

Optimal Multigrid Algorithms for Calculating Thermodynamic Limits

A. Brandt, M. Galun & D. Ron

*Dept. of Applied Mathematics & Computer Science
The Weizmann Institute of Science
Rehovot 76100, Israel*

Abstract

Beyond eliminating the critical slowing down, multigrid algorithms can also eliminate the need to produce many fine-grid configurations in order to reduce statistical errors. Thermodynamic limits can be calculated to accuracy ε in just $O(\varepsilon^{-2})$ computer operations. Examples described in details and with results of numerical tests are the calculation of the susceptibility and the average energy in Gaussian models, and also the determination of the critical temperature in a two dimensional Ising spin model. Extension to more advanced models are outlined.

KEY WORDS: Multigrid; Gaussian model; Ising spin model; Monte Carlo; thermodynamic limit; approximate coarsening.

1. Introduction

The aim in statistical physics is to calculate various average properties of configurations governed by the Boltzman distribution. This is usually done by measuring these averages over a sequence of Monte Carlo iterations. Unfortunately, such processes tend to suffer from several independent inefficiency factors that multiply each other and thus produce very expensive computations.

The best known of these inefficiency factors is the *critical slowing down* (CSD). This is the phenomenon, typical to simulations of critical systems, that with the increase in lattice size there also comes an increase in the number of Monte Carlo passes over the lattice needed to produce a new configuration which is statistically “useful”, i.e., substantially independent of, or only weakly correlated to, a former configuration. Considerable efforts have been devoted to overcome this difficulty. For simple enough cases with real-state variables and at most mild nonlinearities, a general method to eliminate CSD is by classical multigrid methods, properly adapted (see §7.1 in ref. and ref.). For models with severe nonlinearities or discrete variables, such as the ϕ^4 or Ising spin models, a number of publications report on simulation techniques that partially (,) or completely (,,,) eliminate CSD. This means that in a work just proportional to the number of gridpoints, a new, substantially independent configuration can be generated.

Optimal as this result is, other, no less important factors of inefficiency still remain intact. To calculate a thermodynamic quantity to a certain accuracy ε , one needs to produce $O(\sigma^2\varepsilon^{-2})$ essentially independent configurations to average out the deviation exhibited by each of them, where σ denotes the standard (i.e., the L_2 average) deviation. Also, the size of the grid must increase as some positive power of ε^{-1} . Multiscale calculations may overcome these additional inefficiency factors as well by introducing more statistical measurements at coarse levels and by other means (,App.B). More directly, what we intend to demonstrate below is that the multigrid structure can be used for measuring meaningful thermodynamic quantities in an *optimal* computational time.

Namely, we will show that thermodynamic limits (quantities obtained at the limit of infinite grids) can be calculated to accuracy ε using only $O(\varepsilon^{-2})$ *computer operations*. This is just the same order of complexity as needed to calculate, by statistical trials, any simple “pointwise” average, such as the frequency of “heads” in coin tossing.

Our prime examples here will be the calculation of the susceptibility and the average energy in the Gaussian model, and of the critical temperature in the two dimensional Ising model. These cases are ideal for developing, testing and demonstrating the new multilevel techniques, because of their simplicity and because analytical solutions are known and can be used for comparing the results and understanding the behavior of the numerical processes.

For the Gaussian model it is shown in Sec. 2 that both the susceptibility (at

the limit of vanishing meshsize or infinite grid) and the average energy per degree of freedom (on arbitrarily large grids) can be calculated to a relative accuracy ε in about $10\varepsilon^{-2}$ random number generations (independently of the size of the grid). It is also shown that the multigrid algorithms for these two calculations (susceptibility and energy) cannot be the same; their “cycle index” must differ.

The generalization from the Gaussian to other models, with continuous state but not quadratic Hamiltonians, is not straightforward, but possible. The different scales are no longer strictly independent of each other, but it is still usually true that large-scale Hamiltonians can be constructed with significantly-computation-consuming couplings only to neighboring scales, not to much finer ones. The general approach is outlined in Sec. 4. An important feature is that it can be used for a direct and simple computational derivation of macroscopic equations for the model at hand.

For *compact-state* models this general approach turns out to lead quite naturally to stochastic (approximate) blocking occurring on sufficiently coarser grids, reminiscent of the Swendsen-Wang (SW) blocking in discrete-state models. The *explicit* SW-type blocking is thus needed only for *discontinuous* models, i.e., few-state models (Ising, Potts, etc.) or, more generally, models whose large-scale probable fluctuations exhibit strong jumps already at the smallest scale (ϕ^4 models, etc.; the two dimensional XY model at critical temperature may also belong here).

It is not clear whether for such discontinuous models *optimal* computations of thermodynamic limits (i.e., obtaining accuracy ε in $O(\varepsilon^{-2})$ operations) is always possible. As an example, we discuss in detail calculations with the two dimensional Ising spin model. It is shown that, due indeed to the discontinuity of the model, the sequence of configurations produced within one multigrid cycle so depend on each other that not much can be gained by introducing statistical measurements at coarse levels. Nevertheless, the results indicate that it may still be possible to calculate thermodynamic quantities, such as the critical temperature, in optimal ($O(\varepsilon^{-2})$) time.

2. One dimensional Gaussian model: fast calculations of thermodynamic quantities

A multigrid algorithm for simple continuous-state models, such as the Gaussian model, has been described by us in ref. [1], and independently, by Goodman and Sokal [2]. The latter also tested it for the Gaussian model and reported that it indeed eliminated CSD. In the present work we show that the multigrid Monte Carlo can be used not only for eliminating the CSD, but also for accelerated calculation of averages, especially those which depend on large-scale fluctuations. Thermodynamic limits can be calculated to accuracy ε in just $O(\varepsilon^{-2})$ computer operations.

The outline of this chapter is as follows. In Sec. 2.1 we have first considered the one dimensional continuous case for which, by Fourier analysis, we could calculate in closed form some thermodynamic quantities. Analogous quantities on finite grids are introduced in Sec. 2.2. In Secs. 2.3 and 2.4 we describe an extremely efficient multigrid algorithm for evaluating the discrete susceptibility and approaching its limit for zero meshsize. Some remarks on parallel processing are briefly given in Sec. 2.5. Numerical experiments and results are finally summarized in Sec. 2.6. The evaluation of the average energy per degree of freedom is described in Sec. 2.7.

2.1 Continuous case

To facilitate theoretical analysis of the algorithms and results we treat the constant coefficient case. But the same algorithms can be used for much more general situations, with a similar efficiency.

The Hamiltonian associated with the continuous case is:

$$\mathcal{H}(u) = \int_0^L u_x^2 dx, \quad (1)$$

where $u = u(x)$ is a function (configuration) defined for $0 \leq x \leq L$. Homogeneous Dirichlet boundary conditions, $u(0) = u(L) = 0$, are used for definiteness, though others could serve as well. Consequently, a general configuration $u(x)$ can be expanded by:

$$u(x) = \sum_{j=1}^{\infty} c_j \sin(j\pi x/L), \quad (2)$$

where the Fourier coefficient c_j are real. By substituting (2) into (1) one gets:

$$\mathcal{H}(u) = \frac{\pi^2}{2L} \sum_{j=1}^{\infty} j^2 c_j^2. \quad (3)$$

The magnetization is given by:

$$M(u) = \frac{1}{L} \int_0^L u(x) dx = \frac{2}{\pi} \sum_{j=1}^{\infty*} c_j / j, \quad (4)$$

where \sum^* , here and below, stands for a summation over *odd* integers. As the probability density of each configuration u is given by the density function of the Boltzman distribution

$$P(u) = e^{-\mathcal{H}(u)/T} / Z(T), \quad (5)$$

straightforward calculations of the averaged magnetization ($\langle M \rangle$), susceptibility ($\langle M^2 \rangle - \langle M \rangle^2$) and energy ($\langle \mathcal{H} \rangle$) can easily be made using the above Fourier expansion, leading to the following results:

$$\langle M \rangle = 0 \quad (6a)$$

$$\langle M^2 \rangle - \langle M \rangle^2 = \langle M^2 \rangle = \frac{4LT}{\pi^4} \sum_{j=1}^{\infty} (1/j^4) \quad (6b)$$

$$\langle \mathcal{H} \rangle = \frac{T}{2} \sum_{j=1}^{\infty} 1 = \infty. \quad (6c)$$

Although the Hamiltonian is not bounded, its differences associated with changing any c_j (or any discrete degree of freedom defined in Sec. 2.2 below) are well defined, hence it yields these well defined statistics.

2.2 Discrete case

In order to measure such statistical averages numerically, it is necessary to discretize the system. On a grid with meshsize $h = L/N$, the discretized Hamiltonian $\mathcal{H}_h(u)$, approximating (1), can be written as:

$$\mathcal{H}_h(u) = \frac{1}{h} \sum_{i=1}^N [u(x_i) - u(x_{i-1})]^2, \quad (7)$$

where $x_i = ih$ ($0 \leq i \leq N$) are the gridpoints. For the simplicity of the multigrid algorithm (see Sec. 2.3) we have assumed $N = 2^k$; the general case could, however, be calculated as well by handling the boundaries differently. Assuming again $u(x_0) = u(x_N) = 0$, a general grid configuration can be represented by

$$u(x_i) = \sum_{j=1}^{N-1} c_j \sin(j\pi x_i/L), \quad (8)$$

leading together with (7) to

$$\mathcal{H}_h(u) = \frac{2L}{h^2} \sum_{j=1}^{N-1} c_j^2 \sin^2(j\pi h/(2L)). \quad (9)$$

The discrete magnetization is given by

$$M_h(u) = \frac{h}{L} \sum_{i=0}^N u(x_i) = \frac{h}{L} \sum_{j=1}^{N-1} c_j \frac{\cos(j\pi h/(2L))}{\sin(j\pi h/(2L))}. \quad (10)$$

By using the probability distribution (5), where $\mathcal{H}_h(u)$ replaces $\mathcal{H}(u)$, one can derive:

$$\langle M_h \rangle = 0 \quad (11a)$$

$$\langle M_h^2 \rangle = \frac{Th^4}{4L^3} \sum_{j=1}^{N-1} \frac{\cos^2(j\pi h/(2L))}{\sin^4(j\pi h/(2L))} \quad (11b)$$

$$\langle \mathcal{H}_h \rangle = T(N-1)/2. \quad (11c)$$

As $N \rightarrow \infty$ (with fixed L , hence $h \rightarrow 0$), which is called here the thermodynamic limit, the results of the discrete case tend to those of the continuum. $\langle M_h^2 \rangle$ exhibits a discretization error of $O(h^2)$. More precisely, Taylor expansion of each term in (11b) and comparison to (6b) yields

$$\langle M_h^2 \rangle = \langle M^2 \rangle - \frac{Th^2}{3\pi^2 L} \sum_{j=1}^{\frac{N}{2}} \frac{1}{j^2} + TLO(N^{-3}), \quad (11d)$$

where terms with $j > N/2$ are omitted from both $\langle M_h^2 \rangle$ and $\langle M^2 \rangle$ since they clearly have only $TLO(N^{-3})$ total contribution. The *relative* discretization error, e.g., $|\langle M_h^2 \rangle - \langle M^2 \rangle|/\langle M^2 \rangle$, is clearly $O(h^2/L^2)$. Clearly, by using a p^{th} order discretization, this relative error can be further reduced to $O(h^p/L^p)$. As will be explained later, the algorithm will not necessarily actually provide this $O(h^p/L^p)$ accuracy for any given gridsize, as this may turn out to be wasteful when statistical errors are taken into account as well. Instead, it will be constructed so as to keep the number of operations optimal with respect to the overall produced accuracy. The details of the algorithm are given next.

