at each Δt interval. With this reformulation, we can now write the 2-point function as

$$<0|O(t)O(0)|0> = <0|O(0)e^{-Ht}O(0)|0>. \tag{3.1.1}$$

Now insert in the above expression a complete set of energy eigenstates

$$1=\sum_{n=0}^{\infty} |n><n|. \tag{3.1.2}$$

Then

$$<0|O(t)O(0)|0>=|<0|O(0)|0>|^2+|<0|O(0)|1>|^2e^{-E_1t}+.... \tag{3.1.3}$$

We are only interested in the physical states above the vacuum, so we will restrict our attention to the connected 2-point correlation function

$$<0|O(t)O(0)|0>_C=<0|O(t)O(0)|0>-|<0|O(0)|0>|^2 \tag{3.1.4}$$

Clearly, as one goes to the limit $t\rightarrow\infty$, all states except $|1>$, the lowest state above the vacuum, will be exponentially damped out. Now choose an operator O which has zero momentum (i.e., covers an entire time slice). Then the rate of fall-off of the correlation function is a measure of the mass of the lowest glueball state above the vacuum. In practice, finite

computer power limits the number of time slices over which correlations can be measured, since the signal tends to be swamped by statistical noise over more than 2 or 3 time slices. On the other hand, we want to extract masses from the tail of the correlation function, where the effects of states above the ground state are dying off. The best compromise is to calculate the mass from the ratio of 2-point functions over the longest and second-longest time slices that give statistically reliable signals. Thus

$$m \approx \frac{1}{t_n - t_{n-1}} \ln\left[\frac{<O(t_{n-1})O(0)>}{<O(t_n)O(0)>}\right] \tag{3.1.5}$$

In order to extract a glueball state of given angular momentum J, parity P, and charge conjugation C, one merely selects an operator O which satisfies these symmetries. Of course, it may not be possible to construct an operator on the lattice which distinguishes between all the states. For example, the signals of the 0^{++} and 4^{++} states will overlap. However, in these cases the states of higher mass should be highly suppressed by the exponential if one goes to high enough t, so in principle one can uniquely identify the mass of the lowest-lying state in the $t \to \infty$ limit.

Unfortunately, this method of calculating glueball masses runs into severe difficulties because of the limitations of current computer power. It turns out that at the couplings used presently, correlation functions must be calculated over at least 4 lattice spacings in order to sufficiently damp out the effect of unwanted excited states. However, at these distances the statistical noise in the ensemble averages is of the same order of magnitude as the signal itself. This problem is especially acute for some of the more interesting glueball candidates such as the $1^{\mp}$ *oddball*, a state that cannot be formed by any combination of quark quantum numbers. The reason that the problem is especially bad here is that terms of roughly equal size occur with opposite signs in the appropriate operator.

One interesting way of tackling these problems is to introduce an external source function. One then measures the response of the correlation functions to this source [2]. In this way

a strong signal can be measured over longer distances than the source-less method can supply. However, there is now an explicit systematic error introduced, since the vacuum is being perturbed in order to obtain information about it. This error can, of course, be controlled by dialing the strength of the source to smaller values, and/or by going to larger correlation distances, both at the expense of decreasing the signal strength.

It is not yet clear which of these two methods is ultimately the best way to perform glueball calculations. We have chosen to attack the standard sourceless method in an attempt to improve its efficiency. Firstly, we have implemented a technique suggested by Wilson [3], and subsequently developed by several authors [4], called the Monte Carlo Variational Method, which basically seeks to improve the signal by maximising the overlap between the chosen operator on our lattice and the state whose properties we wish to measure. Secondly, we have applied the *smearing* techniques of [5] to improve the signal at the expense of introducing opertors of small but non-zero momentum. Admittedly, this trick re-introduces a small amount of systematic error, but it lends itself well to the other improvements we have chosen, and we have implemented it carefully to see that the error is well-controlled. Thirdly, we have applied the variance-reduction idea of Parisi et al.[6], which was outlined in Section 2 of Chapter 2, and which was used extensively in the calculation of the static quark potential. This latter improvement has not been attempted before in the case of glueballs. The nature of the glueball observable makes slight modifications of the trick necessary in the glueball case. Finally, we perform the calculations on large lattices of size $12^3 \times 16$. This reduces the finite size effects, and allows us to obtain good statistics by using the translational invariance of the glueball operator. We will now describe the progress we have made using these improvements.

2. Monte Carlo Variational Method

The basic idea here is to use a Rayleigh-Ritz style variational principle. Suppose we want to measure the properties of the eigenstate $|s>$, which is known to be the ground state of a certain Hamiltonian, and that the operator we have chosen produces a general state $|\psi>$ which is a linear combination of $|s>$ and the other eigenstates of the Hamiltonian:

$$|\psi>=a_s|s>+\sum_{i\neq s} a_i|i> \tag{3.2.1}$$

Then it is straightforward to convince oneself that

$$\frac{<\psi|H|\psi>}{<\psi|\psi>} \geq E_s \tag{3.2.2}$$

for a Hamiltonian with a positive spectrum (E_s is the energy of the ground state $|s>$). One therefore attempts to find the state $|\psi>$ which minimises $<\psi|H|\psi>$. This is actually done by *maximising* $<\psi|e^{-Ht}|\psi>$, the 2-point function.

In our work, we constructed operators O from linear combinations of 4 different shaped loops, shown in Fig. 3.1 ,

$$O=\sum_i c_i A_i \tag{3.2.3}$$

where the A_i are the loops shown in the figure. The coefficients c_i were varied in such a way as to maximise the expression [7]

$$m_g(t)=\frac{-1}{t}\ln(\max\frac{\sum_{ij} c_i c_j C_{ij}(t)}{\sum_{ij} c_i c_j C_{ij}(0)}) \tag{3.2.4}$$

where the C_{ij} are the connected correlation functions defined by

$$C_{ij}(t) = <0|O_i(0)O_j(t)|0> - <0|O_i|0><0|O_j|0> \tag{3.2.5}$$

The details of the method can be found in Appendix 2.

The method was applied to the calculation of the 0^{++} mass for the SU(3) gauge group. The variational method was applied to obtain the optimal coefficients c_i from the correlation functions over the first time slice, where the statistical noise is minimal. The c_i were then used to calculate the masses from the correlation functions measured over higher time slices (up to distance 4 lattice links). Unfortunately, we found that the masses did not converge toward a single value as more statistics were collected. Contamination of the signal by unwanted states in the first time slice is giving spurious values for the optimal coefficients c_i. We then tried finding the optimal coefficients from the correlation functions over the second time slice, but still convergence problems persisted, since the signal of the correlations over the second time slice had substantially more statistical fluctuation than the signal over the first time slice.

As a result, we were forced to discard the MCVM for our calculation. We then looked at the fall-off of each individual loop operator A_i. At large distances, the signal of the plaquette operator was consistently stronger than that of any other loop, so in the rest of our glueball work we restricted our attention entirely to this operator.

3. Smearing

Since we are calculating particle masses, ideally we want our operators O to extend over an entire time slice; that is, we want them to generate states of zero momentum. However, in order to reduce finite size effects, it is desirable to perform calculations on lattices that are as large as possible. With only one operator per time slice, the statistical noise on such lattices is very large considering the big CPU investment required to generate the configurations.

Therefore, in order to be able to do computations on large lattices, we have adopted the strategy of using operators that extend over finite spatial areas of the lattice [5]. By averaging over the correlations between all such areas, a significant increase in signal can be achieved for a given number of configurations generated. For a glueball with quantum numbers J^{PC}, we constructed operators which transform as J^{PC} about a particular site. The finite spatial operator, called a *smear*, is then obtained by summing these individual operators over a block of

sites in the spatial directions of the lattice (see Fig. 3.2).

Of course, states of non-zero momentum are now contributing to the observable, up to some limit Δp which varies inversely with the smear size. However, provided that Δp is small compared to the mass, each momentum state is expected to contribute according to the free particle dispersion relation, so the total energy can be calculated as a function of mass and smear size. Indeed, by using smears of different sizes, we have checked that the energy does decrease as one goes to higher smear sizes. Asymptotic behaviour of the energy occurs at the higher smear sizes, as can be seen in Table 3.1, so we are confident that our estimates are an accurate indication of the glueball masses themselves.

4. The PPR trick for glueballs

Recall from the discussion in Chapter 2, that the PPR trick consists essentially of averaging more over the links that occur in the observable itself than over the other links in the lattice [6]. Its power derived mainly from the fact that the static quark observable contains straight lines of links, in which the plaquettes surrounding each link in a given Wilson line do not contain any other links in the Wilson line. This situation dissolves for glueballs, because the links within smear operators impinge greatly on each others' environments. One could conceivably use the Parisi trick to average over every fourth link in a smear, so that environments did not overlap with other links, but the procedure is awkward, and only some of the variables are integrated out in this manner.

However, since the actual observable consists of the correlation function between 2 smears separated in the time direction, a more straightforward improvement goes as follows. Suppose one wants to find the correlation between a particular pair of smears. One can freeze all the links that do not contain any links in the 2 smears, and then perform normal updates on the links in *only* the 2 smears, updating each link in a smear just once at each turn, in order to satisfy detailed balance (of course, many such mini-sweeps can be performed in order to improve the statistics). This whole procedure is then repeated over all possible pairs of

smears. One merely requires that the smears in each pair be separated by at least 2 lattice spacings, in order to satisfy the overlap criterion.

The procedure is illustrated in Fig. 3.3 for a given pair of smears. The links which must be used in order to update the links in a smear, but which are not links in the smear operator O itself, are denoted by E and F for the 2 smears. Our improvement then corresponds to the equation

$$<O(t)O(0)> = <\overline{O(t)}\ \overline{O(0)}> = <\frac{1}{N}\sum_{i=1}^{N} O_i(E)\ \frac{1}{N}\sum_{i=1}^{N} O_i(F)>, \qquad (3.4.1)$$

where there are N sweeps performed for each smear: that is, $O_i(E)$ is obtained by updating each link in the smear with fixed environment E once, satisfying detailed balance at each step.

