Numerical investigation of the 3D Hubbard model on a Linux cluster

Isabel Campos *, James W. Davenport

Brookhaven National Laboratory, Center for Data Intensive Computing, Upton, NY 11973, USA

Abstract

We investigate numerically the magnetic properties of the 3D Isotropic and Anisotropic Hubbard model at half-filling on a Linux cluster. The behavior of the transition temperature as a function of the anisotropic hopping parameter is qualitatively described. In the Isotropic model we measure the scaling properties of the susceptibility finding agreement with the magnetic critical exponents of the 3D Heisenberg model. We describe several particularities concerning the implementation of our simulation in a cluster of personal computers paying special attention to the issues related with the parallelization of the algorithm.
© 2003 Elsevier Inc. All rights reserved.

1. Introduction

Many electron systems have been for many years the playground for testing models of superconductivity. In this framework the Hubbard model [1] is expected to reproduce the essential features that purely electronic degrees of freedom are accountable for. In the Hubbard Hamiltonian thermal agitation is modeled via an inter-orbital hopping term, while electrostatic interaction is taken into account via an *effective* Coulomb coupling proportional to the charge density.

The simplicity of this formulation is only apparent. The Hubbard model has (partial) exact solutions only in one dimension. In two or more dimensions only approximate techniques can be applied. Ground state properties and approximate phase diagrams have been analytically derived in the limits $U \rightarrow 0$ and $U \rightarrow \infty$ by using Random Phase Approximation and Strong Coupling expansions, respectively. It is clear that the range of validity of analytical approaches is always an issue, moreover we are interested in the physics of the system at the (more physical) intermediate values of the coupling. In those regions where strong and weak coupling expansions break down Quantum Monte Carlo (QMC) techniques are useful,

---

* Corresponding author. Present address: Leibniz-Rechenzentrum (High Performance Group), Barer Strasse 21, D-80333 Munich, Germany. Tel.: +49-089-289-28721.
  E-mail address: isabel.campos@lrz.muenchen.de (I. Campos).

not only for testing the validity of other analytical methods such as Mean Field, but as a powerful tool to obtain first-principle results.

The numerical investigation of the Hubbard model is however very costly. Away from half-filled bands the path integral measure is no longer positive-definite and Monte Carlo averages of physical observables suffer from large fluctuations. Extracting meaningful results requires extremely large statistics. Even in half-filled bands the complexity of the numerical simulation is high, as we shall discuss throughout the paper. One has to keep in mind that we are dealing with fermionic particles in a region of parameter space in which the Coulomb interaction is big enough for the fermion propagator to be close to singular. Inverting the propagator is a computationally very expensive operation. The most popular algorithm [2], the determinantal method, has a computational complexity which grows with the cube of the system size. Therefore while the system has been studied extensively in two dimensions, in three dimensions the information we have is more restricted.

We decided to investigate in this work the magnetic properties of the three dimensional Hubbard model at half-filling around the Néel phase transition. We concentrate on two issues: the transition temperature, and the universality class of the transition. Pioneering works [3,4] on the phase diagram were carried out in the 1980s on small lattices. The model presents a phase transition line in the plane $\beta - U$ separating a region in which the system is in a disordered paramagnetic phase from another region in which the electron spins are aligned in a staggered way. The actual value of the transition temperature as a function of U remains an open problem. Recent numerical simulations [5] have shown that the phase transition takes place at temperatures much lower than expected from earlier works. Another question to be discerned numerically is the universality class of the phase transition line. Close to the Quantum Critical Point at $T = 0$ the transition should be mean field like (4D Heisenberg exponents). Everywhere else it is expected to be in the universality class of the three dimensional quantum Heisenberg antiferromagnet [6].

In recent years the development of clusters of commercial processors has boosted the computing capabilities at research institutes. This low weight approach to high speed computing is becoming a real alternative to more conventional approaches based on parallel supercomputers traditionally developed by industry. While price is an obvious advantage in favor of *hand made* clusters of PCs, effectiveness is both a model and algorithm dependent issue.

Obvious cases for massive cluster simulation are problems that, while not requiring a huge amount of memory, do require the simulation of many copies of the same program. The speeding up is achieved by setting up different starting conditions for the different copies of the program. On single processors, high performances can be achieved by using the capabilities of recently developed processors, such as vectorization [7]. In this situation the so-called *farm method* on clusters of PCs is a cost-effective, relatively easy to use, source of computer power. The previous statement is only partially true in our case. It is not obvious that the *farm method* is the best solution for the Hubbard model. Warm up times in large lattices might be very large and to have thermalization effects under control it is desirable to have a long single Monte Carlo history.