2.3 Description of the multigrid cycle

Consider the following generalized Hamiltonian, which includes an external magnetic field of density ϕ_i at gridpoint x_i (actually, on the finest grid $\phi_i = 0$ for $0 \leq i \leq N$):

$$\mathcal{H}_h(u) = \frac{1}{h} \sum_{i=1}^N (u_i - u_{i-1})^2 + h \sum_{i=1}^{N-1} \phi_i u_i, \quad (12)$$

where $u_i = u(x_i)$.

The coarse grid with meshsize $H = 2h$ is constructed by taking every other fine-grid point: see Fig. 1. The coarse-grid function $u^H = (u_0^H, \dots, u_I^H, \dots, u_{N/2}^H)$ is interpreted as a *correction* to the fine-grid function $u^h = (u_0, \dots, u_i, \dots, u_N)$, updating the latter through *interpolation and addition*:

$$u^h \leftarrow \tilde{u}^h + I_H^h u^H, \quad (13)$$

where \tilde{u}^h is the fine-grid configuration at the stage of switching to the coarse-grid and I_H^h denotes interpolation from grid H to grid h (we will use the *linear* interpolation, which is optimal here, see conclusion C in Sec. 2.4).

The fine-grid Hamiltonian $\mathcal{H}_h(u^h)$ resulting from that interpolation can be written as follows:

$$\mathcal{H}_h(\tilde{u}^h + I_H^h u^H) = \mathcal{H}_h(\tilde{u}^h) + \mathcal{H}_H(u^H), \quad (14)$$

where $\mathcal{H}_h(\tilde{u}^h)$ is given by (12) and $\mathcal{H}_H(u^H)$ is:

$$\mathcal{H}_H(u^H) = \frac{1}{H} \sum_{I=1}^{N/2} (u_I^H - u_{I-1}^H)^2 + H \sum_{I=1}^{N/2-1} \phi_I^H u_I^H, \quad (15)$$

with

$$\phi_I^H = \frac{-\tilde{u}_{i-2}^h + 2\tilde{u}_i^h - \tilde{u}_{i+2}^h}{2h^2} + \frac{\phi_{i-1}^h + 2\phi_i^h + \phi_{i+1}^h}{4}, \quad (I = i/2 = 1, \dots, N/2 - 1) \quad (16)$$

representing fine-to-coarse induced field-like terms. These coarse terms are calculated from the details of the fine-grid configuration at coarsening and are fixed throughout the processing on the coarser level. The variables of the coarse grid u_I^H are initially set to zero, corresponding to zero initial corrections.

The entire algorithm can be described by a sequence of multigrid cycles for the finest level. A cycle for any given (“current”) level is recursively defined by the following five steps.

1. ν_1 Monte Carlo sweeps are first made on the current level. Then, if this level is the coarsest, go to 5.
2. The next coarser level is created from the current one by determining the coarse field-like terms (16).
3. γ multigrid cycles for the coarse level are performed. (γ may change from one cycle of the current level to another in some periodic manner. The *cycle index* is the average value of γ , and for convenience it will also be denoted by γ . Thus γ need not be an integer, and may be smaller than 1.)
4. Update the current level by performing the linear interpolation (13) from the coarse level.
5. Additional ν_2 Monte Carlo sweeps are finally made on the current level.

The Monte Carlo sweeps are performed by changing each variable in its turn randomly according to its associated distribution, regarding its neighbors as fixed.

The values of ν_1 , ν_2 and γ are discussed below.

The described cycle, even with $\gamma = 1$ (which is called a *V* cycle), would generate a new configuration substantially independent of the pre-cycle one in a

work just proportional to the number of gridpoints; it would thus eliminate the critical slowing down. By the term *substantially independent configuration* we mean that the correlation between any quantity of interest in the initial arbitrary configuration and in the one produced after k cycles decays like $e^{-k/\tau}$, where τ , the cycle auto-correlation time, is independent of the lattice size N . In fact τ is very small, so there is almost no correlation between any quantities before and after even one cycle.

The crucial issue, however, turns out to be different; it is addressed next.

2.4 Fast sampling of susceptibility

Any observable which is bounded as $h \rightarrow 0$ (e.g., magnetization and susceptibility, but *not* energy; see (6c)), must be dominated by contributions from large-scale fluctuations (low-frequency Fourier components; cf., e.g., (6b)). The main issue of the Monte Carlo simulation is therefore to sample quickly as many such fluctuations as possible. For this purpose a cycle index γ substantially larger than 1 may be used, and averages will be calculated over as many measurements as one could make within each cycle, especially at its coarsest level stages.

Consider, for instance, the calculation of $\langle M_h^2 \rangle$. Observe first that M_h can be evaluated on any level. Indeed, denoting by \tilde{u}^h the fine-grid configuration at coarsening, (10) and (13) imply that

$$M_h = \frac{h}{L} \sum_{i=0}^N (\tilde{u}^h + I_{2h}^h u^{2h})_i = \frac{h}{L} \sum_{i=0}^N \tilde{u}_i^h + \frac{2h}{L} \sum_{I=0}^{N/2} u_I^{2h}. \quad (17)$$

Employing this recursively one gets that when working on level ℓ , if its current configuration is u^{h_ℓ} , then the (fine-grid) magnetization corresponding to it is

$$M_h = \sum_{j=0}^{\ell-1} \tilde{M}_{h_j} + M_{h_\ell}, \quad (18)$$

where generally $h_j = 2^j h$, $M_{h_j} = \frac{h_j}{L} \sum_I u_I^{h_j}$, $\tilde{M}_{h_j} = \frac{h_j}{L} \sum_I \tilde{u}_I^{h_j}$, and \tilde{u}^{h_j} is the j^{th} level configuration at the stage of switching to the next coarser level. Thus, many measurements of M_h^2 can be made within a cycle, and their average $\overline{M_h^2}$ can be used as an estimate for the discrete susceptibility $\langle M_h^2 \rangle$. In practice, measurements need be taken only on the coarsest level, and in fact after each relaxation sweep there, because only that is when substantial changes in M_h are introduced.

Let us now estimate the number m_i of relaxation sweeps the algorithm needs to perform on level i , i.e. on the grid with meshsize $h_i = 2^i h$, ($i = 0, 1, \dots, \ell$)

$\log_2(N/2)$), in order to achieve relative accuracy ε in the estimation of the susceptibility. The relaxation sweep on level i strongly affects, hence effectively samples, only those Fourier coefficients c_j (cf. (2)) for which $j = O(L/h_i)$. Hence m_i depends on the contribution of these components to the deviations in measuring $\langle M^2 \rangle$. By (4)

$$M^2 = \frac{4}{\pi^2} \sum_{j,k}^* c_j c_k / (jk). \quad (19)$$

Consider first a term (j, k) in (19) for which both j and k are $O(L/h_i)$, hence the term is effectively sampled $O(m_i)$ times in a cycle. Since the standard deviation of the term is

$$\frac{4}{jk\pi^2} \left(\langle (c_j c_k)^2 \rangle - \langle c_j c_k \rangle^2 \right)^{1/2} = O(j^{-2} k^{-2} LT) = O(h_i^4 L^{-3} T),$$

the standard deviation of its average over the $O(m_i)$ samples is $O(m_i^{-1/2} h_i^4 L^{-3} T)$. There are $O(h_i^{-2} L^2)$ such terms, and their deviations can be considered independent of each other, hence their total contribution is $O(m_i^{-1/2} h_i^3 L^{-2} T)$. In case $j = O(L/h_{i-r})$ and $k = O(L/h_i)$ where $r \geq 1$ (i.e., $h_i > h_{i-r}$), the term (j, k) in (19) is effectively sampled as follows: in an inner loop, for a (nearly) fixed value of c_j , the value of c_k is averaged $O(m_i/m_{i-r})$ times, yielding an average whose deviation is of the order

$$O \left(\frac{c_j}{jk} \left(\frac{m_i}{m_{i-r}} \right)^{-1/2} \langle c_k^2 \rangle^{1/2} \right) = O \left(\frac{c_j}{jk^2} \left(\frac{m_{i-r}}{m_i} LT \right)^{1/2} \right).$$

Then, in an outer loop, the c_j in this average is averaged over $O(m_{i-r})$ samples, giving results with deviations of order

$$O \left(\frac{m_{i-r}^{-1/2}}{jk^2} \left(\frac{m_{i-r}}{m_i} LT \right)^{1/2} \langle c_j^2 \rangle^{1/2} \right) = O \left(\frac{m_i^{-1/2}}{j^2 k^2} LT \right) = O(m_i^{-1/2} h_i^2 h_{i-r}^2 L^{-3} T).$$

There are $O(h_i^{-1} h_{i-r}^{-1} L^2)$ such terms, effectively independent, hence their total deviation is $O(m_i^{-1/2} h_i^{3/2} h_{i-r}^{3/2} L^{-2} T)$, which, when summed over integers $r \geq 0$ gives again $O(m_i^{-1/2} h_i^3 L^{-2} T) = \langle M^2 \rangle O(m_i^{-1/2} h_i^3 L^{-3})$. Hence the *total relative* expected error in measuring $\langle M^2 \rangle$ is

$$\varepsilon = O \left(\sum_{i=0}^l m_i^{-1/2} h_i^3 L^{-3} \right) + O(h^p L^{-p}) \quad (20)$$

where the last term added here is the discretization error (cf. Sec. 2.2). The total work (operations) on all the levels is clearly

$$W = \sum_{i=0}^l m_i O(L/h_i). \quad (21)$$

The optimal choice for m_i (yielding either minimal ε for a given W or minimal W for a given ε) is obtained when $\frac{\partial \varepsilon}{\partial m_i} + \lambda_1 \frac{\partial W}{\partial m_i} = 0$, which, by (20) and (21) yields

$$m_i = \lambda_2 h_i^{8/3} = \lambda_3 2^{8i/3}, \quad (22)$$

where λ_1, λ_2 and λ_3 are independent of i . Relation (22) is realized by the cycle index $\gamma_{\text{opt}} = 2^{8/3} \simeq 6.35$.