This improvement by itself is of limited value, because with the lattice size that we have chosen, there are 12^3 different smear operators in each time slice, one for each possible position of a smear in the time slice. For a smear of size s, each link will appear in s^3 different smear operators. If there are N updates per smear, each link will therefore be updated $s^3 \times N$ times. However, a little thought reveals that it is perfectly legal to update the links of an entire time slice at a time, instead of updating each smear separately. After such an update, which we call a *hypersweep*, all 12^3 smears are measured. Hence, during N hypersweeps, each link need now only be updated N times, a factor of s^3 lower than that needed to obtain the same number of measurements using the naive improvement.

The method is illustrated in Fig 3.4. For a given smear, the part of its environment which lies in the spatial directions is now changing with each hypersweep, so the environments E_i and F_i now have subscripts. The correlation of 2 particular smears is now written as

$$<O(t)O(0)> = <\frac{1}{N}\sum_{i=1}^{N} O_i(E_i)\ \frac{1}{N}\sum_{i=1}^{N} O_i(F_i)>. \qquad (3.4.2)$$

There are still no links in either smear's environment impinging on the other smear provided that $t \geq 2a$, so the procedure is free of systematic error. Of course, at the end of N

hypersweeps on a given time slice, it is necessary to return to the original configuration of the time slice before proceeding to the next slice, or the detailed balance condition *would* be violated.

5. Results for the O^{++} glueball

Our calculation was performed at the coupling β=5.7, on a lattice of size $12^3\times 16$ [9]. No attempt was made to try to find out about the scaling of the observable by measuring it at different couplings. This is ultimately, of course, an important question, but the computational demands of the problem are already severe enough. The simplest operator that satisfies the 0^{++} quantum numbers is the rotationally invariant combination shown in Fig 3.2. As described in the previous MCVM section, suitably transforming operators were constructed for 4 different loops. However, only the operator based on the simple plaquette gave a reliable signal over large distances, so our final data analysis is based solely on this operator.

We calculated the correlation functions between smears of sizes 2, 3, 4 and 5. Measurements were made over 40 gauge configurations, each separated by 10 sweeps to reduce correlations. Since we were looking particularly at the importance of the hypersweep improvement, we used 3 different values of the number N of hypersweeps: N=0,N=10 and N=50. The error bars were obtained by blocking the connected correlation functions into 4 sets of 10, and propagating the resulting standard deviation through the mass formula.

The results are shown in Table 3.1. They show, firstly, that there is a monotonic decrease in the energy estimate as the smear size is increased, which shows that we are converging towards the mass of the glueball state. Secondly, for each smear size, there is a monotonic behaviour in t of the mass estimate. This is important because it indicates that at small values of t the mass estimates are being severely affected by radial excitations of the 0^{++} state. It demonstrates that one must indeed go to large t values to extract reliable mass estimates.

Finally, the table shows the most important feature of the method, which is, that in the N=50 case, we have been able to obtain a reliable signal for the correlation functions out to

the 4^{th} time slice. There was even a positive signal over the 5^{th} time slice, but it fluctuates significantly as more measurements are obtained. It is not reliable enough to use for a mass estimate, and we do not present it here. This is an encouraging result because all previous estimates of the SU(3) 0^{++} mass have been unable to obtain accurate estimates from beyond the 3^{rd} time slice, since the correlation functions have been lost in the statistical noise at large distances. More precisely, this is true of the methods that are free of systematic errors. Source methods do get signals over larger distances, but at the cost of the systematic errors discussed earlier.

Obviously, there will be an ideal ratio of the number of hypersweeps to normal sweeps in order to minimise the statistical error for a given amount of CPU time. For instance, although the N=50 results show a strong signal over 4 time slices, the time required to perform this many hypersweeps dominates the generation of a usual gauge configuration by a factor of 10. By measuring the statistical errors for different numbers of hypersweeps to greater precision, one could determine the optimal value of N. However, the statistical errors in our calculation are not precise enough to be able to perform such an analysis.

We will take our value for the mass gap obtained at the largest smear size of 5, and for the largest number of hypersweeps N=50, to be an upper bound on the 0^{++} mass. We obtain

$$m(0^{++}) = (300 \pm 70) \cdot \Lambda_L \tag{3.5.1}$$

where Λ_L is the dimensional renormalisation group invariant, related to the lattice spacing by Eq. (4.2.1). A well-documented review by Berg [10] indicates that the different methods of calculating the 0^{++} glueball mass on a lattice are all producing estimates roughly in agreement at

$$m(0^{++}) = (280 \pm 50) \cdot \Lambda_L \tag{3.5.2}$$

These figures are certainly consistent with our results. Using the scale that was set in the

string tension discussion of Chapter 1, our result corresponds in physical units to a mass

$$m(0^{++}) \approx 1400\ MeV \tag{3.5.2}$$

Although none of the glueball spectrum calculations are yet accurate enough to warrant a detailed comparison with experiment, this value agrees with other lattice calculations performed at weaker (bare) couplings, and at couplings off the Wilson axis [10,11].

References

[1] see e.g., J. Kogut, Rev. Mod. Phys. 51 (1979) 659.

[2] C. Michael, I. Teasdale, Nucl. Phys. B215 (1983) 433;

A. Konig et al., Phys. Lett. 138B (1984) 410;

P. de Forcrand et al., Phys. Lett. 148B (1984) 140;

P. de Forcrand et al., Phys. Lett. 160B (1985) 137.

[3] K.G. Wilson, Remarks at Abingdon Meeting on Lattice Gauge Theories (March, 1981).

[4] B. Berg, A. Billoire, C. Rebbi, Ann. Phys. 142 (1982) 185;

M. Falcioni et al., Phys. Lett. 110B (1982) 295;

K. Ishikawa, G. Schierholz, M. Teper, Phys. Lett. 110B (1982) 399;

A. Patel, R. Gupta, Phys. Lett. 138B (1984) 296.

[5] K. Ishikawa et al., Z. Phys. C21 (1983) 167.

[6] G. Parisi, R. Petronzio, F. Rapuano, Phys. Lett. 128B (1983) 418.

[7] B. Berg A. Billoire, Nucl. Phys. B226 (1983) 405.

[8] B. Berg, A. Billoire, Nucl. Phys. B221 (1983) 109.

[9] S. W. Otto, P. Stolorz, Phys. Lett. 151B (1985) 428.

[10] B. Berg, in: *Progress in Gauge Field Theory* (Cargese, 1983), eds. G. 't Hooft et al. (Plenum, 1984).

[11] A. Patel et al., USCD/HUTP/UW preprint UCSD-10P10-260 HUTP-86/A035 UW/40048-06 P6.

Figure Captions

(3.1) The different loop operators A_i used in the construction of the glueball wave function (see Eq. (3.2.3)) for the Monte Carlo Variational Method.

(3.2) The rotationally invariant operator for the simple plaquette, which transforms as 0^{++} about the site n. A smear is obtained by summing this operator over all the sites in the smear.

(3.3) Sequence of configurations corresponding to Eq. (3.4.1), for a given pair of 2-dimensional smears $O(E)$ and $O(F)$ with fixed environments E and F. The parts of the environments in the time direction are not shown.

(3.4) Sequence of configurations corresponding to Eq. (3.4.2), in which the spatial part of each smear's environment now changes. The parts of the environments in the time direction (not shown) remain constant.