These and related questions will be addressed here. Simulations have been carried out on the PC cluster at the Center for Data Intensive Computing in Brookhaven National Laboratory. It consists of about 150 Pentium III processors. Clock speeds range from 500 MHz to 1 GHz, and the available DRAM is 1 Gbyte per processor. Communication between the processors is achieved via a commercial Fast Ethernet switch which provides a one-way bandwidth per channel of up to 100 Mbits/s.

The paper is organized as follows: we first briefly review the model and algorithm in Section 2 in order to fix our notations; the numerical simulation, observables and dynamics of the Monte Carlo process is discussed in Section 3. Our results for the Isotropic and Anisotropic Hubbard model are described in Sections 4 and 5, respectively. Section 6 contains conclusions and comments about future perspectives. We left for the Appendix A some particularities related to our implementation of the simulation in a cluster of PCs.

2. Model and numerical algorithm

In the following, we summarize the standard procedure to perform Monte Carlo simulations on the Hubbard Model [2,8,9] and fix our notation. Consider the Hubbard Hamiltonian at half-filling

$$\hat{H} = -t \sum_{\langle ij \rangle, \alpha} \left(c_{i\alpha}^{\dagger} c_{j\alpha} + c_{j\alpha}^{\dagger} c_{i\alpha} \right) + U \sum_{i} \left(n_{i+} - \tfrac{1}{2} \right)\left(n_{i-} - \tfrac{1}{2} \right) \equiv \hat{K} + \hat{V}. \tag{1}$$

Here, the index $i = 0, \ldots, V \equiv L^D$ labels the sites of a lattice of side L in D spatial dimensions on which periodic boundary conditions have been imposed; $c_{i\alpha}^{\dagger}$ and $c_{i\alpha}$ are, respectively, the creation and annihilation operators for electrons with a z-component of spin α at the site i; $n_{i\alpha} = c_{i\alpha}^{\dagger} c_{i\alpha}$ denotes the usual number operator. The sum $\langle ij \rangle$ is over all pairs of nearest neighbors on the lattice. The first term models the thermal agitation, whose strength is characterized by the hopping parameter t; the second term corresponds to an electrostatic Coulomb repulsion of intensity U.

The path-integral representation is obtained by introducing an imaginary time coordinate τ, and considering the action of all possible configurations of the fields between $\tau = 0$ and $\tau = \beta$. The partition function of the equivalent statistical model reads

$$\mathscr{Z} \equiv \mathrm{Tr}\left(\mathrm{e}^{-\hat{S}} \right) = \mathrm{Tr}\left(\mathrm{e}^{-\beta \hat{H}} \right). \tag{2}$$

In order to perform a numerical simulation the theory is defined on a lattice in space and time dimensions. The partition function of the discretized theory can be written as

$$\mathscr{Z} = \mathrm{Tr}\left(\mathrm{e}^{-\Delta\tau \hat{K} - \Delta\tau \hat{V}} \right)^{N_t}, \tag{3}$$

where we defined $\beta \equiv \Delta\tau N_t$, with $\Delta\tau$ lattice spacing in the temporal direction and N_t number of time slices. The role of β is analogous to the inverse of the temperature T of a classical statistical model in $D + 1$ dimensions.

In order to have a well defined relative probability for each configuration in phase space, fermions must be integrated out analytically in the partition function (3). In a first step the kinetic and potential terms are separated in the partition function by the splitting

$$\mathscr{Z} = \mathrm{Tr}\left(\mathrm{e}^{-\Delta\tau \hat{K}} \mathrm{e}^{-\Delta\tau \hat{V}} \right)^{N_t} + \mathcal{O}\left(\Delta\tau^2 \left[\hat{K}, \hat{V} \right] \right). \tag{4}$$

We will neglect the contributions coming from the second term on the r.h.s of Eq. (4). The leading order in the error introduced by the so-called Trotter approximation is proportional to the square of $\Delta\tau$. This systematic error has to be kept under control in actual numerical simulations by choosing $\Delta\tau$ small enough (i.e., smaller than the statistical error).