For *any* fixed cycle index γ we have $m_i = m\gamma^i$, where m is the total number of cycles performed. Since $h_i L^{-1} = O(2^{i-l})$, we can perform the summations in (20) and (21) and obtain

$$\varepsilon = O\left(m^{-1/2} \frac{\gamma^{-l/2} - 2^{-3l}}{1 - 2^{-3}\gamma^{1/2}}\right) + O(2^{-lp}) \quad (23)$$

and

$$W = O\left(m \frac{\gamma^l - 2^l}{1 - 2\gamma^{-1}}\right) \quad (24)$$

for any $2 < \gamma < 2^6$. Actually, by choosing γ and the approximation order p so that γ is significantly smaller than 2^{2p} , the second term in (23) can be ignored, yielding $W = O(\varepsilon^{-2})$ and $m = O(\varepsilon^{-2}\gamma^{-l})$. While $\gamma = \gamma_{\text{opt}}$ indeed minimizes $W\varepsilon^2$, the value of $W\varepsilon^2$ (when the second term in (23) is negligible) is only 15% different from that minimum for any $4 \leq \gamma \leq 11$. The four main conclusions from this analysis are therefore as follows.

A. Cycle index γ . Generally, any $2 < \gamma < \min(64, 2^{2p})$ yields $W = O(\varepsilon^{-2})$. *Asymptotically* (for $\varepsilon \rightarrow 0$), the minimal value of $W\varepsilon^2$ is attained for $\gamma = 2^{8/3} \simeq 6.35$ and values very close to the minimum are obtained for $4 \leq \gamma \leq 11$. *In practice* (for realistic values of ε), the smaller values of γ in this range are better, since for them the influence of the second term in (23) is smaller.

B. Discretization order p . There is little advantage in raising the order beyond $p = 2$. It would allow the use of cycles with larger γ , but the dominating coarsest-grid work will remain essentially the same. The only slight advantage may be the smaller storage requirement, which is $O(N) \geq O(\varepsilon^{-1/p})$ (cf. Sec. 2.5).

C. Interpolation order. Linear interpolation (second order interpolation) I_{2h}^h is good enough, as any smooth configuration u^h has an approximate configuration $I_{2h}^h u^{2h}$, such that the two configurations has almost the same energy. This means that the physical density function does not change much and thus the efficiency of the algorithm is not reduced. Thus, not much could be gained by using interpolation orders higher than 2 (meaning I_H^h higher than linear). Such higher interpolation order would also make the coarse grid Hamiltonian substantially more complicated. A relatively little complication, together with almost all

the possible gain, can be obtained by using third order interpolation (based either on the function $x(L - x)$ or on $\sin(\pi x/L)$) only at the transition to the *coarsest* grid (which has only one degree of freedom).

On the other hand, the order of the interpolation operator I_{2h}^h should be *at least* 2, i.e., linear interpolation. This is because the coarser levels should accurately sample all the components slow to change under the current-level Monte Carlo process. It means that every slowly changing configuration u^h must have an approximate configuration of the form $I_{2h}^h u^{2h}$, and the two configurations should have approximately the same energy. The linear interpolation satisfies this requirement. A border case is the *first order (constant) interpolation* I_{2h}^h . For any smooth configuration u^h approximated by a function u^{2h} on a coarse grid with meshsize $2h$, the energy of $I_{2h}^h u^{2h}$ is about twice that of u^h . Hence, a very smooth component on the finest grid, u^h , that is effectively changed only on a coarse grid with meshsize $2^q h$, will be represented on that coarse grid by an approximation which has an energy 2^q times its own energy. Thus, the changes introduced to the amplitude of such an approximation on that coarse grid will be only $O(2^{-q/2})$ times the typical fluctuations of that amplitude. It follows that roughly 2^q visits to that coarse grid will be needed to accumulate a typical fluctuation. This means that a *cycle index* $\gamma \geq 2$ *must be used* to eliminate the CSD. Thus, for two reasons *constant interpolation* is not used. Firstly, the work per W cycle ($\gamma = 2$) is $O(N \log N)$, so one cannot eliminate the CSD. Secondly, and much more importantly, a whole W cycle is needed to produce a single useful measurement, whereas in the algorithm above each additional movement on the coarsest grid generates another useful measurement.

D. Number of cycles m . *Asymptotically*, to avoid an error ε dominated by the discretization error (the second term in (23)), one should choose $m \leq O((16/\gamma)^\ell)$. Any larger number of cycles would do useless work of reducing the statistical error, because this error is already smaller than the discretization error. *In practice*, for realistic ε , the smallest m possible, i.e., $m = 1$, is preferable for minimizing the influence of the discretization error. This means that whenever the desired accuracy ε is reduced (or the available amount of processing W is increased) the work is increased not by increasing m , but by raising l , i.e., introducing new finer levels (processed very rarely, of course) only.

In summary, in computing susceptibility one can use second order discretization and second order interpolation, any cycle index in the range $4 \leq \gamma \leq 11$ (with some preference for the lower values), and any number of cycles $1 \leq m \leq O((16/\gamma)^\ell)$ (with some preference to $m = 1$): the overall computational work will always be dominated by the $m\gamma^\ell = O(\varepsilon^{-2})$ work on the coarsest grid, with ε being the relative accuracy that will be obtained in calculating the thermodynamic limit $\langle M^2 \rangle$.

2.5 Parallel processing

The algorithm described above can use very high degree of parallel processing. On each grid, the Monte Carlo sweeps can proceed simultaneously at half the grid-points: first the odd, then the even. The calculation of the coarse-grid functions (16) can be done at all points in parallel. More important, the employment of the γ coarse grid cycles can all proceed in parallel to each other, and so can the γ^2 cycles of the still coarser level, etc. So, if an unlimited number (or actually $O(\varepsilon^{-2})$) processors are available, the algorithm can in principle be performed in $O(\log \frac{1}{\varepsilon})$ parallel steps. This of course ignores processor communication considerations.

For keeping the total communication, as well as the total storage, at minimum, it may be desired to have as small N (number of points on the finest grid) as possible. For this, a discretization order $p = 4$ may be preferred, since $N = O(\varepsilon^{-1/p})$. Still higher order discretization would not help, since components with $O(h)$ wavelength contribute $O(h^4)$ to $\langle M^2 \rangle$, so for an $O(\varepsilon)$ accuracy a grid must be used whose meshsize is no larger than $O(\varepsilon^{1/4})$.

2.6 Numerical results

We have tested the multigrid algorithm with $\gamma = 1, 2, 3, 4, 6, 12$ and 24 on grid of sizes up to 128 . Our main aim was to show that for appropriate values of γ optimal behavior is achieved; i.e., the average error in the approximation for $\langle M^2 \rangle$ (6b) produced by a multigrid cycle is reduced by a factor $\sqrt{\gamma}$ upon using a (finest) grid twice finer (N twice larger), which increases the work by a factor γ . The susceptibility has been measured over just *one* cycle ($m = 1$). In that cycle, M_h^2 has been measured, using (18), after each relaxation step on the coarsest level, hence $(\nu_1 + \nu_2)\gamma^\ell$ times altogether. The average of these measurements, $\overline{M_h^2}$, is the approximation for $\langle M_h^2 \rangle$ (11b), which in turn is also an approximation for the desired thermodynamic limit $\langle M^2 \rangle$. The measured relative error is defined as $\varepsilon = |\overline{M_h^2} - \langle M^2 \rangle| / \langle M^2 \rangle$. We also define α to be the expected value of $\# \text{RAN} \cdot \varepsilon^2$, where $\# \text{RAN}$ is the amount of work spent per cycle, measured by the number of times a random number is generated, which (for $\gamma > 2$) is dominated by $(\nu_1 + \nu_2)\gamma^\ell$, the number of relaxation steps on the coarsest level. Thus, α should turn out constant if indeed the algorithm solves to accuracy ε in $O(\varepsilon^{-2})$ operations. For each value of N , α was estimated by averaging ε over an ensemble of 2000 to 5000 runs.

In Fig. 2 we present α for $\gamma = 1, 2, 3, 4, 6, 12$ and 24 vs (\log) system size N . As expected, for $2 < \gamma \leq 16$ the value of α tends to a constant as N increases. Also, the cycle with smaller values of γ is slightly more efficient (has smaller α) as expected (see conclusion A in Sec. 2.4). The graphs for $\gamma = 4, 6$ are indistinguishable.

2.7 σ -Susceptibility and its computation

We introduced a new type of observable, χ_σ , defined in the continuum limit by

$$\chi_\sigma = \left\langle \frac{1}{L} \int_0^L (u(x))^2 dx - \left(\frac{1}{L} \int_0^L u(x) dx \right)^2 \right\rangle,$$

and approximated in the discrete model by

$$\begin{aligned} \chi_\sigma^h &= \left\langle \frac{1}{N} \sum_{i=0}^N u_i^2 - \left(\frac{1}{N} \sum_{i=0}^N u_i \right)^2 \right\rangle \\ &= \left\langle \frac{1}{N} \sum_{i=0}^N u_i^2 \right\rangle - \langle M^2 \rangle. \end{aligned}$$

The significance of this observable is that it corresponds to the susceptibility of the Ising model (in the cases where $\langle M \rangle = 0$), and to susceptibility of other sigma models at the limit of small fluctuations. Indeed, assuming $u_i - \bar{u}$ to be small, where $\bar{u} = (\sum_{j=1}^N u_j)/N$, it is easy to see that

$$\frac{1}{N} \left\langle \left(\sum_{i=1}^N \cos u_i \right)^2 + \left(\sum_{i=1}^N \sin u_i \right)^2 \right\rangle \simeq N - \frac{1}{N} \chi_\sigma^h,$$

the expression on the left-hand side being the susceptibility of the XY model with u_i as the angle parameters. We therefore call χ_σ the σ -susceptibility.