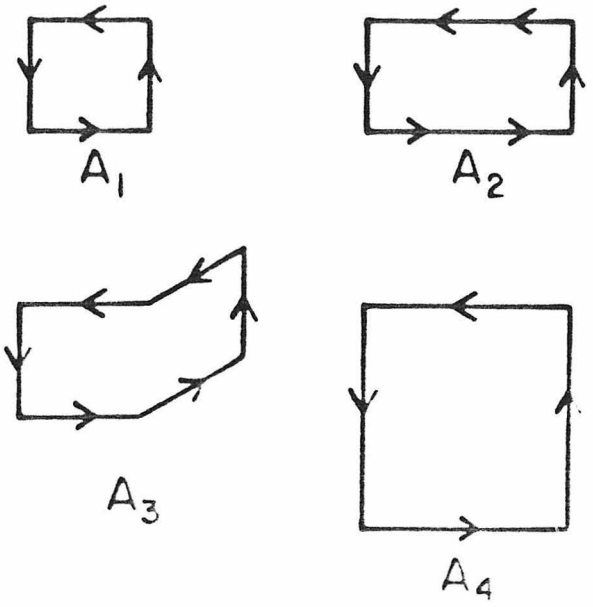

Fig. 3.1

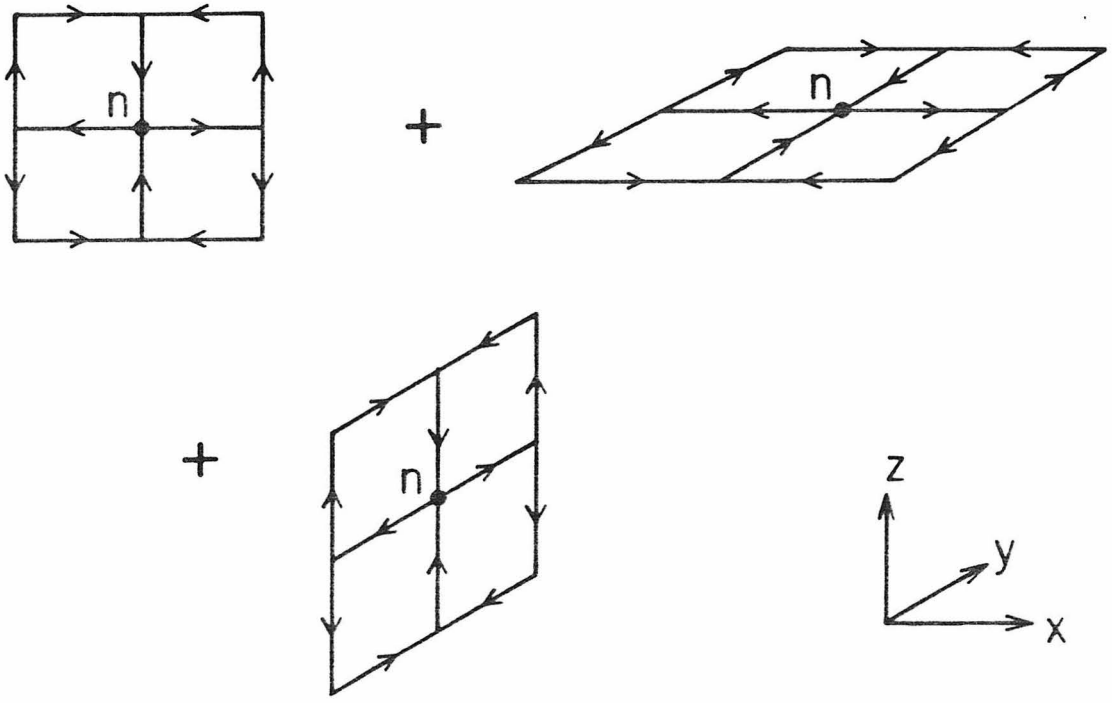

Fig. 3.2

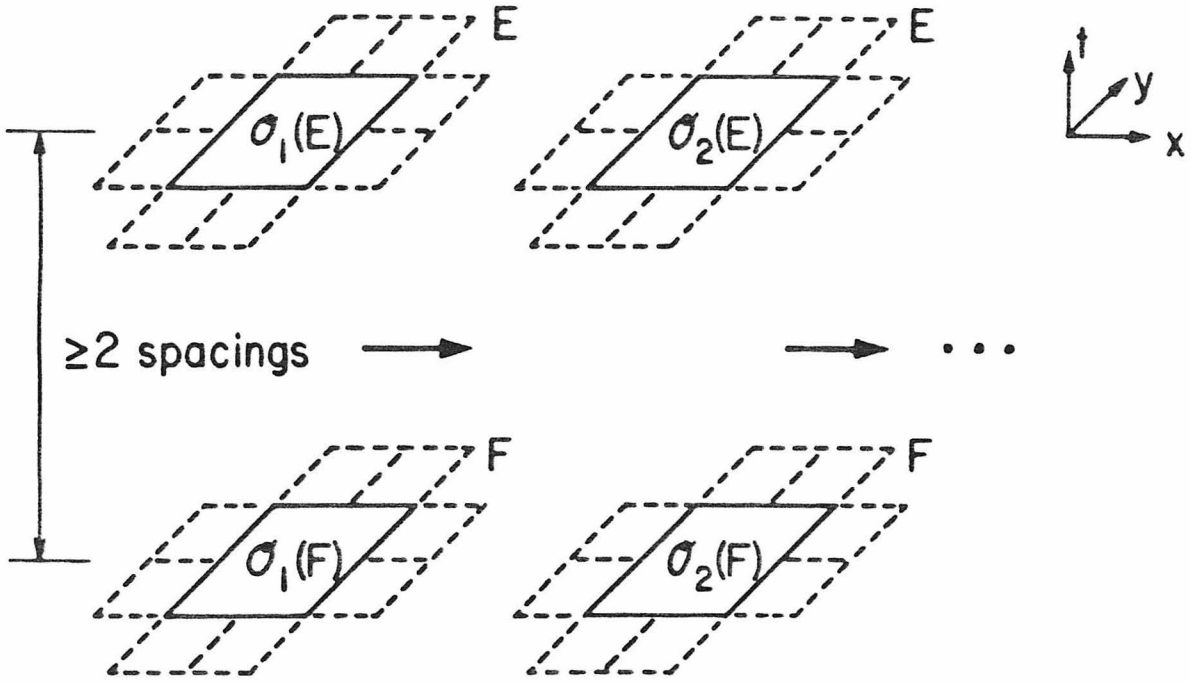

Fig. 3.3

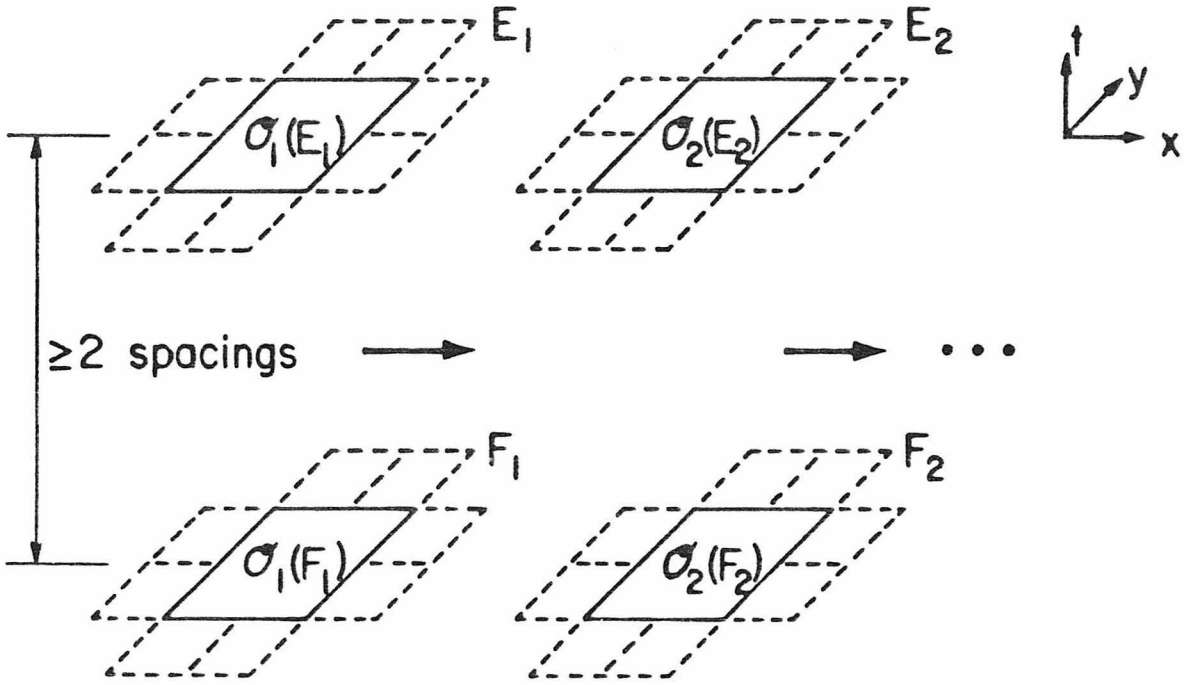

Fig. 3.4

Estimates for the mass gap ma at β=5.7 for 3 different values of hypersweep number N. The dashes indicate that the results at those values were not statistically significant. The first set of values for N=0 and N=10 are identical because they were measured on the same computer run, and the averaging trick described in the text cannot be applied for correlation distances less than 2 lattice spacings.

N	Time slice	smear size 2	smear size 3	smear size 4	smear size 5
0	1	2.33 ± 0.01	2.10 ± 0.01	1.97 ± 0.01	1.90 ± 0.01
	2	1.73 ± 0.03	1.55 ± 0.03	1.35 ± 0.03	1.25 ± 0.04
	3	-	-	-	-
	4	-	-	-	-
10	1	2.33 ± 0.01	2.10 ± 0.01	1.97 ± 0.01	1.90 ± 0.01
	2	1.69 ± 0.03	1.51 ± 0.03	1.40 ± 0.03	1.32 ± 0.04
	3	1.49 ± 0.02	1.34 ± 0.02	1.24 ± 0.03	1.16 ± 0.04
	4	-	-	-	-
50	1	2.36 ± 0.01	2.15 ± 0.01	2.04 ± 0.01	1.98 ± 0.01
	2	1.73 ± 0.02	1.54 ± 0.02	1.42 ± 0.02	1.34 ± 0.04
	3	1.57 ± 0.04	1.40 ± 0.04	1.26 ± 0.05	1.17 ± 0.05
	4	2.05 ± 0.22	1.33 ± 0.25	1.18 ± 0.25	0.99 ± 0.23

Table 3.1

Chapter 4: The Microcanonical Renormalisation Group

Introduction

As we have seen in Chapter 1, a continuum quantum field theory (QFT), as defined by a Euclidian path integral, corresponds to the critical point of a statistical system, the latter being obtained by regularising the QFT by defining it on a lattice. When regularised on a lattice of given spacing, the theory is defined by a partition function which in general takes the form

$$Z=\int dU e^{-S(U)}$$

where

$$S(U)=\sum_{\alpha}\beta_{\alpha}S_{\alpha}(U)$$

the β_α being a set of couplings for the different actions $S_\alpha(U)$. The central issue in the critical phenomena approach to field theories is to understand how these parameters $[\beta_\alpha]$ of the theory must be tuned in order to approach the critical point. This behaviour is governed by the renormalisation group (RG), which is the focus of investigation in this chapter.

The renormalisation group approach to critical phenomena has already been remarkably successful at calculating important physical properties, such as the critical exponents of thermodynamic functions, that are not calculable by any other technique [1]. The method is based upon the realisation that the huge number of degrees of freedom that interact in a statistical system can be reduced by averaging them over small areas in configuration space to give a smaller set of effective degrees of freedom. This process is called a renormalisation group (RG) transformation. The interactions between these new effective degrees of freedom are described by a new set of couplings β_α^*. The partition function at this larger distance scale is thus given by

$$Z=\int dU^* e^{-S^*(U^*)}$$

where

$$S^*(U^*)=\sum_\alpha \beta_\alpha^* S_\alpha(U^*)$$

For simple theories, the new effective couplings are related to the original ones by a calculable scaling function. If one then assumes that the correlation length remains constant in physical units under the RG transformation, the critical exponents can be easily extracted.

Unfortunately, the RG apparatus is not as easy to apply in the case of gauge theories as it is for traditional statistical physics problems. The scaling relations for gauge theories are not calculable exactly. Nevertheless, the basic approach is still fruitful [2]. Starting with an action containing only local interactions defined at some lattice spacing a, one can apply an appropriate RG transformation to generate a new effective action on a lattice that is larger by some factor $\lambda>1$. Applying this RG transformation many times, one obtains a trajectory in the multi-parameter coupling space

$$[\beta_\alpha^{(0)}] \rightarrow [\beta_\alpha^{(1)}] \rightarrow [\beta_\alpha^{(2)}] \rightarrow \ldots.$$

Eventually, after many transformations, this process of renormalisation will produce non-zero couplings for the non-local interactions, so one could in principle discover the behaviour of QCD at large distance scales, given its behaviour at small distances. Unfortunately, Monte Carlo methods are the only way of implementing this renormalisation process at *all* couplings, and the problem is too vast in scope to be practicable on present-day computers.