The kinetic part of the Hamiltonian is a quadratic form in fermion fields, calculating the trace is therefore trivial. In order to write the interaction term as a quadratic form as well one introduces a set of auxiliary boson fields [10]

$$\mathrm{e}^{-\Delta\tau U(n_{i+} - 1/2)(n_{i-} - 1/2)} = \frac{\mathrm{e}^{-\Delta\tau U/4}}{2} \sum_{\sigma_i(l) = \pm 1} \mathrm{e}^{-\Delta\tau \sigma_i(l)\lambda(n_{i+} - n_{i-})}, \tag{5}$$

where $\{\sigma\}(l)$ denotes an Ising field defined in the spatial lattice at the time slice $l = 1, \ldots, N_t$. The constant λ is related to the parameters of the Hamiltonian by the equation $\cosh(\Delta\tau \lambda) = \exp(\Delta\tau U/2)$ for positive λ.

After all these manipulations it is possible to perform the trace in (2) yielding

$$\mathscr{Z} = \sum_{\{\sigma(l)\}} \det \hat{M}^+ \det \hat{M}^-, \tag{6}$$

with

$$\hat{M}^\alpha = \mathbb{1} + \hat{B}^\alpha_{N_t} \hat{B}^\alpha_{N_t-1} \cdots \hat{B}^\alpha_1 \equiv \mathbb{1} + \hat{A}^\alpha(N_t), \tag{7}$$

where we have defined the matrices

$$\hat{B}^\alpha_l = e^{\mp\lambda\Delta\tau\alpha\delta_{ij}\sigma_i(l)} e^{-\Delta\tau\hat{K}}, \tag{8}$$

$$\hat{A}^\alpha(l) = \hat{B}^\alpha_l \hat{B}^\alpha_{l-1} \cdots \hat{B}^\alpha_1 \hat{B}^\alpha_{N_t} \cdots \hat{B}^\alpha_{l+1}. \tag{9}$$

Summarizing, we have substituted the local fermionic interaction in the partition function by an Ising model with a complicated multi-spin interaction. The remaining sum over the Ising field configurations in (6) can be computed by standard Monte Carlo techniques. We adopt here the approach proposed by Blanckenbecler [8], extended to the Hubbard model at low temperatures in [2,9].

We refer to the above cited works for a discussion on the details of the algorithm. In essence the update mechanism is based on the updating of the equal-time Green's function for an electron of spin α propagating through the field created by $\sigma(l)$

$$\hat{G}^\alpha(l)_{ij} \equiv \left\langle \mathscr{T}\left[c_{i\alpha}(l\Delta\tau)c^\dagger_{j\alpha}(l\Delta\tau)\right]\right\rangle = [\mathbb{1} + \hat{A}^\alpha(l)]^{-1}_{ij}, \tag{10}$$

with $\mathscr{T}$ denoting the temporal ordering operator.

The Green's function turns out to be the fundamental object of the simulation since it contains the information needed to update the field $\sigma(l)$. Besides, along the simulation, observables like the energies and the local magnetic moment are calculated as expectation values of certain matrix elements of $\hat{G}(l)$.

The computation of the Green's function is unfortunately also the most expensive part of the algorithm in terms of computing time. The numerical evaluation of Eq. (10) requires performing N_t multiplications of matrices of dimension V. That requires order $N_t \times V^3$ operations plus the inversion of the resulting matrix, which takes of order V^3 operations.

Timings get worse at low temperature since the matrices $\hat{B}_l$ get more and more ill-conditioned when increasing β. The computation of the product in Eq. (10) is then plagued with round-off errors. Obtaining a meaningful result requires intermediate re-orthogonalizations in order to isolate the divergent scales in the matrix product [9]. For very large values of β the Green's function cannot even be calculated in a computer due to finite precision problems.

The situation can be partially alleviated by realizing that Eq. (10) immediately implies

$$\widehat{G^\alpha}(l+1) = \widehat{B^\alpha}(l+1), \widehat{G^\alpha}(l), \widehat{B^{\alpha-1}}(l+1), \tag{11}$$

which can be used to "advance" the Green's function from time slice l to $l + 1$. The significant reduction in number of operations, comes at the price of increasing the round-off errors. Due to this fact Eq. (11) can be used a limited number of consecutive times, say till round-off errors become of the order of the statistical ones. One then has to recompute $\hat{G}(l)$ according to Eq. (10). For the reasons discussed above, at low temperatures, say $\beta > 6$, the range of applicability of Eq. (11) is very limited.

It is clear at this point that to accelerate the simulation we have to concentrate efforts in speeding up matrix operations, in particular the matrix multiplication. We address this point on more quantitative grounds in Appendix A, where our implementation of the algorithm in a cluster is also discussed.