The observables χ_σ and χ_σ^h can be computed analytically as follows. By Parseval Identity and by (4)

$$\chi_\sigma = \left\langle \frac{1}{2} \sum_{j=1}^{\infty} c_j^2 - \left(\frac{2}{\pi} \sum_{j=1}^{\infty} c_j/j \right)^2 \right\rangle.$$

Hence (using the density function of the Boltzman distribution (5)) the continuous σ -susceptibility is

$$\chi_\sigma = \frac{LT}{2\pi^2} \sum_{j=1}^{\infty} \frac{1}{j^2} - \frac{4LT}{\pi^4} \sum_{j=1}^{\infty} \frac{1}{j^4} = \frac{LT}{24}.$$

Similar calculations in the discrete case lead to

$$\begin{aligned} \chi_\sigma^h &= \left\langle \frac{1}{2} \sum_{j=1}^{N-1} c_j^2 - \left(\frac{h}{L} \sum_{j=1}^{N-1} c_j \frac{\cos(j\pi h/(2L))}{\sin(j\pi h/(2L))} \right)^2 \right\rangle \\ &= \frac{Th^2}{8L} \sum_{j=1}^{N-1} \frac{1}{\sin^2(j\pi h/(2L))} - \frac{Th^4}{4L^3} \sum_{j=1}^{N-1} \frac{\cos^2(j\pi h/(2L))}{\sin^4(j\pi h/(2L))}. \end{aligned}$$

We claim that the same multigrid algorithm, which is described in Sec. 2.3, can achieve optimal results, in measuring the σ -susceptibility, using appropriate cycle index γ . Fourier analysis shows that any $2 < \gamma < 8$ yields $W = O(\varepsilon^{-2})$. Numerical experiments for $\gamma = 3, 5$ demonstrates that the algorithm solves to accuracy ε in $O(\varepsilon^{-2})$ operations, i.e. α tends to a constant as N increases. See Fig. 2.a.

2.8 Computing average energy per degree of freedom

From (6c) and (11c) it is evident that the average energy per degree of freedom is exactly $T/2$ both in the discretization, with any meshsize, and in the continuum limit (where that is exactly the average energy of each Fourier component). We study now the fast Monte Carlo calculation of this quantity.

Whereas the calculation of susceptibility has been shown to be heavily dominated by the coarsest level, the sampling of \mathcal{H} presents the other extreme. Since most Fourier components are substantially affected by relaxation on the finest level, a measurement of \mathcal{H} should be done after each relaxation sweep on that level, and the work should be dominated by the finest-grid sweeps.

Consider the calculation of the average energy on a given *fixed* finest grid, with meshsize h_0 and $N = L/h_0$ intervals. Since, $\mathcal{H}(u) = \frac{\pi^2}{2L} \sum_{j=1}^{\infty} j^2 c_j^2$ (see (3)), in each measurement of \mathcal{H} each Fourier component contributes the following deviation:

$$\frac{j^2 \pi^2}{2L} \left(\langle (c_j^2)^2 \rangle - \langle c_j^2 \rangle^2 \right)^{1/2},$$

which is of the order $O(T)$. For any level with meshsize h_i , the number of components with wavelength $O(h_i)$ is $O(h_i^{-1}L)$, and their total deviation in each measurement is therefore $O(h_i^{-1/2}L^{1/2}T)$. If grid h_i is averaged $m_i \geq 1$ times, this deviation drops to $O(m_i^{-1/2}h_i^{-1/2}L^{1/2}T)$. To obtain relative accuracy ε , this deviation should be less than $\varepsilon \langle \mathcal{H}_h \rangle = O(\varepsilon h_0^{-1}LT)$, hence it is necessary that $m_i \geq O(\varepsilon^{-2}h_0^2h_i^{-1}L^{-1})$. To guarantee that the deviations contributed from all levels do not accumulate unboundedly, the slightly stronger condition

$$m_i \geq O(\varepsilon^{-2}h_0^{2-\delta}h_i^{\delta-1}L^{-1}), \quad (24)$$

where δ is any (small) positive number, may be required.

In particular, $m_0 \geq O(\varepsilon^{-2}N^{-1})$. On the other hand, for the total work to be *at most* $O(\varepsilon^{-2})$ it is necessary that $Nm_0 \leq O(\varepsilon^{-2})$, hence

$$m_0 = O(\varepsilon^{-2}N^{-1}). \quad (25)$$

Both (24) and (25) can be satisfied by any multigrid cycle with index $\gamma > .5$. The total work will still be $O(\varepsilon^{-2})$ and independent of N iff $\gamma < 2$.

Thus, the energy can be calculated to $O(\varepsilon)$ relative accuracy in $O(\varepsilon^{-2})$ operations by a multigrid cycle with index $.5 < \gamma < 2$. Effectiveness diminishes as γ approaches either endpoints (.5 or 2). The V cycle, i.e., $\gamma = 1$, is of course the most convenient.

Results are presented in Fig. 3. \mathcal{H}_h has been measured, using (14), after each relaxation sweep on *each* level. The average of these measurements, $\overline{\mathcal{H}_h}$, is an approximation for $\langle \mathcal{H}_h \rangle$. The measured relative error is defined as $\varepsilon = |\frac{\overline{\mathcal{H}_h}}{N-1} - \frac{T}{2}| / \frac{T}{2}$. The values of α , defined as in Sec. 2.6, are shown for $\gamma = 0$ (simple Monte Carlo), 1, 2 and 3 (with $\nu_1 = \nu_2 = 1$) as a function of the number m_0 of sweeps over a grid with $N = 64$ points. For clarity of results, the work in first equilibrating the system by 10 V cycles (actually a much smaller number would suffice) is not taken into account in calculating α . Indeed, for small ε this work should be negligible.

The results show the expected slowing down for $\gamma = 0$ and the optimality of $\gamma = 1$. The work on grid $h_i = 2^i h_0$ is $Nm_0(\gamma/2)^i$, so the total work (or #RAN) is $2Nm_0$ for $\gamma = 1$, $Nm_0 \log_2 N$ for $\gamma = 2$ and $Nm_0((\gamma/2)^{\log_2 N} - 1)/(\gamma/2 - 1)$ for $\gamma > 2$. Hence, for $N = 64$ the work is $2Nm_0$, $6Nm_0$ and $20.8Nm_0$ for $\gamma = 1, 2$ and 3 respectively. The relative values of α seen in Fig. 3 almost exactly correspond to this increase of work with γ , showing in comparison to it only a slight initial decrease (with the initial increase of γ for fixed m_0). The slight decrease is due to the increased accuracy of smoother components, which are the only ones to benefit from higher γ (initially; when γ increases further, the additional benefit is negligible).

Generally, for any fixed N α is approximately a constant if $\gamma > .5$, but for $\gamma \geq 2$ this constant increases with N , being $O(N^{\log \gamma / \log 2 - 1})$ for $\gamma > 2$ and $O(\log N)$ for $\gamma = 2$. Figures 4a and 4b demonstrate this.

3. Toward optimal algorithms for Ising models

3.1 The multigrid cycle

A multigrid Monte Carlo method, based on a stochastic coarsening procedure, has been successfully applied to the two dimensional Ising model to inexpensively produce a sequence of statistically independent configurations. See refs. , and for details; the following is a brief description.

Consider the ferromagnetic Ising model Hamiltonian associated with a spin configuration s :

$$\mathcal{H}(s) = - \sum_{\langle i,j \rangle} J_{ij} s_i s_j \quad , \quad (J_{ij} > 0) \quad (3.1)$$

where s_i is the (1 or -1) value of the spin at site i and where $\langle i, j \rangle$ runs over all pairs of nearest neighbor sites on a square doubly periodic lattice. The basic strategy of the algorithm is the stochastic generation of a new, coarser Hamiltonian:

$$\mathcal{H}^1(s^1) = - \sum J_{ij}^1 s_i^1 s_j^1 \quad , \quad (J_{ij}^1 \geq 0) \quad (3.2)$$

with a decreased number of degrees of freedom. Each new “spin” s_i^1 is actually a *block* of one or more spins s_α , so that a spin s_α may belong to one and only one particular block s_i^1 . The process of creating Hamiltonian \mathcal{H}^1 from \mathcal{H} is referred to as going from fine to coarse level, or *coarsening*. Given \mathcal{H}^1 , usual Monte Carlo sweeps can be performed to generate transitions in the phase space of the new level. The process of restoring finer scale degrees of freedom, i.e., interpreting each flipped s_i^1 as a simultaneous flip of all the spins s_α in that block, and then returning to work with \mathcal{H} , is called *uncoarsening*. The description of the coarsening procedure and the organization of the coarsening/uncoarsening steps are given next.

The stochastic blocking is performed by scanning the fine grid interactions $(J_{ij}s_i s_j)$ one by one, in any convenient order, each interaction in its turn being either *kept alive* or *terminated*, according to a certain criterion (given below). If it is “kept alive”, then it actually just stays unchanged. If it is “terminated” then with a probability $P_{ij} = \exp[-J_{ij}(s_i s_j + 1)/T]$ it is *deleted*, i.e., the interaction between the spins is omitted from the coarse Hamiltonian; and with probability $1 - P_{ij}$ it is *frozen*: its two interacting spins are blocked together so that both are flipped simultaneously in all subsequent coarser level moves (including still-coarser level moves subsequently made), until uncoarsening takes place. This particular way of terminating bonds has been introduced by Swendsen and Wang [1], and we will denote it SW. Its particular prescription for P_{ij} guarantees detailed balance and is also such that a bond connecting anti-parallel spins (for which $s_i s_j = -1$) is deleted with probability 1, and thus only parallel spins can be blocked together.

Consequently, by freezing more bonds, blocks of increased sizes are created. The potential size of the produced block, i.e., the number of spins s_α it would include upon freezing a particular candidate bond, has been used as a criterion for deciding whether to terminate that bond or keep it alive; e.g., by adopting the rule of keeping the bond alive iff its freezing would produce a block of more than, say, 4 spins.

When the process of stochastic coarsening is completed, all spins are grouped into two different kinds of blocks: the *disconnected blocks*, each of which is separated from all others by boundaries of deletions and thus have no remaining alive interactions, and the *interacting blocks*, which still have alive bonds between them. The coupling J_{ij}^1 between two interacting blocks s_i^1 and s_j^1 is calculated by summing up all alive bonds connecting them. Starting with $J_{ij} = 1$ on the fine level, stronger couplings ($J_{ij}^1 > 1$) may appear between the interacting blocks. Thus, a new Hamiltonian \mathcal{H}^1 is constructed.

The entire process can be repeated recursively: to be effective, the Monte-Carlo simulation of \mathcal{H}^1 itself includes both conventional Monte Carlo sweeps and stochastic coarsening/uncoarsening steps. The next stochastic blocking is employed by regarding the blocks s_i^1 of the coarse level as now being the spins from which the blocks of the next coarser level are constructed. Each new, coarser block represents a block of blocks s_i^1 and in turn can be referred to as a block of spins s_α of the finest level. Repeating this recursively, a *sequence* of increasingly coarser *levels* is created. Each level k consists of a list of blocks, into which the original set of spins s_α (which in this notation are also s_α^0) is uniquely decomposed. This list is actually a union of two sub-lists: the disconnected blocks created at all finer levels up to and including the current one, and the interacting blocks, denoted s_i^k , which are coupled by the k -level Hamiltonian

$$\mathcal{H}^k(s^k) = - \sum J_{ij}^k s_i^k s_j^k \quad (k = 0, 1, 2 \dots). \quad (3.3)$$

Progressing to increasingly coarser levels, the number of alive bonds keeps decreasing until at the coarsest level none exists. At that stage all bonds are either frozen or deleted, so the original sets of spins s_α is completely decomposed into disconnected blocks.