Instead, we look at the more limited, but still important, problem of choosing an action with which to conduct Monte Carlo simulations that is as deep in the continuum region (i.e., as close to the critical point) as possible. The most naive way to do this is to conduct simulations at high values of β. The problem here is that the correlation length increases as one goes to higher β. The variables therefore show more long-range order, and it takes much longer to adequately traverse phase space in the simulation. This phenomenon is known as critical slowing down. In order to avoid this, we go back to the renormalisation process and look at the

couplings after just a few RG steps have been applied. At each step, the correlation length in units of the lattice spacing is reduced by a factor of λ, while the long-range properties of the original action are preserved. Hence, if we use the effective action obtained after N renormalisation steps to do simulations, our results will correspond to those obtained with the original action, which has a correlation length λ^N times bigger. We are therefore probing more deeply into the continuum limit with the effective action without paying the price of critical slowing down, and the scaling behaviour of the observables willl be improved.

The most powerful way of implementing renormalisation on a lattice is with the Monte Carlo Renormalisation Group [2,3,4,5]. In principle, all the renormalisation properties of QCD could be found by this method, given unlimited computer power, although of course, this is not really possible in practice. The method consists of generating a set of lattice configurations at a given set of couplings $\beta_\alpha^{(0)}$ using the usual Monte Carlo technique, based upon the canonical ensemble. The RG transformation then consists of averaging out, or "blocking," small groups of links to leave configurations of fewer links. This has the effect of removing short-distance fluctuations, i.e., integrating out high momenta. These new blocked configurations will be distributed according to a new effective action, with unknown effective couplings $\beta_\alpha^{(1)}$. The expectation values of various terms in the action are measured for these new blocked configurations.

Our work consists of now using a method based on the microcanonical ensemble to directly *measure* the new effective couplings, using these blocked expectation values as input. In this way the renormalised couplings are tracked directly under RG transformations, enabling improved actions to be found.

In Section 1, we describe some of the features of the the renormalisation group and critical phenomena that are relevant to the problem of renormalising a gauge theory like QCD. In Section 2, we discuss the basic principles of the Monte Carlo Renormalisation Group. In Section 3, we describe the various methods of simulation based on the microcanonical ensemble. In Section 4, we discuss the differences in the results produced by microcanonical and

canonical simulations. In Section 5, we outline our renormalisation calculation, and in Section 6 we present our results for the "improved" action.

1. Critical phenomena and the renormalisation group

Let's first review the important aspects of critical phenomena that relate to our calculation. We have already described how RG transformations generate a sequence of actions in a multi-dimensional coupling space. In the theories of interest, this sequence has a fixed point $\beta_\alpha^{(*)}$. If one is in the region of attraction of this fixed point, the effect of an RG transformation is always to drive the point in coupling space closer to the fixed point. The correlation length ξ must either diverge or be zero at such a point. We are interested in the cases where $\xi=\infty$: that is, where the fixed point is a critical point. This corresponds to the continuum limit in the case of field theories.

In general, there is actually a critical surface in the coupling space upon which the correlation length diverges (see Fig. 4.1). If one begins at a point slightly off this surface, i.e., starting with some small non-zero lattice spacing a, the effect of RG transformations is to reduce the correlation length (in lattice units) as the lattice spacing grows. Thus one moves further away from the critical surface. Some examples of such trajectories are shown in Fig. 4.1. At first the trajectories will move almost parallel to the critical surface, but after many RG transformations, each trajectory will turn and converge along a direction orthogonal to the critical surface (there may be more than one such direction).

The interactions which correspond to these directions of instability are called relevant. In field theories they are the renormalisable interactions. The couplings corresponding to the remaining interactions are either attracted by the fixed point, in which case they are called irrelevant, or ignore the fixed point, in which case they are called marginal. The irrelevant interactions have no effect on processes at large (i.e., physical) distances, since after several RG transformations (integrations of short-distance behaviour), they play no further role in the RG trajectory of the action.

If there are n relevant operators, then n parameters must be adjusted in order to approach the critical surface. In the case of QCD, it is known that only one parameter, the bare coupling g, need be adjusted in order to approach the critical point at $g=0$. Therefore, in this case there must be just one direction of instability, which we call the renormalised trajectory (RT). An action which sits on this trajectory is the ideal action to use in Monte Carlo simulations, since it contains no irrelevant short-distance operators at all!

It is our goal in this chapter to locate this trajectory by measuring the flow of various actions towards it under RG transformations. Of course, the exact position of the fixed point, and indeed the exact form of the RT, depend on the form of the RG transformations, and are therefore not in themselves physical characteristics of the theory. Nevertheless, the different RG transformations between them supply us with an envelope of RT's within which the actions are significantly closer to the continuum behaviour than naive ones.

2. Monte Carlo Renormalisation Group

This powerful numerical technique has been developed by several authors as a way of studying the rescaling of theories in regions of intermediate coupling [2,3,4,5]. To see how it works, consider a general lattice action involving a family of couplings β_α defined on a lattice of spacing a:

$$S(U)=\sum_\alpha \beta_\alpha S_\alpha(U) \tag{4.2.1}$$

The partition function is

$$Z=\int dU e^{-S(U)} \tag{4.2.2}$$

One generates a set of configurations, using Monte Carlo methods, with this partition function. The degrees of freedom are then thinned by averaging the link variables of these configurations over small blocks of variables. This corresponds to integrating out the original variables U to leave a new set of variables U^* as follows:

$$\int dU P(U^*,U) e^{-S(U)} = e^{-S^*(U^*)} \tag{4.2.3}$$

where

$$S^*(U^*) = \sum_\alpha \beta_\alpha^* S_\alpha(U^*) \tag{4.2.4}$$

The function $P(U^*,U)$ describes how the new variables U^* are to be constructed from the old ones U. One now has a theory defined on a blocked lattice by the "new" partition function

$$Z = \int dU^* e^{-S^*(U^*)} \tag{4.2.5}$$

This will be the same as the old partition function at large distances, and hence describe the same physics, provided the function $P(U^*,U)$ satisfies the constraint

$$\int dU^* P(U^*,U) = 1 \tag{4.2.6}$$

More precisely, what happens in the MCRG is that the new configurations generated by the blocking transformations will be distributed according to the new action Eq. (4.2.4). This means that the ensemble averages of the blocked configurations will automatically be the correct estimates for the partition function with the new action (of course, one does not yet know what the values of the β_α^* specifying the new action are!). In typical applications of the MCRG, one performs several blockings on sets of successively smaller lattice configurations. The idea then is to extract information about the flow of couplings at each blocking level.

There are many different forms that the blocking transformations can take. The original work on the MCRG used a scale factor of two [2,3,4,5]. Subsequent authors have developed a method which uses a finer scale factor of $\sqrt{3}$ [6]. In this chapter we will measure the effective actions generated as this $\sqrt{3}$ blocking transformation is applied. Briefly, the transformation works as follows. Consider a 3-dimensional cube within the lattice, as shown in Fig. 4.2, with sites i and j at opposite corners of the cube. A new link variable on the blocked lattice is associated with the diagonal ij, shown emboldened in the figure. The value of the new

blocked link is obtained by selecting a 3-link path between i and j, and taking the product of the link matrices along that path. An example of one such path is shown in Fig. 4.2. One finds this product for each of the 6 such paths connecting i and j, calculates their sum, and then projects the resulting matrix back onto the SU(3) manifold.

The details of the method can be found in [6]. Notice that the transformation has the feature that each of the sites in the blocked lattice is a site in the old lattice, and that the blocked links are formed by continuous paths of the old links between the blocked sites. This relieves us of the need to fix a gauge, since it means that the transformation is gauge-invariant.

By applying this blocking procedure to a set of configurations generated at some initial couplings, one obtains a set of configurations which are distributed according to some new effective action. The computable quantities on these blocked lattices are the ensemble averages of various operators, not the new effective couplings themselves. To extract the values of these renormalised couplings from the blocked ensemble averages, we use a technique based upon the microcanonical ensemble. We now describe how this method works.

3. Microcanonical simulation

Recall that the microcanonical ensemble is obtained by constraining the system of interest to have a constant energy, in contrast to the canonical ensemble, which operates by allowing the system to exchange energy with a heat reservoir at a constant temperature. Of course, the microcanonical ensemble offers a perfectly good way, in its own right, of conducting all Monte Carlo simulations. In the thermodynamic limit it is essentially equivalent to the canonical ensemble, and hence must obtain the same physical results in this limit. In the particular problem that we want to look at, it just happens to give a particularly efficient way of computing the renormalisation flow of our theory. Let's look, then, at the general principles underlying this method of simulation, and at how they differ from those based on the canonical ensemble in the non-thermodynamic region to which computers have access.

One possible method has been developed and used extensively in molecular dynamics simulations, and has recently been applied to the case of gauge theory simulations [7], especially those involving fermionic degrees of freedom [8]. One introduces a constant to the action which is independent of all the link variables U_l. Thus, the usual canonical ensemble average for some operator O is rewritten as

$$<O>=\frac{1}{Z}\int\prod_l dU_l\int\prod_l dp_l\, O(U_l)e^{-\beta(S+K)} \tag{4.3.1}$$

where

$$K=\frac{1}{2}\sum_l p_l^2 \tag{4.3.2}$$

and

$$Z=\int\prod_l dU_l\int\prod_l dp_l\, e^{-\beta(S+K)} \tag{4.3.3}$$

In these expressions, S is the standard action that would be used with the canonical ensemble. The new constant term K can be thought of as the kinetic energy of a classical system with potential energy S. Now consider constraining the total energy H of this classical system to be a constant value E. The system can then be allowed to evolve in a new dimension which we label by τ. The p_l can now be thought of as the canonically conjugate momenta of the link variables U_l:

$$\frac{dU_l}{d\tau}=p_l \tag{4.3.4}$$

Hamilton's equations in this fictitious time τ now give

$$\frac{d^2U_l}{d\tau^2}=\dot{p}_l=-\frac{\partial S(U)}{\partial U_l} \tag{4.3.5}$$

In a computer simulation, the new procedure therefore consists of using the new partition

function

$$Z=\sum_{C}\sum_{P}\delta[S(C)+K(P)-E] \tag{4.3.6}$$

where each configuration of the link variables C and of the canonically conjugate momenta P is arrived at by numerically solving the ordinary differential equations Eq. (4.3.5). The "temperature" to which the energy E corresponds is then given by twice the average value of the kinetic energy, as in the standard treatment of classical statistical mechanics (the equipartition principle).