3. The Monte Carlo simulation

We have run numerical simulations on the general anisotropic Hubbard model in $d = 3$

$$\hat{H} = -t \sum_{\langle ij \rangle, \alpha} \left(c_{i\alpha}^{\dagger} c_{j\alpha} + c_{j\alpha}^{\dagger} c_{i\alpha} \right) - tz \sum_{\langle ij \rangle, \alpha} \left(c_{i\alpha}^{\dagger} c_{j\alpha} + c_{j\alpha}^{\dagger} c_{i\alpha} \right) + U \sum_{i} \left(n_{i+} - \tfrac{1}{2} \right) \left(n_{i-} - \tfrac{1}{2} \right). \tag{12}$$

In this notation tz and t represent, respectively, the inter-planar and the in-plane hopping parameters. Varying tz the system undergoes a crossover between the purely two dimensional behavior at $tz = 0$ and the three dimensional isotropic case at $t = tz$. Intermediate values of tz model situations in which the material is better represented by a weakly coupled set of two dimensional layers, than by a three dimensional isotropic lattice.

The phase diagram of the Hubbard model in $d = 3$ contains a phase transition line in the plane $\beta - U$. The high temperature phase is paramagnetic while in the low temperature region the ground state is an antiferromagnet with the electron spins oriented staggered wise in all spatial directions. The limiting behavior of the model for large U corresponds to the three dimensional quantum Heisenberg model. At $U = 0$ the system becomes a gas of non-interacting electrons which shows no transition at all.

Along the numerical simulation we measure the kinetic and the Coulomb energies via the expectation value of the following operators

$$e_{k} = \frac{\mathcal{N}}{4D} \cdot \sum_{i, \hat{\mu}} \left\langle c_{i\alpha}^{\dagger} c_{i+\hat{\mu}\,\alpha} \right\rangle, \tag{13}$$

$$e_{c} = \mathcal{N} \cdot \sum_{i} \left\langle n_{i+} n_{i-} \right\rangle, \tag{14}$$

where the index $\hat{\mu} = 0, \ldots, 5$ denotes the six spatial directions and the normalization factor $\mathcal{N} = 1/(VN_t)$ accounting for the sum over spatial and temporal lattices has been used. We also measure the local magnetic moment defined as

$$S^2 = \mathcal{N} \tfrac{3}{4} \left\langle \left(n_{i+} - n_{i-} \right)^2 \right\rangle. \tag{15}$$

Both, energies and magnetic moment, can be expressed as appropriate combinations of matrix elements of the Green's function.

The Ising variables are coupled to the z-component of the electron spin at each site. Taking this into account we construct the order parameter in the following way. At each time slice l we define

$$m_{\text{stag}}^{l} = \frac{1}{V} \sum_{i} (-1)^{x+y+z} \cdot \sigma_{xyz}, \tag{16}$$

where (x, y, z) are the coordinates of site i. The average over configurations is defined by

$$M_{\text{stag}}^{l} = \left\langle \sqrt{\left(m_{\text{stag}}^{l} \right)^2} \right\rangle. \tag{17}$$

Adding up the contributions of all time slices we get

$$M_{\text{stag}} = \frac{1}{N_t} \sum_l M_{\text{stag}}^l. \tag{18}$$

Along the simulation we have set $t = 1$. In this paper we concentrate on the results for the isotropic case $t = tz$. Some exploratory studies at $tz < t$ have also been performed. In all cases we work at fixed U and sweep through β looking for the value at which the magnetic susceptibility has a maximum. To keep systematic errors under control we use always $\Delta\tau \sim 0.125$. We are interested in the scaling properties of the susceptibility at the phase transition in order to measure the magnetic critical exponents. For this purpose we concentrate most of our statistics at a single value of the Coulomb interaction, $U = 6$. Here we run lattice sizes $L = 4, 6, 8, 10$ with Monte Carlo times ranging from 10^5 for $L = 4$ to 10^4 for $L = 10$ at each β value. We discard between 10% and 20% of the data as thermalization time, depending on the parameter space point and the lattice size. The total computing time spent is the equivalent of about 240 months in a Pentium III processor at 1 GHz.