The entire algorithm can be described as a sequence of multigrid cycles for the finest level, where a multigrid cycle for any given level k (the current “fine” level, assumed not to be the coarsest level) is recursively defined as consisting of the following five steps.

1. N_1 Monte Carlo sweeps are first made on the fine level, using the Hamiltonian (3.3).
2. The next coarser level $(s^{k+1}, \mathcal{H}^{k+1})$ is created from the fine level by the above stochastic coarsening process.
3. γ multigrid cycles for the coarse level are performed. If, however this coarse level is the coarsest, do nothing.
4. Each interacting block s_i^{k+1} whose final value is different from its initial value is translated into flipping all the fine-grid spins s_α^k which belong to it. The spins s_α^k in each disconnected block are flipped simultaneously with probability 1/2.
5. N_2 Monte Carlo sweeps are finally being made on the fine level.

We have used cycles with $N_1 = N_2 = 1$ and $\gamma = 2$. The parameter γ is called the cycle index, and cycles with $\gamma = 2$ are called *W*-cycles.

The obtained results were in a certain respect very satisfactory: the CSD seem to have been completely eliminated, meaning that in a work proportional to the gridsize a new, substantially independent configuration was created. It was later proved [ref to Li & Sokal] that some similar algorithms (in which, however, the

ordering of bond termination is the same in all cycles) must suffer a (very slight) slowing down. This raises the suspicion that the same may be true here, although the proof does not strictly apply to our described algorithm (which generates different bond termination ordering in each cycle: see the above rule for keeping bonds alive). At any rate, the practical behavior exhibited no trace of slowness.

3.2 Dependency of configurations within a cycle

Using a straightforward calculation of measuring the desired averages (observables) on the sequence of configurations produced by multigrid cycles on the finest level converges, to be sure, faster than measurements on the simple (only single-spin) Monte Carlo simulations. But still, as expected, the obtained statistics were slow in averaging out the deviation (from the observable average) exhibited by each configuration. If a standard deviation σ is contributed by the features of some scale, these features have to completely change $O((\sigma/\varepsilon)^2)$ times in order to obtain accuracy ε . The trouble is presumably mostly related to large scale deviations, because only few samples of them are contained in each produced configuration, while small scale fluctuations, one could hope, are effectively averaged out in each given configuration.

The main idea for overcoming this difficulty was to average as many configurations per cycle as possible. Take, for example, the calculation of the mean of the squared magnetization per site $\langle M^2 \rangle$, where $M^2 = \frac{1}{N} \left(\sum_{i=1}^N s_i \right)^2$ and N is the number of spins in the finest level. Instead of measuring M^2 once per cycle, compute it each time the algorithm visits the coarsest level, that is, $2^{\ell-1}$ times per W -cycle, where ℓ is the number of levels. Moreover, (as pointed out in the Appendix of ref. 8), on each visit to the coarsest level it is possible to immediately average over all the different spin configurations referring to their decomposition into disconnected blocks. More precisely, it is easy to show that if there are μ disconnected blocks consisting of n_1, n_2, \dots, n_μ spins, respectively, (where $\sum_{i=1}^\mu n_i = N$), then the average of M^2 taken over the 2^μ equally probable different spin configurations allowed by these blocks is given by:

$$\langle M^2 \rangle = \frac{1}{N} \sum_{i=1}^\mu n_i^2 \quad (3.4)$$

Experiments showed that, disappointingly, near the critical temperature the convergence with such many-per-cycle measurements was *not* substantially faster than with once-per-cycle measurements. We have compared the standard deviation exhibited by the once-per-cycle measurements with that of the average of the $2^{\ell-1}$ measurements (3.4) in one cycle, and found that on grids upto 128×128 the

latter was not even twice as small as the former. Even though this may improve on extremely large grids, not much can be gained in the practical range.

To understand this behavior one should first observe that the average (3.4) of M^2 taken over the 2^μ equally probable configurations is heavily dominated by the largest block, hence they are very strongly correlated. This is particularly true at the critical (or lower) temperatures, where the size of the largest block far exceeds all others.

Consider next the relations between the $2^{\ell-1}$ coarsest configurations obtained within a single cycle. What the experiments show is that they, too, are correlated to each other. Each of them, in other words, strongly depends on both the finest configuration at which the cycle begins and on the first level of stochastic coarsening performed on it. A detailed study showed that the reason is that the stochastic coarsening tends to produce many more deletions along boundaries (between identical-sign regions of the current configuration) than elsewhere. As a result, in all subsequent coarse-grid calculations, there appears a statistical bias to retain many of these boundaries. In particular, the largest block is likely to remain roughly the same.

Thus, in reducing the number of degrees of freedom one loses not just fine-scale fluctuations, but large-scale ones as well.

3.3 Reducing the dependency

At first, this strong correlation between scales appears to be a necessary property of discrete-state models. But then, a similar situation is encountered when *constant* interpolation is used even for the Gaussian model (cf. Sec. 5.3 in [91]). Since the SW termination resembles constant interpolation, the question now is whether a better coarsening technique, capturing some features from *linear* interpolation, can be devised for Ising spins, so as to reduce the dependency between the configurations produced within the same cycle, i.e., by the same coarsening. Let us denote by $\chi_0 = \langle M^2 \rangle$ the true susceptibility, by $\sigma_0 = \langle (M^2 - \chi_0)^2 \rangle^{1/2}$ the standard deviation from χ_0 of M^2 at any single configuration, by χ_1 the average of M^2 for the Hamiltonian \mathcal{H}^1 and by $\sigma_1 = \langle (\chi_1 - \chi_0)^2 \rangle^{1/2}$ the standard deviation of χ_1 from χ_0 . The above coarsening, based on the SW termination, produced $.5\sigma_0 < \sigma_1 < \sigma_0$. The question then is whether a better coarsening technique can produce σ_1 much smaller than σ_0 .

One obvious difference between constant and linear interpolation is that the latter relates a given variable to *two* neighbors, not one. Thus, our first attempt at a linear-like interpolation is to replace the two-spin SW coarsening with the following *three spin coarsening* (3SC).

For simplicity we describe (and have developed and tested) only the case of uniform bonds (constant J_{in}); this is not essential, but introduces simplifying

symmetries. Denote by $\beta = J_{ij}/T$ the uniform thermal binding between neighbors. Consider a spin s_0 with two neighbors, s_- and s_+ say. The current Hamiltonian has the form

$$\frac{1}{T}\mathcal{H} = -\beta s_0 s_- - \beta s_0 s_+ - \dots$$

where the dots stand for any other terms. Three other Hamiltonians are offered as alternatives:

$$\frac{1}{T}\mathcal{H}_1 = -\infty s_0 s_- - a s_0 s_+ - \dots$$

$$\frac{1}{T}\mathcal{H}_2 = -a s_0 s_- - \infty s_0 s_+ - \dots$$

$$\frac{1}{T}\mathcal{H}_3 = -b s_- s_+ - \dots$$

The ∞ value in H_1 (H_2) means that s_0 and s_- (s_+) are blocked together. Note that in H_3 the two bonds between s_0 and its two neighbors are deleted, but a new direct bond is introduced between the neighbors themselves. One selects H_i with probability P_i ($i = 1, 2, 3$), where $P_1 + P_2 + P_3 = 1$. To obtain detailed balance, these probabilities are taken to depend on the current value of s_- , s_0 and s_+ according to Table 1 — plus the obvious rule that $P_i(-s_-, s_0, -s_+) = P_i(s_-, s_0, s_+)$ — and the value of a and b are taken so that

$$e^{2a} = (e^{2\beta} - e^{-2\beta})/(2 - 2p_*)$$

$$e^{2b} = e^{-2\beta}/p_*,$$

p_* being a small positive parameter. We chose $p_* = .15$, but other values in the range $.05 \leq p_* \leq .2$ are perhaps as good.

Table 1

s_-	s_0	s_+	P_1	P_2	P_3
+	+	+	$\frac{1}{2}(1 - e^{4\beta})$	$\frac{1}{2}(1 - e^{4\beta})$	$e^{4\beta}$
+	-	+	0	0	1
+	+	-	$1 - p_*$	0	p_*
+	-	-	0	$1 - p_*$	p_*

The detailed balance of this, and also that of SW and other coarsening schemes, is a special case of the following theorem, which generalizes the Kandel-Domany [51] Theorem.

Detailed Balance Theorem. Let us denote a configuration of a model, $\mathcal{H}(u)$ its Hamiltonian, and $\mathcal{H}_1(u), \mathcal{H}_2(u), \dots$ some alternative Hamiltonians, where

the use of the Hamiltonian $H_i(u)$ also means restriction of the configurations u to a subset where some functionals $F_{i1}(u), F_{i2}(u)$, etc. are frozen. Then, in a Monte Carlo process with current configuration \tilde{u} , replacing $\mathcal{H}(u)$ by $\mathcal{H}_i(u)$ in probability $P_i(\tilde{u}) \geq 0$ maintains detailed balance provided

$$P_i(\tilde{u}) = f_i(F_{i1}(\tilde{u}), F_{i2}(\tilde{u}), \dots) e^{\mathcal{H}(\tilde{u}) - H_i(\tilde{u})}, \quad (3.5)$$

where f_i are arbitrary functions and where

$$\sum_i P_i(\tilde{u}) = 1 \quad \text{for any } \tilde{u}. \quad (3.6)$$

Note that if no real freezing is done with a certain \mathcal{H}_i , i.e., if the only frozen functionals $F_{ij}(u)$ are constant (independent of u), then the coefficient f_i can only depend on i , and not in any way on \tilde{u} (which is the special case proved in [51]).