The stochastic methods based on the canonical ensemble have hence been replaced by a deterministic procedure, in which randomness is effectively generated by the complexity of the configuration space. This method has been shown to be a successful way of conducting Monte Carlo simulations in general [7,8]. However, for our purposes it has the drawback of introducing the coupling as a function of the average kinetic energy, which is a global property of the canonically conjugate momenta. It is therefore not possible to measure the coupling for more than a single term in the action, which is what we want to do in MCRG investigations.

The solution is to use a different version of the microcanonical method, which could be called pseudo-deterministic. The technique, developed by Creutz [9], has previously been applied to renormalisation studies of both the Ising [10,11] and O(3) [12,13] models. The method consists of a random walk upon a constant energy surface in an enlarged configuration space. This space is comprised of the lattice system of interest, plus an extra degree of freedom, called a *demon*, that jumps around the system transferring energy between the actual link variables of the lattice. It is the energy of the combined lattice/demon system that is kept constant. Thus, the partition function is defined by

$$Z=\sum_{C}\sum_{E_d}\delta[S(C)+E_d-S^t] \tag{4.3.7}$$

where S is the action of a lattice configuration C, E_d is the demon energy, and S^t is the total action to which the lattice/demon system is constrained. This single demon degree of freedom thus takes the place of all the canonically conjugate momenta described in the molecular dynamics method of simulation.

Now, the entropy $\sigma(S)$ of the lattice system at some action S is defined to be the logarithm of the number of states available to the system with that action. Hence, the number of states in the microcanonical ensemble that contain a demon of energy E_d is given by

$$N(E_d)=e^{\sigma(S^t-E_d)} \tag{4.3.8}$$

Provided that E_d is small compared to S^t, the right-hand side of this equation can be expanded in a Taylor series,

$$N(E_d)=e^{\sigma(S^t)-\frac{\partial\sigma}{\partial S}]_{S^t}\cdot E_d+\frac{1}{2}\frac{\partial^2\sigma}{\partial S^2}]_{S^t}\cdot E_d^2+\cdots} \tag{4.3.9}$$

But the coefficient of E_d in the above expression is precisely the inverse temperature β of a lattice system whose internal energy is S^t:

$$\beta=\frac{\partial\sigma}{\partial S}]_{S^t} \tag{4.3.10}$$

The coefficient of E_d^2 is essentially the inverse specific heat of the lattice system, since

$$\frac{\partial^2\sigma}{\partial S^2}]_{S^t}=\frac{\partial\beta}{\partial S}]_{S^t}=-\frac{1}{T^2}\cdot\frac{\partial T}{\partial S}]_{S^t} \tag{4.3.11}$$

where the temperature T is $1/\beta$. Thus, the demon's distribution will go as

$$P(E_d)=e^{-\beta E_d-\frac{\beta^2E_d^2}{2CV}+O(1/V^2)} \tag{4.3.12}$$

where C is the specific heat of the system per unit volume V, namely

$$C=\frac{1}{V}\cdot\frac{\partial S}{\partial T}]_{S^t} \tag{4.3.13}$$

Of course, this is essentially the way in which the canonical ensemble is usually derived from the microcanonical. The difference here is that the lattice system itself now forms the heat reservoir for the demon, rather than the lattice being at thermal equilibrium within a large heat reservoir. The demon is allowed to become in thermal equilibrium with the lattice, and its distribution then gives a measure of the inverse temperature of the larger lattice system. This is the important feature of the demon method, for it allows us to probe the temperature of a lattice system whose observable averages have already been found by other methods. Notice that if the demon energy is not kept small, it can remove a substantial amount of energy from the lattice system, and will then be sampling the temperature of a system which is actually somewhat cooler than S^t. This is reflected in the $1/V$ corrections in Eq. (4.3.12), which must be taken into account unless $d \ll S^t$.

It turns out for our lattices that a suitable demon domain is [-1,1]. The $1/V$ terms in Eq. (4.3.6) can then be neglected, and the inverse temperature can be easily extracted from the ensemble average of the demon energy,

$$<E_d> = \frac{\int_{-1}^{1} E_d e^{-\beta E_d} dE_d}{\int_{-1}^{1} e^{-\beta E_d} dE_d} = \frac{1}{\beta} - \frac{1}{\tanh\beta} \tag{4.3.14}$$

Our interest currently is in sampling the temperature couplings of certain lattice systems. However, it is perfectly possible to use the method to generate many lattice configurations in order to measure average correlation functions as conventional canonical methods do. A microcanonical simulation conducted at action S^t will correspond to a canonical one conducted at the inverse temperature β given by Eq. (4.3.10). There will, however, be finite volume corrections which will make results differ slightly. Some of these corrections occur even in the limit of zero demon energy, where energy is transferred directly from one part of the lattice to another. These corrections are discussed in more detail in Section 4.

The particular beauty of the demon technique is that one can use it to measure many different couplings simultaneously. This is not possible for the molecular dynamics technique discussed previously. We can introduce a demon variable, labelled by the index i, for each operator of interest. The total partition function is then

$$Z = \sum_{E_{d_i}} \sum_{C} \prod_{i} \delta(S_i(C) + E_{d_i} - S_i^t) \tag{4.3.15}$$

where the subscript i has been introduced to label the lattice action $S_i(C)$, the demon energy E_{d_i}, and the constant constraining action S_i^t for each demon. Each demon's energy is then related to the corresponding coupling β_i by the obvious generalisation of Eq. (4.3.14). The method is therefore ideal for our purposes, in which even a simple initial action can generate many interactions after a few renormalisation steps.

Briefly, the algorithm proceeds by looking at each link variable in the lattice in turn, and attempting to randomly change the value of the variable. Every such attempt results in a new set of tentative demon values E_{d_i}, in order to keep the S_i^{tot}s constant. The attempted change is accepted only if all the new demon values remain within their predetermined bounds. In our case, all the demons were kept within the interval [-1,1]. The demon variables were then averaged over to give the various couplings via Eq. (4.3.14).

It should be pointed out here that there is a potential difficulty associated with this method of simulation. In the stochastic methods described earlier, it is possible to prove rigorously that a Markov chain of configurations, plus detailed balance, guarantees that the ensemble generated follows the desired stable Boltzmann distribution, so that ensemble averages are physically meaningful. The microcanonical demon method, however, is essentially a deterministic technique. In the case of the simulation of the Ising model, for example, the lattice/demon configuration at a given time completely determines the next lattice/demon configuration in the sequence. It is hoped that the complexity of the configuration space is enough to ensure the ergodicity of the simulation: i.e., that all regions of the available state

space are evenly scanned by the process. This ergodicity has not been proved, although some work has been done towards it [14].

In order to deal with this problem, numerical simulations have been performed using both stochastic and deterministic methods [10,12]. In the case of the Ising model, it is even possible to compare both the simulations with analytic results [10]. All the work done so far has produced results in which the two methods agree to within statistical errors. On this basis it seems reasonable that the demon method can be relied upon to give sensible estimates of observables.

4. Microcanonical versus canonical averages

With the finite size lattices that must be used in computer simulations, the canonical and microcanonical ensembles will produce different ensemble averages. It is important to understand these discrepancies between the two ensembles, especially since we are switching between the two in the course of the calculation. To look at these various corrections, let's suppose we have computed a set of various ensemble averages in the canonical ensemble, for a system of size V, at inverse temperature β. There are basically 3 ways in which the microcanonical ensemble averages will differ from these quantities.

Firstly, since the demon energy E_d is non-zero, and since the volume is finite, the demon will not be distributed in an exactly Boltzmann manner, but rather in the form Eq. (4.3.12) discussed previously. This leads to inverse volume corrections to the formula Eq. (4.3.14) used to relate the demon averages to the coupling β [10,13,15]. Label with the subscript ∞ all the demon averages obtained when Eq. (4.3.12) is taken to only the first power in E_d (i.e., the infinite volume case). Then the demon averages that one obtains when going to the second power in E_d (i.e., the finite volume case) are given [13] by

$$<E_d>_{finite}=<E_d>_\infty+\frac{1}{2CV}(<E_d>_\infty<E_d^2>_\infty-<E_d^3>_\infty \quad) \tag{4.4.1}$$

Secondly, suppose we were interested instead in the average values of lattice observables, rather than the lattice temperature. If a canonical average $<O>_{can}$ is used as the constant total energy S^t for a microcanonical simulation, then the lattice average $<O>_{micro}$ will, of course, differ from the desired value by the average demon energy $<E_d>$.

Thirdly, there is a correction unrelated to the demon size. It would occur even in a simulation in which energy was transferred directly from one part of the lattice to another, and applies to observables other than the energy observable used to define the microcanonical ensemble itself. An observable O_{can} calculated in the canonical ensemble at inverse temperature β is related to the observable O_{micro} calculated in the microcanonical ensemble on a lattice held at energy E by

$$O_{can}(\beta)=\frac{\int dE\,\rho(E,V)O_{micro}(E)e^{-\beta E}}{\int dE\,\rho(E,V)e^{-\beta E}} \tag{4.4.2}$$

where $\rho(E,V)$ is the density of states at energy E and volume V. The integrands in this expression are strongly peaked around the average energy $<E>_{can}$ in the canonical ensemble. Therefore, expand O_{micro} about this energy:

$$O_{can}(\beta)=O_{micro}(<E>_{can})+\frac{1}{2}<E^2-<E>^2_{can}>_{can}[\frac{\partial^2 O_{micro}(E)}{\partial E^2}]_{E=<E>_{can}}+.... \tag{4.4.3}$$

The coefficient of the second term is, again, essentially C, the specific heat per unit volume. If we rewrite this expression in terms of the energy density $e=E/V$, we have [10,15]

$$O_{can}(\beta)=O_{micro}(<e>_{can})+\frac{C}{2\beta^2 V}[\frac{\partial^2 O_{micro}(e)}{\partial e^2}]_{e=<e>_{can}}+O(\frac{1}{V^2}) \tag{4.4.4}$$

Hence these corrections also go as the inverse volume (and higher orders).

As mentioned before, an extensive study has been performed of these three errors in the case of the 2-dimensional Ising model, where analytic expressions are available for the O_{can}'s and their derivatives. Numerical simulations confirm the expected discrepancies between the

ensembles [10], so one can be confident of the ergodicity of the simulations.

In the case of our renormalisation investigations, only the first source of error need be considered at all, since we are merely looking at the couplings of a certain Boltzmann distribution: the averages of the lattice observables themselves in the microcanonical simulation are of no interest. Further, our volumes are such that this error is negligible compared to the statistical errors of the calulation.