Proof. Denote by $P(u^1 \rightarrow u^2)$ the probability of obtaining the configuration u^2 at any stage after starting with $\tilde{u} = u^1$. Then clearly

$$P(u^1 \rightarrow u^2) = \sum_k^{\{1,2\}} P_k(\tilde{u}) e^{\mathcal{H}_k(u^1) - \mathcal{H}_k(u^2)} / z_k \quad (3.7)$$

where $\sum_k^{\{1,2\}}$ sums only over such k for which $F_{kj}(u^1) = F_{kj}(u^2)$, ($j = 1, 2, \dots$). But for each such k , by (3.5),

$$\frac{P_k(u^1)}{P_k(u^2)} = e^{\mathcal{H}(u^1) - \mathcal{H}_k(u^1) - \mathcal{H}(u^2) + \mathcal{H}_k(u^2)},$$

hence, by (3.7),

$$P(u^1 \rightarrow u^2) = P(u^2 \rightarrow u^1) e^{\mathcal{H}(u^1) - \mathcal{H}(u^2)},$$

which is the desired detailed balance. ■

We have tested 3SC on an $L \times L$ periodic grid by applying the coarsening step for all triplets s_- , s_0 and s_+ at grid positions $(j, 2k-1)$, $(j, 2k)$ and $(j, 2k+1)$ respectively such that $j+k$ is even. We compared it with an SW coarsening that terminated all the corresponding (s_0, s_-) and (s_0, s_+) bonds. Results at the critical temperature are summarized in Table 2. They show that for 3SC, unlike SW, the ration σ_1/χ_0 decreases with L . This means that if the susceptibility is measured on the first coarse grid, without ever returning to the fine, the average error is small: it tends to 0 as L increases.

Table 2

L	χ_0	σ_0	σ_1 <i>SW</i>	σ_1 <i>3SC</i>
4	12.2		1.8	.7
8	41.4		7.2	1.5
16	139.5	56.8	25.6	4.0
32	470.2	192.5	81.6	10.6

The observation that has led to the construction of 3SC is that the basic flaw in the SW coarsening is the introduction of many deletions, usually clustered along well-defined lines: the lines of current boundaries of spin alignment. These lines therefore exhibit in \mathcal{H}^1 weakened couplings, and are thus likely to persist as boundaries of spin alignment also on coarse grids. This means strong correlation between different coarse grid configurations. In 3SC the introduction of such weakened-coupling lines is minimized.

This is just a first attempt; it all may well be done better. Observe that the blocks created by 3SC are not necessarily *contiguous*: the Hamiltonian E_3 create a bond between s_- and s_+ , so they later may be blocked together without having the points in between, such as s_0 , included in the block. More general schemes may create blocks that are not necessarily *disjoint*. And so forth: the possibilities are many.

It is not clear at this point whether the ideal statistical efficiency is always attainable. What *has* been established, we believe, is that it is possible to benefit greatly from making many measurements at the coarse levels of a multilevel Monte-Carlo algorithm, even in discrete-state models, if a suitable coarsening scheme is used.

3.4 Optimal calculation of T_c

Even though the SW coarsening is not optimal, as explained above, it can still be used in an optimal calculation of certain thermodynamic quantities. As the simplest example of such a quantity we chose the critical temperature T_c itself. The tests reported below indicate that the above multigrid cycle (Sec. 3.1) can directly be used for a very inexpensive, in fact optimal determination of T_c . More precisely, a sequence of increasingly better approximations to T_c is obtained on increasingly larger grids by performing only a few cycles on each. To achieve an accuracy ε in T_c , the amount of computational work turns out to be $\alpha\varepsilon^{-2}$, where the average value of α is about 100.

The algorithm is based on the following measurements. In every W cycle, performed with some temperature T , the algorithm visits the coarsest level many times. At each such visit the domain is completely decomposed into disconnected blocks, and the ratio r between the number of spins in the largest of these blocks and the total number N of spins in the lattice is measured. Denote by \bar{r} the average of these r 's within one W -cycle; clearly $0 < \bar{r} \leq 1$. Generally, if $\bar{r} < \eta$ (η chosen as described below), then we say that the cycle “indicates” a supercritical temperature: $T > T_c$, (or $\beta = 1/T < \beta_c$). Similarly, if $\bar{r} > \eta$ then the cycle indicates a subcritical temperature: $T < T_c$ (or $\beta > \beta_c$).

This definition of T_c is suitable for fairly large grids if η is chosen reasonably small. In fact, in the range we have calculated (grids upto 128×128), $\eta = .5$ has already proved to be small enough. In principle, for much larger grids the definition of T_c should be modified to allow for the fact that \bar{r} becomes small near T_c , even in the subcritical range. One possible modification is to replace r by r' , defined as the ratio between the *length* of the largest block and the length L of the domain. This quantity r' , or its average within a cycle, does not become small near T_c . (The length of a block B can for example be defined as $1 + \max_B |i_1 - j_1|$, where the max is taken over all pairs of points $i = (i_1, i_2)$ and $j = (j_1, j_2)$ such that both i and j are in B . For locating T_c on very large grids it is enough to calculate the *approximate* length of the largest block. The coarsening process can very inexpensively incorporate a procedure that supply each block B on some coarse level with its approximate $\min_B i_1$ and $\max_B i_1$, from which similar quantities can be calculated for all blocks at all coarser levels.) In the practical range of our calculations, however, this more elaborate quantity r' proved unnecessary. What our simpler procedure, based on r , calculates is in fact another thermodynamic quantity, $T(\eta)$, which is the temperature for which the average magnetization per site is η , i.e., the temperature for which $\lim_{N \rightarrow \infty} \langle |\sum_{i=1}^N s_i|/N \rangle = \eta$. But the difference $|T_c - T(.5)|$ is below the accuracy one can obtain with grids upto 128×128 , so we had no motivation to run tests with r' instead of r .

The experiments show that using an $L \times L$ lattice, an interval of roughly $1/L$ around β_c is the best approximation one could get for the critical value; i.e., within that interval, the criterion does not correctly distinguish between sub and super-critical temperatures. To get an approximation twice more accurate, it is therefore necessary to switch to a 4 times larger grid: $2L \times 2L$. (This observation is of course in agreement with the known critical exponent $\nu = 1$, i.e., with the correlation length being proportional to $(T - T_c)^{-1}$.) In order to save work, the algorithm is constructed so that much of the search for T_c on any given grid is carried over from smaller grids. Since each interval is being further corrected by a larger grid's interval, it is not important to check the criterion precisely, and for practical purposes it is sufficient to perform only *one* W -cycle at each temperature. In this way, on each of the grids a computational work equivalent to just few Monte Carlo passes is enough for determining T_c to within an interval roughly as narrow as can ever be obtained on that grid. The interval will get narrower and narrower

as the grid becomes larger and larger, until a desired accuracy is obtained. The details of the algorithm are as follows.

Initialization. Start from a very small random grid — level 0 (say 4×4). Set an initial temperature $T_0 = T_0^0 = 1/\beta_0^0$ (the subscript stands for the level, the superscripts for the sequence of temperatures within each level). For instance $\beta_0^0 = 0$, which is an infinite temperature and hence, for sure, supercritical. Finally, choose some $\Delta\beta_0 > 0$, the step in which the temperature is lowered. (The values used by us are shown in Fig. 5 below.)

For each level $i = 0, 1 \dots$ do the following three steps.

1. Perform one W -cycle (with $\beta = \beta_i^0$) to reach near equilibrium (erasing in particular lower-level periodicity (see 3 below)).
2. Make one W -cycle for each $\beta_i^k = \beta_i^0 + k\Delta\beta_i, k = 0, 1, \dots$ until either $\bar{r} > \eta$ for $\Delta\beta_i > 0$ or $\bar{r} < \eta$ for $\Delta\beta_i < 0$ is obtained for some k . If this condition is already satisfied at the first step (W -cycle) then $\Delta\beta_i \leftarrow -\Delta\beta_i$, i.e., switch the direction of the search.
3. The β_i^k for which \bar{r} has first passed η will be denoted β_i and will serve as our final approximation to β_c on level i . Switch to level $i+1$: its grid is four times larger (factor two in each direction), and its initial configuration is the current configuration on level i extended periodically in each direction (exploiting its doubly periodic boundary conditions). Set $\beta_{i+1}^0 = \beta_i$ and $\Delta\beta_{i+1} = -\Delta\beta_i/2$. Go to step 1 with $i+1$ replacing i .

The step ratio $\Delta\beta_{i+1}/\Delta\beta_i = -1/2$ is reasonable due to the known ratio, mentioned above, between the accuracy obtainable on the corresponding grids. If this ratio were not known (i.e., if we did not know that in this particular model $\nu = 1$), the value of $\Delta\beta_{i+1}$ would have to be determined in an adaptive way, decreasing it faster whenever too many steps were required in the former lattice size.

This algorithm produces an open ended sequence of increasingly better approximations to $\beta : \beta_0, \beta_1 \dots$, as long as required or allowed by computing resources.

It should be noticed that the algorithm saves a lot of work by not insisting on exact measurements of $\langle \bar{r} \rangle$, or even on exact equilibration, at each temperature. It just senses the approximate point at which β_c , the critical temperature, is passed. More accuracy will not help since on each grid that point is only fuzzily defined.

Numerical results. The accuracy ε of β_i , the i^{th} approximation for β_c , is defined by the difference $|\beta_i - \beta_c|$, where $\beta_c = .4406868$ is the known thermodynamic limit. The amount of work, denoted by $\#RAN$, has been measured by us by the number of times a random number is generated. This effectively counts the work in the Monte Carlo sweeps and in the stochastic coarsenings, which are

indeed the most time consuming processes. For an algorithm to be optimal, the quantity $\alpha = \epsilon^2 (\# \text{RAN})$ should be roughly constant, or at least bounded. In Fig. 5 the values of α obtained in our experiments are shown for grid sizes up to 128×128 . Each shown value of α is averaged over an ensemble of about 100 systems.

The results show that the algorithm is not sensitive, within limits, to changes in $\Delta\beta_0$: asymptotically, T tends to the correct value of T_c , and at approximately the same rate α . The behavior of α as a function of the system size can, however, serve as a good indicator for fixing the stepsize $\Delta\beta_0$. Large values of $\Delta\beta_0$ yield fast localization of the first β_i 's, hence small values of α for small L . But these β_i 's are less accurate, hence more localization work is needed later, thus yielding a typical rise in α for large L . Too small values of $\Delta\beta_0$, on the other hand, show much increased values of α for small L . This is explained as follows. A small $\Delta\beta_0$ necessarily implies many steps at each grid size. Not only the amount of work is thus increased, but so does, on small enough grids, the probability of *accidentally* crossing the threshold η too early, implying a bad approximation for T_c , i.e., a large α . For larger grids the test becomes more reliable (accidental crossing is much less likely), hence a smaller step provides a better approximation for T_c and in turn a smaller α .

4. Extensions

The multilevel computational methods for eliminating CSD and for fast calculation of thermodynamic limits, described in the previous Sections, are being extended to increasingly more complicated models. Some initial steps and results are briefly described below. Fuller discussion will appear elsewhere.

4.1 Variable coefficient

4.2 Higher dimensional Gaussian models

The extension of the above algorithm for the Gaussian model from one to higher dimensions is straightforward. Domains of arbitrary shapes and any type of boundary conditions can be treated, with accuracy ε still obtained in $O(\varepsilon^{-2})$ operations.

Local mode analysis, quite similar to the Fourier analysis above, show that in dimension d , for such optimal determination of susceptibility, the cycle index γ should be in the range $2^d < \gamma \leq \min(2^{2p}, 63)$. This implies that for $d \geq 4$ second

order discretization ($p = 2$) is no longer enough.

Experiments were conducted on a rectangular two dimensional domain with homogeneous Dirichlet boundary conditions, using second order discretization ($p = 2$) and cycle index $\gamma = 6$. In this case Fourier analysis can again be used, yielding closed-form expressions for $\langle M^2 \rangle$ and $\langle M_h^2 \rangle$, so the average relative error ε in our statistical determination of these quantities could again be measured conveniently. Denoting again $\alpha = \# \text{RAN} \cdot \varepsilon^2$ (cf. Sec. 2.6), our experiments show, here too, that α is very nearly a constant, independent of ε .

4.3 Approximate coarsening

Whenever the discrete Hamiltonian $\mathcal{H}_h(u^h)$ depends polynomially on u^h (e.g., quadratic dependence, as in (7)) the coarse grid Hamiltonian $\mathcal{H}_H(u^H)$ can exactly be derived, yielding again a polynomial of the same order (e.g., cf. (15)). Simple expressions for \mathcal{H}_H , generally similar in form to \mathcal{H}_h , can also be derived in many other cases by using a low order interpolation I_H^h (e.g., first order; i.e., blocking several fine-grid variables so that their changes by the coarse grid, $I_H^h u^H$, are identical). Such low orders are not always optimal, even when they yield substantial multigrid acceleration. In the Gaussian case, for example, and in many other cases as well, the interpolation order, like the discretization order p , should satisfy $p > d/2$ for optimal results (cf. Sec. 4.1. In the border case $p = d/2$ the results are nearly optimal, with accuracy ε obtained in $O(\varepsilon^{-2} \log(\varepsilon^{-1}))$ operations). Even for removing the critical slowing down, simple blocking is often not enough, and should be replaced by stochastic blocking (e.g., cf. Sec. 3.1). Generally, *approximate* coarsening should then be used. The main principles developed for doing this, and the relation of this to stochastic blocking, are explained below.

For definiteness, the description will be in term of a Hamiltonian \mathcal{H}_h which promotes smoothness in the most usual sense. That is, in the absence of external fields $\mathcal{H}_h(u^h)$ can be written in terms of local differences of u^h , and it attains its minimum when these differences vanish. As a result, the probable changes in the configuration which the usual (local) Monte Carlo process fails to produce efficiently are smooth ones. They can therefore be represented as interpolants $I_H^h u^H$ from a coarser grid function u^H . Many models, such as XY and other $U(n)$ or $SU(n)$ models written in terms of phases, belong to this class. Generalization to gauge fields and other cases are mentioned in Secs. 4.2.1 and 4.2.2 below.

The fine-grid function $I_H^h u^H$, describing the *changes* introduced to the current fine-grid configuration \tilde{u}^h by a coarse-grid configuration u^H , will be called the *displacement* function. Differences of neighboring displacements $((I_H^h u^H)_i - (I_H^h u^H)_j)$ for neighboring sites i and j will be called *strains*. Since the interpolation operator I_H^h is linear, each strain will be a linear combination of *coarse strains* (differences $u_I^H - u_J^H$ for neighboring coarse sites I and J).

Since what we want from the coarse grid is to accelerate *smooth* components, we can aim at a smooth displacement field, hence small strains. The Hamiltonian increment $\mathcal{H}_h(\tilde{u}^h + I_H^h u^H) - \mathcal{H}_h(\tilde{u}^h)$ can then be expanded in a Taylor series in terms of the strains. Truncating this expansion (e.g., retaining only linear and quadratic terms) will give us a functional depending polynomially (e.g., quadratically) on the strains, hence also polynomially (e.g., quadratically) on the coarse strains. The coefficients of the polynomial depend of course on \tilde{u}^h , but this remains unchanged throughout the coarse level processing. The Taylor expansion will also give us *strain limits*, i.e., bounds on each strain—hence bounds on each *coarse* strain—under which the truncation yields a certain accuracy ε_t .

The coarse-strain polynomial approximation to the fine-grid Hamiltonian increment is what we actually define as the coarse-level Hamiltonian $\mathcal{H}_H(u^H)$. Since it is polynomial, its own coarsening, to still coarser levels, is straightforward. In many cases, however, one like the coarse-level Hamiltonian to preserve the form and the topological properties of \mathcal{H}_h . For example, in models where each u_i^h is a phase (or an angular variable) and the Hamiltonian \mathcal{H}_h is given in terms of certain trigonometric functions periodic in these phases, this property does not carry over to the polynomial-type Hamiltonian $\mathcal{H}_H(u^H)$. To correct this, the polynomial-type coarse Hamiltonian \mathcal{H}_H can itself be approximated back in terms of functions (e.g., trigonometric functions) with the original topology. The strain limits (or rather the coarse strain limits) derived above would guarantee that this approximation, too, has $O(\varepsilon_t)$ accuracy. If, through this process, the coarse Hamiltonian is not a polynomial, then its further coarsening should of course use once again the approximation method described herein. Often, indeed, the coarse Hamiltonian has the same form as the fine one, and the same programs can be used at all levels, both for Monte Carlo passes and for coarsening. (This similarity of the levels, and indeed further generality of the coarsening process, sometimes requires an additional shift of the coarse-grid variables to a “Full Approximation Scheme”; cf. Sec. 4.2.1).

Since $\mathcal{H}_H(u^H)$ is only an approximation to $\mathcal{H}_h(\tilde{u}^h + I_H^h u^H) - \mathcal{H}_h(\tilde{u}^h)$, to maintain *detailed balance* the correction $I_H^h u^H$ should have been rejected by the Metropolis-type probability

$$P_{\text{reject}} = 1 - \min[1, e^{\mathcal{H}_H(u^H) - \mathcal{H}_H(0) - \mathcal{H}_h(\tilde{u}^h + I_H^h u^H) + \mathcal{H}_h(\tilde{u}^h)}]$$

(see, e.g., ref.). Instead of executing such rejections, the algorithm only measures P_{reject} and ensures its vanishing in the thermodynamic limit. This is done by having ε_t (the above mentioned bound on the truncation error, which controls the strain limits) suitably decreased together with ε (the measure of approaching the thermodynamic limit, which controls, as in Sec. 2, the size of the finest (largest) level and the amount of samplings performed on each level). The lowering of ε_t can be done without running into slowing down because a given move on a *fixed* level becomes smoother and smoother on the scale of the finest grid as the latter becomes finer and finer, hence this move will be acceptable for smaller and

smaller ε_t as ε is reduced further and further. Other reasons for avoiding slowing down are the updates, the Hamiltonian dependent interpolation and the stochastic appearance of disconnections, discussed next.

The Monte Carlo (e.g., Metropolis) process on each level is constrained by the strain limits, inherited from all the coarsening stages leading from the finest level to the current one. Hence, if the strain limits are approached too closely at some points of some intermediate levels, the process on coarser levels will be completely paralyzed. To avoid this, the algorithm uses *updates*. An update is a return to the next finer grid, introducing there the displacements implied by the current grid solution, and then coarsening again (with displacements and strains being now defined with respect to the *updated* fine grid solution). This updates the current grid equations and “relieves” the solution from strains too close to their limits. An update can be done locally, wherever strain limits are approached. Or, more conveniently and sometimes more effectively, it can be done globally, e.g., after each full Monte Carlo sweep. In principle, while introducing the displacements on the finer grid during an update, some of the finer grid strains may approach *their* limits, requiring an update at still finer level. This, however, seldom happens and cannot cascade to ever finer levels, because moves on any level are very smooth on the scale of much finer levels, and will therefore affect their strains very little. In some models, once in a (long) while a coarse level may cause a “break”, a discontinuous change that requires updating all the way to the finest scales, but such updates are likely to be very local and sufficiently rare.

Due to such updates, large moves are permissible on coarser grids. To make such moves also *probable*, the interpolation I_H^h at each level should be so constructed that its potential displacements are as probable as possible. This implies, for example, that if a certain variable u_i^h is coupled by \mathcal{H}_h more strongly to its neighbors in one direction than in another, then the interpolation to site i should have a proportionately larger weight in that direction. (More generally, if \mathcal{H}_h is quadratic, I_H^h should be constructed by principles similar to those in ref. . If \mathcal{H}_h is not quadratic, the same principles should be applied to the quadratic part of \mathcal{H}_h , which is well defined since the anisotropic \mathcal{H}_h has been obtained from a finer-level Taylor expansion.) It is true that on the finest level the Hamiltonian is usually isotropic and hence I_H^h can be isotropic too; but on coarser levels, due to stochastic variations in \tilde{u}^h at each coarsening stage, the produced Hamiltonians are no longer isotropic, hence anisotropic interpolations should correspondingly be introduced. This tends to create even stronger anisotropy at corresponding points of still coarser levels.

Thus, stochastically, at points of sufficiently coarse grids, the interpolation may become heavily one-sided, similar to the “*blocking*” familiar in discrete-state systems (as described in Sec. 3.1 above). Any such (near) blocking to one side comes together with (near) disconnection to other sides. The interpolation becomes (nearly) trivial and (nearly) entails no strain limits. This stochastically frees sufficiently coarse levels to perform moves that yield changes in the topology

of the configurations.

Note that I_H^h never depends on the current configuration \tilde{u}^h of the current level h . It only depends on the coefficients of \mathcal{H}_h . This is necessary for approaching detailed balance (as $\varepsilon_t \rightarrow 0$). On the other hand the coefficients of \mathcal{H}_h at any intermediate level do depend on the current configuration $\tilde{u}^{h/2}$ at the next *finer* level (as well as on $\mathcal{H}_{h/2}$). This enables the coefficients and I_H^h to depend on stochastic configuration development at all finer levels. In this sense the coarser level forms of move, such as the (near) blocking, develop stochastically. This is essential for making them associated with enough free energy, hence probable.

Similar to the Hamiltonian (or see Sec. 4.2.1 below), other functionals (e.g., the magnetization) can also be approximated on coarser levels, with this approximation, too, improving as ε_t is lowered. This permits averaging these functionals over the many coarse-level samples produced within one cycle, thus greatly reducing statistical errors at little computational cost (cf. Sec. 2).

4.2.1 Functionals not expandable in terms of strains. Many functionals one wants to approximate on coarser levels are not given directly in terms of differences of neighboring variables, and their increment cannot be directly expanded in terms of strains. Additional steps are then needed in the approximate coarsening. The only basic property that should be preserved is that the approximation converges in the limit of “smooth” displacements, where “smoothness” can be controlled explicitly, by generalized “strain limits”, and where the “smooth” limit describes the class of probable large scale moves.