5. Microcanonical renormalisation group for SU(2)

We have chosen to look at the renormalisation flow of the couplings for the group SU(2) when the MCRG is applied to it [20]. This doesn't tell us anything directly about the behaviour of QCD, which is based upon the group SU(3). However, the theories based on the 2 groups are known to have similar phase structures, and both have the important property of asymptotic freedom. Therefore it is expected that the renormalisation flows for the 2 groups in the intermediate coupling region will behave in qualitatively similar ways. In particular, what we are trying to do is to find ways of more closely approaching the continuum behaviour of SU(2) with numerical techniques. The improvements that are made in this area can then certainly be applied to the case of SU(3). However, this is computationally a much more intensive task, so we have concentrated in this work on looking at the efficiency and power of the renormalisation process itself.

The specific calculation that we describe here builds upon work done by Gupta and Patel [16], in which the renormalised couplings were measured after a single application of the $\sqrt{3}$ blocking transformation. In later work by Gupta, Patel and Otto [17], a set of 210 lattice configurations were generated on 18^4 lattices at several different initial couplings. The $\sqrt{3}$ transformation was then applied twice to each set of configurations, and the block expectation values of various Wilson loops on the resulting twice-blocked 6^4 lattices were calculated. These block expectation values formed the input to our microcanonical simulation, in which we measured the couplings to which they correspond.

The calculation has three main aims. Firstly, we want to know whether the demon method remains stable as several different couplings are measured. Conceivably, one could get stuck in small areas of configuration space because of the number of demon bounds that must be satisfied simultaneously. The resulting averages would fluctuate wildly and would be physically useless. We have measured up to 4 couplings simultaneously (the limit is due to the constraints of computer time), and have found this pathological behaviour to be absent.

Secondly, it is important to see how sensitive the values of the renormalised couplings are to the number of couplings that are measured. Gupta and Patel found, using a more indirect method, that the couplings do change as the number included in the analysis is varied. Briefly, we find that the same is also true when the couplings are measured using the microcanonical demon method.

Thirdly, we want to be able to measure the renormalised couplings after more than one blocking transformation. In their work, Gupta and Patel [16] assumed a linear transformation relating the observables on the initial and once-blocked lattices. Using this transformation, they were able to determine the couplings after one blocking step. The microcanonical method, on the other hand, is capable, in principle, of measuring the couplings straightforwardly over an arbitrary number of blocking steps. In our calculation the statistics allowed us to easily measure the coupling flow over 2 blocking steps, and hence we were able to obtain an improved action, and to confirm the general behaviour observed by Gupta and Patel.

The calculation was implemented by choosing each S_i^{tot} to be the ensemble average of the action S_i over all 210 blocked 6^4 lattices, as calculated in [17]. The computation was then carried out on 6^4 lattices also, so that finite size effects in the two cases remained the same. Each lattice/demon system was started off with all the energy being contained in the demons, i.e., with $E_{d_i}=S_i^{tot}$, and with $<S_i>=1-<O_i>=0$, $<O_i>$ being the averages of the Wilson loop operators. The system was thermalised by only allowing changes that transferred energy from the demon to the lattice, until all the demons were within the interval [-1,1]. All the demons were then allowed to roam at will throughout this interval, and their ensemble averages $<E_{d_i}>$

were measured.

We looked at 2 subspaces of the general lattice action

$$S=\beta_F\sum TrU_p+\beta_{6p}\sum TrU_{6p}+\beta_A\sum\left[4/3(TrU_p)^2-1/3\right]+\beta_{3/2}\sum\left[2(TrU_p)^3-TrU_p\right] \quad (4.5.1)$$

In this expression, U_p denotes the product of link matrices forming a simple plaquette (in the fundamental representation). U_{6p} denotes the product of matrices forming a planar loop of perimeter 6 links (also in the fundamental representation). The other operators included are the $J=1$ (adjoint) and $J=3/2$ representations of the simple plaquette (for SU(N) groups, these can be easily expressed in terms of the fundamental representation).

The reason for including higher-spin representations in the action is that previous renormalisation group studies have shown them to play an important role in the coupling flow in the strong coupling region. Chief among these is the Migdal-Kadanoff approximate renormalisation scheme, which has been used to map out the coupling flow at strong (bare) couplings [18] (see Fig. 4.3). The trajectories have been found to converge to a line which is roughly described [18] by

$$\beta_A=-0.24\beta_F. \quad \beta_{3/2}=0.12\beta_F \quad (4.5.2)$$

However, the scheme is flawed, because it always projects onto simple plaquette actions, and therefore misses any long-distance interactions. It fails, for instance, to give the order of phase transitions correctly. Attempts have therefore been made to approach the problem from the weak coupling region, by expanding the lattice action in powers of the lattice spacing a, and by then cancelling off the terms with powers larger than a^4. This perturbative improvement program, initiated by Symanzik [19], requires the introduction of longer-range interactions, of which the planar perimeter-6 loop is an example. With more computer time, it would of course be desirable to also look at the effects of larger loop operators also. Hopefully, our general action is flexible enough to be able to take account of both non-perturbative (higher-spin representations) and perturbative (larger loops) effects in one fell swoop.

6. Results

We first measured the couplings in the $[\beta_F,\beta_A,\beta_{6p}]$ subspace. All the initial couplings on the 18^4 lattices were on the Wilson axis, with the exception of one coupling along the Migdal-Kadanoff (MK) line $\beta_A=-0.24\beta_F$ in Eq. (4.5.2). The results are shown in Table 1. From the table, it can be seen that the couplings have settled along the trajectory $\beta_A/\beta_F=-0.16$. The especially encouraging thing here is that this has happened for both the MK and Wilson couplings. This is important, because these two actions come from different sides of the RT, so we have good evidence that two blockings have been sufficient to locate the RT. It can also be seen that at weaker couplings, we approach the line $\beta_{6p}/\beta_F=-0.04$. This is slightly higher than the ratio of -0.05 which the perturbative Symanzik program obtains at weak couplings [19].

We then measured the couplings for the enlarged subspace $[\beta_F,\beta_A,\beta_{6p},\beta_{3/2}]$. The results can be found in Table 2. They show that the introduction of the extra operator has a strong effect on the trajectory of the couplings. We now obtain $\beta_A/\beta_F\approx-0.21$, a substantially different ratio than that obtained with the smaller subspace. This new ratio also lies close to the MK line. Notice that the importance of the new $J=3/2$ operator cannot be judged by the value of the $\beta_{3/2}/\beta_F$ ratio, which is itself quite small. Notice also that the $6p$ ratio is unchanged by the addition of the 3/2 operator. This decoupling behaviour has been observed previously [16,17].

Let's now see how our results compare with the couplings obtained after 1 blocking step. The once-blocked couplings have been included in the tables. They show that in the $[\beta_A,\beta_F]$ subspace, the trajectories from all the starting actions have converged more strongly towards a single line, the RT. The change in most of the ratios is of the order of 10%. However, the trajectories in the $[\beta_{6p},\beta_F]$ subspace are not at all stable, becoming somewhat lower after the second blocking step in the case of the stronger couplings. Why is this so?

In order to examine this question, we need to digress for a moment to see how the effects of the RG transformation are described. Assuming we are close to the fixed point of

the theory, the couplings at the $n+1^{th}$ blocking level can be related to the couplings at the n^{th} level by the linearised relation

$$\beta_i^{(n+1)}=\beta_i^* + \sum_j T_{ij}(\beta_j^{(n)}-\beta_j^*) \tag{4.6.1}$$

where

$$T_{ij} = [\frac{\partial \beta_i^{(n+1)}}{\partial \beta_j^{(n)}}]_{\beta^*} \tag{4.6.2}$$

Hence the difference of a coupling from its fixed point value when the RG transformation is applied is determined by the transformation matrix T. In the case of QCD, there is only one eigenvector of the matrix T with an eigenvalue $\lambda>1$. This means that there is just a single RT (specified in coupling space by this eigenvector), since after n applications of T, all other eigenvectors will be suppressed by the n^{th} power of their eigenvalues. These other eigenvectors correspond to the irrelevant operators discussed in Section 1, the largest eigenvector to the relevant operator.

The point now is that the work in [16] found the $6p$ operator to be the largest irrelevant operator. Since it is irrelevant, its effect will vanish after many blocking steps. However, its eigenvalue is about 1/3. Therefore, it will take several blockings before it converges to the RT. Despite this slow convergence at the stronger couplings, though, we do find that at weaker couplings the coupling stabilises to $\beta_{6p}/\beta_F \approx -0.04$, which is close to the value of -0.05 obtained in the Symanzik improvement program [19].

In conclusion, the demon method has been found to be a highly effective way of directly measuring renormalised couplings in MCRG calculations. At intermediate couplings, the RT is very similar to the line produced by a Migdal-Kadanoff analysis, while at weaker (bare) couplings, it shows the perturbative behaviour obtained in Symanzik's program. Our calculations indicate that the ideal action with which to conduct Monte Carlo simulations is along the line

$$\beta_A/\beta_F=-0.21 \quad \beta_{3/2}/\beta_F=-0.04 \tag{4.6.3}$$

It is not clear what the ideal value is for the ratio β_{6p}/β_F , due to the convergence problem mentioned above, although we do know that it is important to include the 6p operator in any improved action.

References

[1] K.G. Wilson, J. Kogut, Phys. Rep. 12C (1974) 75.

[2] K.G. Wilson, in: *Recent Developments in Gauge Theories* (Cargese, 1979), eds. G. 't Hooft et al. (Plenum, New York, 1980);

[3] L.P. Kadanoff, Rev. Mod. Phys. 49 (1977) 267.

[4] R.H. Swendsen, Phys. Rev. Lett. 42 (1979) 859;

[5] S.K.Ma, Phys. Rev. Lett. 37 (1976) 461.

[6] R. Cordery, R. Gupta and M. Novotny, Phys. Lett. 128B (1983) 425.

[7] D.J.E. Callaway and A. Rahman, Phys. Rev. Lett. 49 (1982) 613.

[8] J. Polonyi, H.W. Wyld, Phys. Rev. Lett. 51 (1983) 2257.

[9] M. Creutz, Phys. Rev. Lett. 50 (1983) 1411.

[10] G. Bhanot, M. Creutz, H. Neuberger, Nucl. Phys. B235 [FS11] (1984) 417.