For example, the magnetization $M^h(u^h)$ in the XY model does not depend on differences of neighboring variables. One can easily construct a similar coarse-grid functional M^H such that $M^H(I_h^H u^h)$ approximates $M^h(u^h)$ for any smooth u^h , where I_h^H denotes some fine-to-coarse transfer (e.g., injection, or some local averaging). This functional cannot be used directly, because u^h is not necessarily smooth. However, the difference $M^h(u^h) - M^H(I_h^H u^h)$ is a functional which does depend on local differences, and therefore *its* increment *can* be expanded in terms of strains.

Note that in this case the coarse Hamiltonian, in addition to a polynomial in the coarse strains, will include the non-polynomial functional $M^H(I_h^H \tilde{u}^h + u^H)$. (It thus may be sometimes more convenient to work on the coarse grid in terms of the full-approximation variable $\bar{u}^H = I_h^H \tilde{u}^h + u^H$ than with the incremental variable u^H . This would be the “Full Approximation Scheme” (FAS); cf. Sec. 8 in ref. and Sec. 7.1 in ref.). As a result, the transition to still coarser grid ($2H$) will have to use once again the approximation method of Sec. 4.2, hence impose its own additional strain limits. At sufficiently coarse levels, however, these limits would stochastically open up at many points due to the (near) blocking described above.

4.2.2 Comments on gauge fields. The coarsening of Hamiltonians of gauge fields can follow the principles described above, except that “strains” here will mean phase changes of $I_H^h u^H$ around plaquettes, and “coarse strains” will mean similar changes of u^H around coarse plaquettes.

To be sure, using some choice of low order I_H^h , strains may equal fixed multiples of coarse strains, and an exact coarsening of the Hamiltonian may be possible. Such a scheme can yield substantial Monte Carlo acceleration, but generally will not be optimal. A pre-assigned, non-stochastic coarsening generally will not produce sufficiently probable coarse-grid movements to eliminate all slowing down. Hence, on coarser levels, when the Hamiltonian is no longer uniform, even a low order interpolation should be weighted by the field strength, to produce probable moves. This will make it very costly to operate exact Hamiltonians on all coarser levels. Thus, approximate coarsening, leading as explained above to stochastic blocking, will generally be needed for obtaining accuracy ε in $O(\varepsilon^{-2})$ operations.

4.4 One dimensional XY model tests

Preliminary tests of the above approximate coarsening principles have been conducted on the one dimensional XY model, where it is easy to implement the anisotropic interpolations and straightforward to view all the topological aspects. Calculations of susceptibility were carried out with the periodic Hamiltonian

$$\mathcal{H}_h = -N \left[\sum_{i=2}^N \cos(u_i - u_{i-1}) + \cos(u_1 - u_N) \right],$$

for $N = \frac{1}{h} = 2^2, 2^3, \dots, 2^{13}$ and for various β in the range $2^{-7} \leq \beta \leq 2^{-2}$. The tests clearly show that when ε is lowered (i.e., as larger grids are employed) ε_t can be lowered so that detailed balance is approached (P_{reject} decreases) without running into critical slowing down. Also, topological changes (large-scale loops) are performed by the coarse levels frequently, completely changing the susceptibility values within each cycle. Detailed study of the results is planned.

4.5 Particle simulations and macroscopic equations

The multilevel algorithm with the approximate coarsening procedure described above can be used equally well when the “finest” Hamiltonian \mathcal{H}_h is defined in terms of *particles* (i.e., u_i^h is a vector describing the location of the i -th particle) rather than grid functions. The next coarser level in that case, though, will still be a *grid* function, describing a field of displacements u^H . This grid’s meshsize H is chosen so that each of its cells typically contains a small number (possibly

even less than one on the average) of particles. I_H^h is now an interpolation from the grid to the (old) particle locations, $(I_H^h u^H)_i$ being the displacement (change in location) of the i -th particle. In this way u^H describes collective movements of the particles.

Further coarsening, to still coarser levels ($2H$, then $4H$, etc.) can be produced in the usual way. On sufficiently coarse levels the resulting Hamiltonians will describe *macroscopic equations* of the system. For example, the (discretized) equations of elasticity emerge in this way when the finest level \mathcal{H}_h describes appropriate atomic interactions (§7.2). Generally, the macroscopic equations obtained this way are expected to be much simpler than those derived by group renormalization methods. This is directly due to the slight “iterativeness” left in the process by the updates described above (which however, as explained, only couple neighboring levels). In many cases the need for such updates may tend to disappear on sufficiently coarse levels. Even when this is not the case, an activation of much finer levels during large-scale (coarse level) simulations will only rarely and locally be needed.

The large-scale equations can extend to domains not covered by the finest (grid or particles) level, as explained next.

4.6 Domain duplication

If the fine-grid Hamiltonian uses periodic boundary conditions, instead of using cycle index γ (cf. Sec. 2.3), the domain can be “duplicated” γ times; i.e., the coarse Hamiltonian can be extended periodically to a domain γ times larger. (To extend equally in all d coordinates, $\gamma = 2^d$ can be chosen. If such γ is too large (cf. Secs 2.4.A, 2.7 and 4.1), alternate coarsening levels can use alternate coarsening directions.) No return to finer levels will ever be needed if they have already provided enough statistics (i.e., if the required number of cycles is 1; cf. Sec. 2.4.D). This is possible to do exactly when needed, i.e., when the size of the domain required to yield accuracy ϵ in some calculation is such that it would contain more than $O(\epsilon^{-2})$ finest-grid sites. The computational cost can still be only $O(\epsilon^{-2})$, since the finest level is not employed over the entire domain.

5. Summary

The calculation of an average quantity Q for an infinite system (a “thermodynamic limit” of finite systems) by a Monte Carlo process within an accuracy of $\pm\epsilon$, usually requires the following three factors.

1. First, one should employ a large enough computational lattice $N \times N \times \dots = N^d$, whose linear dimension N should usually increase as ϵ decreases:

$N = N(\varepsilon)$; presumably N grows like $\varepsilon^{-\rho}$, where ρ is positive.

2. On this lattice one carries out a Monte Carlo process which produces a sequence of configurations, each configuration (from a certain point on) appears in its physical probability. Many of these configurations add nothing to the statistical measurement of Q because they depend on each other. The process requires $O(N^z)$ Monte Carlo sweeps, hence $O(N^{d+z})$ computer operations, to create each new, statistically independent configuration. The critical exponent z is of course non negative. The critical slowing down is the case where z is positive.
3. It is not enough to create one independent configuration, because any such configuration has a deviation from Q . If the standard deviation is σ , one would need $O(\sigma^2/\varepsilon^2)$ independent configurations in order to measure Q to the desired accuracy ε .

Taking these three factors together, one would overall need

$$O(N^{d+z}\sigma^2/\varepsilon^2) = O(\sigma^2/\varepsilon^{2+\rho(d+z)}) \quad (29)$$

computer operations in order to obtain an error smaller than ε while measuring Q .

The purpose of the cluster algorithms is to reduce z as much as possible. The purpose of the multigrid techniques presented in this work is to eliminate the entire exponent $\rho(d+z)$ from (29), i.e., to *obtain an error smaller than ε in only $O(\sigma^2/\varepsilon^2)$ overall computer operations.*

This is especially good news for higher (e.g., four) dimensional problems: the work increase with accuracy is essentially independent of the dimension.

It is shown above in detail, especially in Sec. 2, how to achieve such optimal results in some simple cases. The extension to more advanced asymptotically free models is discussed in general terms in Sec. 4.2.

In the d -dimensional model the finest grid contains roughly N^d sites. Thus, the relative discretization error in the susceptibility is $O(N^{-p})$, where p is the discretization order on the finest grid. In order to get relative accuracy ε , in the estimation of the susceptibility, it is necessary that $O(N^{-p}) \leq \varepsilon$, i.e. $O(N^p) \geq 1/\varepsilon$. When this finest grid is also the finest grid in the multigrid Monte Carlo simulation, the work on the coarsest grid is $O(\gamma^{\log_2 N})$ or $O(N^{\log_2 \gamma})$. Hence, to get relative accuracy ε , it is necessary that $N^{\log_2 \gamma} \leq O(1/\varepsilon^2) \leq O(N^{2p})$ or $\gamma \leq 2^{2p}$. On the other hand, in order to keep the work on the coarsest grid dominant, γ should satisfy $2^d < \gamma$. It follows that in the multigrid algorithm the discretization order p should satisfy $p > d/2$. Thus, for optimal performance in a dimension higher than three, one should use a discretization order higher than 2. In four dimensions, the second order discretization is still nearly optimal (taking $\gamma = 2^d$).

The parameters of the multigrid algorithm, such as the cycle index γ and the coarse-to-fine interpolation order, depend not only on the involved model and its

discretization, but also on the measured quantity Q . In d -dimensional models and 1:2 coarsening ratio, for calculating quantities dominated by large scale fluctuations (e.g. susceptibility), $2^d < \gamma \leq 2^{2p}$ must be used, where p is the order of discretization. For quantities dominated by small scale fluctuations (e.g., the energy per degree of freedom), $\gamma < 2^d$ is needed to obtain accuracy ε in $O(\varepsilon^{-2})$ computational work.

Second order (e.g., linear-polynomial) coarse-to-fine interpolation is optimal in the Gaussian model, for measuring either susceptibility or energy. To obtain the same accuracy ε by first order (e.g., constant) interpolation, the work is larger than $O(\varepsilon^{-2})$.

Acknowledgments

We are grateful to R. Benav for useful discussions, and to E. Domany and D. Kandel for a basic correction to our method in Sec. 3. Work supported in part by the Air-Force Office of Scientific Research, United States Air Force under Grants AFOSR-86-0126 and AFOSR-86-0127, by the National Science Foundation under Grant NSF DMS-8704169 and by the Basic Research Foundation of the Israeli Academy of Sciences and Humanities.

References

Figure legends

- Figure 1: Fine grid points are denoted by \bullet and coarse grid points by \times .
- Figure 2: Performance in measuring susceptibility. Each curve shows α (measuring computational work times the square of the obtained accuracy) as a function of the system size N for the indicated value of the cycle index γ .
- Figure 3: Performance in measuring average energy. Each curve shows α (measuring computational work times the square of the obtained accuracy) as a function of the number m_0 of sweeps over a system with $N = 64$ points for the indicated value of the cycle index γ .
- Figure 4a: Performance in measuring average energy. Each curve shows α (measuring computational work times the square of the obtained accuracy) as a function of $1/\varepsilon$, for the indicated values of the system size N and the cycle index γ .
- Figure 4b: same as in Figure 4a with different cycle indexes.
- Figure 5: Each curve shows α as a function of the system size L , for the indicated value of $\Delta\beta_0$.