[11] G. Bhanot, Phys. Lett. 154B (1985) 63.

[12] M. Creutz et al., Phys. Rev. Lett. 53 (1984) 875.

[13] M. Okawa, Phys. Rev. Lett. 54 (1985) 963.

[14] V.I. Arnold, A. Avez, *Ergodic problems of classical mechanics* (Benjamin, New York, 1968);

P.R. Halmos, *Lectures on ergodic theory* (Chelsea, New York, 1956).

[15] M. Creutz, in: *Gauge Theory on a Lattice*: 1984 (Argonne National Lab., April, 1984).

[16] A. Patel and R. Gupta, Nucl. Phys. B251 [FS13] (1985) 789.

[17] A. Patel, S. Otto and R. Gupta, Phys. Lett. 159B (1985) 143.

[18] K. Bitar, S. Gottlieb and C. Zachos, Phys. Rev. D26 (1982) 2853.

[19] K. Symanzik, in: *Proceedings of the Trieste Workshop on Nonperturbative Field Theory and QCD* (Dec. 1982), World Scientific 61.

[20] P. Stolorz, Phys. Lett. 172B (1986) 77.

Figure Captions

(4.1) Typical coupling flows under renormalisation transformations in a multi- dimensional coupling space $[\beta_1,\beta_2,\beta_3]$. Actions beginning *on* the critical surface converge to the fixed point (2 such trajectories are shown). Other actions move along trajectories (marked by an x after each transformation) that eventually converge along a single line, the renormalised trajectory (RT).

(4.2) The $\sqrt{3}$ blocking transformation. The diagonal ij is a link in the new blocked lattice. A typical contribution to it is the link product $U_1U_2U_3$. The new link ij is obtained by averaging this product with that of the 5 other 3-link paths from the site i to the site j.

(4.3) The SU(2) MK renormalisation trajectories projected onto the $[\beta_F,\beta_A]$ subspace (from ref [18]).

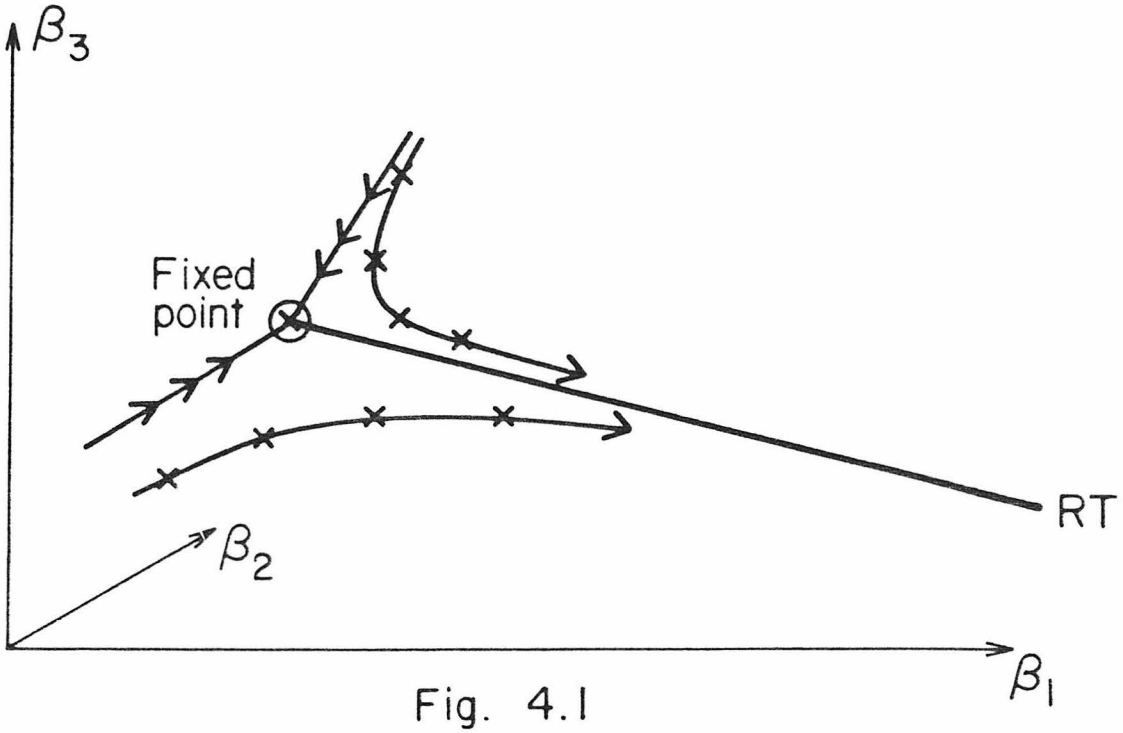

Fig. 4.1

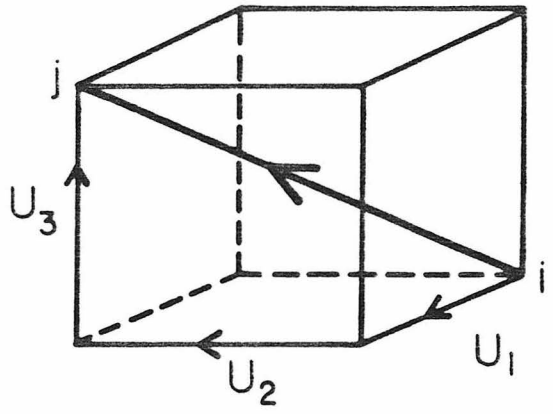

Fig. 4.2

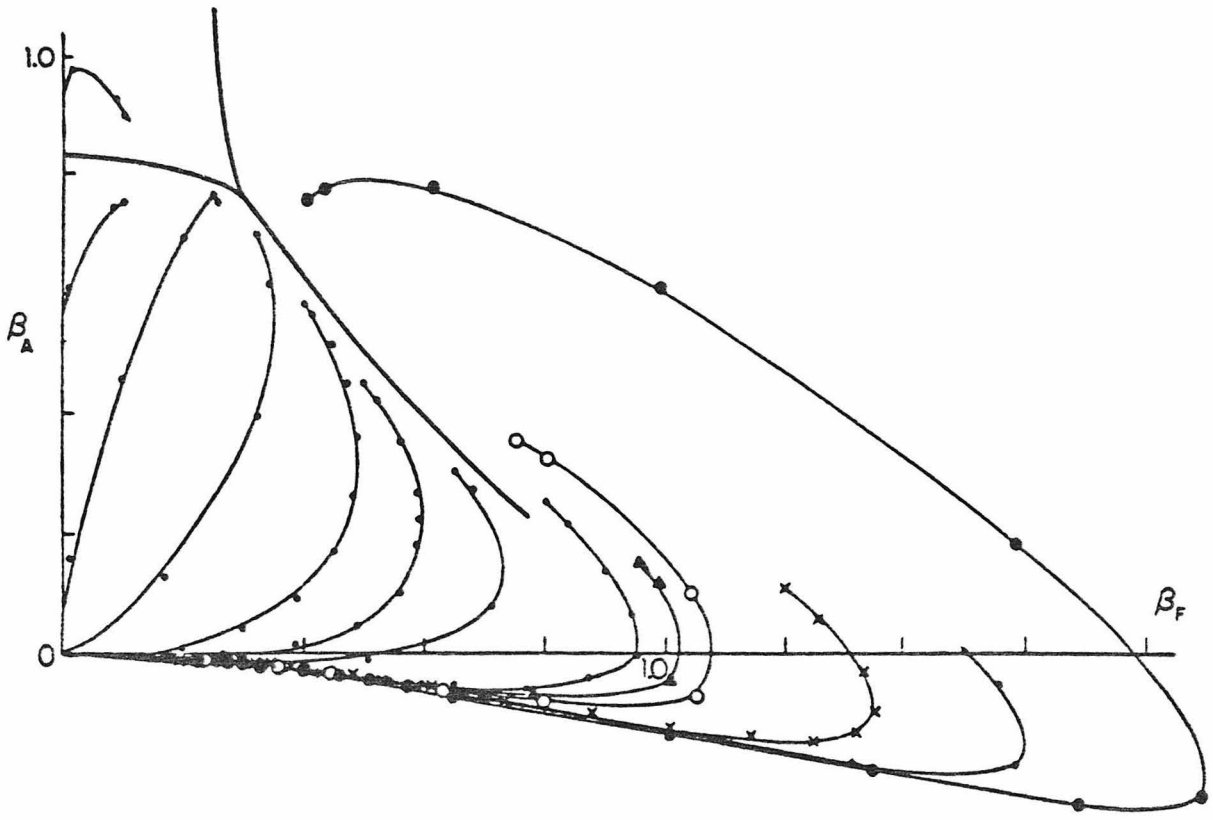

Fig. 4.3

Projection of the renormalised SU(2) action onto the $[\beta_F, \beta_A, \beta_{Cr}]$ space for several starting actions. For each starting action, the first row of figures shows the couplings after 1 blocking transformation on starting lattices of size 9^4 (calculated in [16]). The second row shows the couplings after 2 blocking transformations on starting lattices of size 18^4, calculated using the microcanonical demon method.

Initial action β	β_F	β_A/β_F	β_{Cr}/β_F
2.50 (W)	2.491 ± 0.005	-0.145 ± 0.001	-0.0041 ± 0.0006
	2.044 ± 0.014	-0.158 ± 0.005	-0.006 ± 0.003
2.75 (W)	3.007 ± 0.006	-0.140 ± 0.002	-0.0216 ± 0.0015
	2.662 ± 0.018	-0.161 ± 0.004	-0.011 ± 0.003
3.00 (W)	3.576 ± 0.009	-0.124 ± 0.001	-0.0335 ± 0.0006
	3.198 ± 0.017	-0.139 ± 0.004	-0.034 ± 0.003
3.25 (W)	3.936 ± 0.007	-0.114 ± 0.001	-0.0392 ± 0.0002
	3.802 ± 0.038	-0.132 ± 0.007	-0.041 ± 0.001
3.50 (W)	4.430 ± 0.010	-0.112 ± 0.001	-0.0427 ± 0.0002
	4.351 ± 0.063	-0.129 ± 0.007	-0.040 ± 0.002
4.35 (MK)	3.225 ± 0.007	-0.144 ± 0.001	-0.0283 ± 0.0007
	2.913 ± 0.022	-0.157 ± 0.003	-0.027 ± 0.002

Table 4.1

Projection of the renormalised SU(2) action onto the $[\beta_F,\beta_A,\beta_{3/2},\beta_{6p}]$ space for several starting actions. For each starting action, the first row of figures shows the couplings after 1 blocking transformation on starting lattices of size 9^4 (calculated in [16]). The second row shows the couplings after 2 blocking transformations on starting lattices of size 18^4, calculated using the microcanonical demon method.

Initial action β	β_F	β_A/β_F	$\beta_{3/2}/\beta_F$	β_{6p}/β_F
2.50 (W)	2.571 ± 0.005	-0.195 ± 0.001	0.043 ± 0.001	-0.0036 ± 0.0003
	2.058 ± 0.006	-0.186 ± 0.006	0.038 ± 0.003	-0.010 ± 0.002
2.75 (W)	3.16 ± 0.01	-0.199 ± 0.003	0.042 ± 0.002	-0.0208 ± 0.0015
	2.815 ± 0.035	-0.214 ± 0.011	0.044 ± 0.006	-0.019 ± 0.004
3.00 (W)	3.69 ± 0.01	-0.190 ± 0.004	0.040 ± 0.002	-0.0314 ± 0.0007
	3.469 ± 0.048	-0.211 ± 0.012	0.039 ± 0.004	-0.032 ± 0.003
3.25 (W)	4.12 ± 0.02	-0.160 ± 0.005	0.025 ± 0.003	-0.0374 ± 0.0004
	4.003 ± 0.037	-0.182 ± 0.010	0.032 ± 0.006	-0.040 ± 0.003
3.50 (W)	4.71 ± 0.02	-0.168 ± 0.005	0.028 ± 0.003	-0.0402 ± 0.0004
	4.396 ± 0.067	-0.150 ± 0.015	0.007 ± 0.006	-0.049 ± 0.002
4.35 (MK)	3.42 ± 0.01	-0.211 ± 0.002	0.044 ± 0.001	-0.0268 ± 0.0011
	3.098 ± 0.033	-0.235 ± 0.012	0.055 ± 0.004	-0.029 ± 0.003

Table 4.2

Concluding Remarks

The lattice regularisation has been markedly successful at allowing us to calculate many of the important quantities in gauge field theories. In fact, the lattice is the *only* regularisation scheme which does not require a perturbative expansion in the coupling in order to perform renormalisation. This crucial feature makes up for the fact that the lattice sacrifices some elegant features such as Lorentz invariance, since many problems of interest in gauge theories require us to look in regions where the coupling is not small. The calculation of the hadron spectrum, and the question of how QCD confines quarks within hadrons, are examples of such problems. Analytic methods in the lattice formalism are not yet, and may never be, capable of evaluating the appropriate path integrals for observables in the regions of intermediate couplings. Numerical methods have therefore been introduced to estimate these integrals, and significant progress has been made, some of which has been described in this thesis.

The important questions that now arise are, firstly, "How necessary are numerical methods, in the long run, in comparison to purely analytic approaches?" Secondly, "What directions will these calculations take?"

In regard to the first question, the most important observation to make here is probably the one that non-perturbative issues hinge on the analogy between quantum field theories and statistical physics. Physical quantities are all determined by the correlation length in regions where this is extremely large compared to the lattice spacing, where the collective behaviour of a vast number of locally coupled degrees of freedom becomes dominant. In the cases of physically relevant gauge theories, this collective behaviour is highly non-trivial. Analytic methods have only been able to completely describe this behaviour in the cases of relatively trivial theories such as the Ising model, despite many years of effort. This fact seems to reserve the computation for large-scale computers. In particular, we have in the MCRG technique an approach which can, in principle, tell us everything about the renormalisation of pure gauge theories, provided enough computing power is at our disposal. At the moment this is definitely not the case. For example, restrictions on the number of lattice sites (of the order

of 10 in each dimension), and on the values of the couplings that can be used, typically produce results on lattices barely a hadron radius in size, with a lattice spacing $1/10^{th}$ of this. Future work will clearly depend on the use of machines that are orders of magnitude more powerful than those we have available now.

Now we turn to the second question. Perhaps the most ambitious, and important, possibility opened up by numerical methods is that of implementing a complete renormalisation program for QCD via an MCRG approach that incorporates dynamical fermions. In this way, we would be able to trace the behaviour of the theory from the reasonably well understood region of small distances, to the collective behaviour that occurs over large distances and is responsible for confinement. It's possible that analytic methods could be developed which would "get around" the fact that the fermion determinant is non-local. Without a breakthrough in this area, brute force techniques requiring huge computers are needed.

Ultimately, of course, we need to be able to calculate physical quantities with precision on lattices that we know are close to the continuum limit. This means using very weak bare couplings, where Monte Carlo methods are severely hampered by the problem of critical slowing down already mentioned in the body of the thesis: i.e., the extremely slow traversal of configuration space near the critical point. In order to conduct reliable calculations on large lattices, it will therefore be necessary to avoid this problem by using improved actions of the type discussed in Chapter 4 of this thesis. The use of these improved actions will be crucial in the calculation of the hadron and glueball spectra. When they are implemented, it will be possible to obtain extremely accurate physical predictions on the more powerful machines that will be available to us in the near future.

Appendix 1: The PPR trick and detailed balance

We wish to show here that the PPR improvement described in Section 2 of Chapter 2 only applies to straight sections of a Wilson loop, and not to corners, because detailed balance is violated in the latter case.

The conditon of detailed balance is that the transition probability between 2 configurations C_i and C_j must satisfy

$$\frac{W(C_i,C_j)}{W(C_j,C_i)} = \frac{e^{-S(C_j)}}{e^{-S(C_i)}}. \tag{A1.1}$$

For simplicity assume a 2-dimensional lattice. Consider first an observable which consists of the trace of the product of 2 links U and V lying in a straight line. The sums of link products which constitute the heat bath environments for each link are denoted by E and F respectively. Starting with a legal configuration with link values U_0 and V_0, the PPR trick consists of applying a heat bath independently to each of the 2 links. Thus we compute

$$\bar{U}=\frac{1}{N}\sum_{i=1}^{N} U_i \tag{A1.2}$$

and

$$\bar{V}=\frac{1}{N}\sum_{i=1}^{N} V_i \tag{A1.3}$$

where the probability of obtaining any given link value from the initial value is given by the distribution

$$P(U_0,U_i)=e^{-S(U_iE)} \tag{A1.4}$$

in the former case, and by

$$P(V_0,V_i)=e^{-S(V_iE)} \tag{A1.5}$$

in the latter case. The final observable $Tr\bar{U}\bar{V}$ thus gives N^2 terms U_iV_j from $2N$ heat baths. Since each of these terms is generated independently by the procedure, the total transition probability is just the product of the 2 distributions (A1.4) and (A1.5). We therefore have

$$W(U_0V_0,U_iV_j)=e^{-S(U_iE)}e^{-S(V_jF)} \tag{A1.6}$$

and

$$W(U_iV_j,U_0V_0)=e^{-S(U_0E)}e^{-S(V_0F)}, \tag{A1.7}$$

so detailed balance is trivially satisfied.

Now consider the case of a corner. The environment E belonging to the U link is now a function of the value of the V link, and vice versa. Once again, the PPR trick generates each of the links in any given term U_iV_j independently, so the total transition probability is the product of the individual probabilities as before:

$$W(U_0V_0,U_iV_j)=e^{-S(U_i\mathbf{E}(V_0))}\,e^{-S(V_j\mathbf{F}(U_0))} \tag{A1.8}$$

Hence, if we were to start from the configuration U_iV_j, the probability to go "backwards" to the configuration U_0V_0 is given by

$$W(U_iV_j,U_0V_0)=e^{-S(U_0\mathbf{E}(V_0))}\,e^{-S(V_0\mathbf{F}(U_0))} \tag{A1.9}$$

However, detailed balance requires that this last transition probability read

$$W(U_iV_j,U_0V_0)=e^{-S(U_0\mathbf{E}(V_j))}\,e^{-S(V_0\mathbf{F}(U_i))} \tag{A1.10}$$

The expressions (A1.9) and (A1.10) differ in the bold-face parts. We conclude that the trick is not applicable to parts of an observable where one link in the observable forms part of the environment of another link in the observable. In our calculation, this means that the trick can be applied to the straight parts of Wilson loops, but not to the corners.

Appendix 2: Maximisation of the glueball wave function

In Section 2 of Chapter 3, we described how the glueball mass can be obtained by maximising an expression of the form (see Eq. (3.2.4))

$$y(c)=\frac{c^T A c}{c^T B c} \tag{A2.1}$$

In this expression, c is a vector whose components c_i are the coefficients of the various operators O_i from which the glueball wave function is constructed. The (symmetric) matrices A and B describe the correlations between these operators across different time intervals. Thus, the element A_{ij} gives the correlation between 2 operators O_i and O_j across a given time interval. In this appendix we outline the best way to maximise y in Eq. (A2.1) as a function of the "coefficients vector" c.

The condition we want to satisfy is

$$\frac{\partial y}{\partial c_i} = 0 \quad \text{for} \quad all \quad i \tag{A2.2}$$

A typical term in this derivative is

$$\frac{\partial c^T A c}{\partial c_i} = \sum_{i \neq j} A_{ij} c_j + 2A_{ii} c_i + \sum_{j \neq i} c_j A_{ji} \tag{A2.3}$$

But we know that $A_{ij}=A_{ji}$ since the matrices are symmetric. Hence

$$\frac{\partial c^T A c}{\partial c_i} = 2(Ac)_i \tag{A2.4}$$

We can therefore rewrite Eq. (A2.2) as

$$\frac{(c^T B c)(Ac)_i - (c^T A c)(Bc)_i}{(c^T B c)^2} = 0 \quad \text{for} \quad all \quad i \tag{A2.5}$$

We now assume that $c^T B c$ has no zeroes, which is true provided that our statistics are good

enough. Then Eq. (A2.5) says that

$$Ac = \lambda Bc \tag{A2.6}$$

where

$$\lambda = \frac{c^T Ac}{c^T Bc} \tag{A2.7}$$

We have now reduced our maximisation procedure to solving the generalised eigenvalue problem. We must find the eigenvectors c and the eigenvalues λ of Eq. (A2.6). The set of coefficents that we want is then given by the eigenvector that correponds to the largest eigenvalue. There are standard routines available to solve this problem.