To assess the statistical quality of our data, following [11] we define the unnormalized autocorrelation function for the observable $\mathcal{O}$

$$C_{\mathcal{O}}(t) = \frac{1}{N-t} \sum_{i=1}^{N-t} \mathcal{O}_i \mathcal{O}_{i+t} - \langle \mathcal{O} \rangle^2 \tag{19}$$

as well as the normalized one

$$\rho_{\mathcal{O}}(t) = \frac{C_{\mathcal{O}}(t)}{C_{\mathcal{O}}(0)}. \tag{20}$$

The integrated autocorrelation time for $\mathcal{O}$, $\tau_{\mathcal{O}}^{\text{int}}$, can be measured using the window method

$$\tau_{\mathcal{O}}^{\text{int}}(t) = \frac{1}{2} + \sum_{t'=1}^{t} \rho_{\mathcal{O}}(t') \tag{21}$$

for large enough t, which is in practice selected self-consistently. We use t in the range $5\tau^{\text{int}}$, $10\tau^{\text{int}}$, and we check that the obtained τ^{int} remains stable as the window in t is increased.

In Fig. 1 we plot the autocorrelation function in $L = 4$ for the staggered magnetization (left side) and for the kinetic energy (right side). We observe the staggered magnetization building up stronger autocorrelations
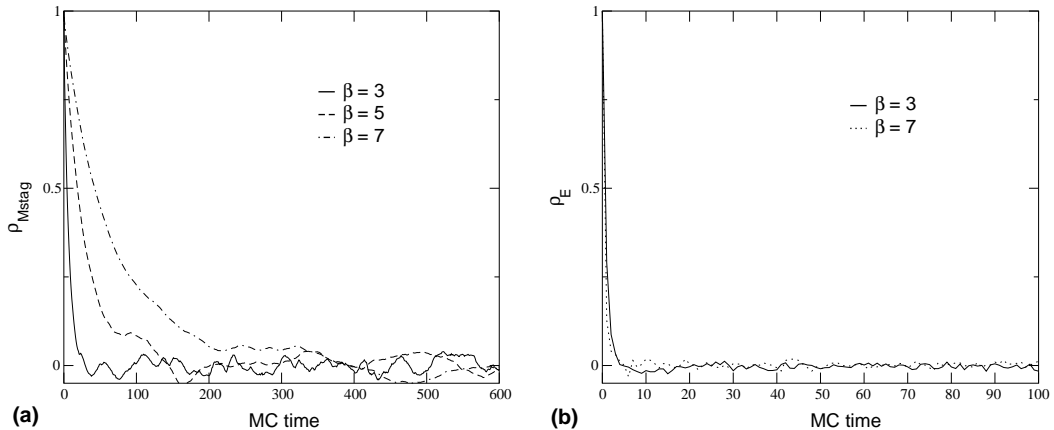

Fig. 1. Autocorrelation function of the staggered magnetization (a) and of the kinetic energy (b). Note the change in the scale between the two plots.

as the phase transition is approached. From this point of view the kinetic energy seems to be insensitive to changes in the temperature. We conclude from here that the order parameter is a better observable to discuss the onset of criticality on the model.

4. Isotropic Hubbard model

We have run numerical simulations in lattices ranging from $L = 4$ to $L = 10$ to investigate the magnetic behavior of the system around the Néel phase transition. For this purpose we have measured the order parameter defined in Eq. (18) and the staggered magnetic susceptibility

$$\chi_{\text{stag}} = V \left\langle \left(\sum_l m^l_{\text{stag}} \Big/ N_t \right)^2 \right\rangle, \tag{22}$$

which is a monotonically increasing function of β for so is $(M_{\text{stag}})^2$.

At low β values, high temperatures, the system is in a disordered paramagnetic phase. The mean value of the order parameter is zero up to corrections proportional to $1/V$. In Fig. 2(a), we plot the MC evolution of M_{stag} at $\beta = 2$ in a $L = 6$ lattice with $U = 6$. As temperature decreases, and the magnetic phase transition is approached, the value of the order parameter increases indicating the tendency of the electron spins to organize themselves in a staggered way. In Fig. 2(b) the MC evolution of M_{stag} at $\beta = 5$ in a $L = 6$ lattice at the same value of U is plotted. We see the system is flipping back and forth between the disordered paramagnetic phase and the staggered ordered one. There is a constant factor $(1 - e^{-\Delta\tau U})^{-1}$ relating the two-point correlation functions expressed in terms of the electron spin with the ones expressed in terms of the Ising fields [10]. Our plots of magnetic variables contain already this factor.

The staggered magnetization is an increasing function of β for fixed U. In fact, for all the lattice sizes we investigated M_{stag} increases with β reaching a maximum at some point, and developing a plateau afterwards. The onset of such plateaus is related to the maximal value that the magnetization can reach in that particular lattice size. In Fig. 3 we display the results for an $L = 4$ lattice for different values of the Coulomb interaction. The values of the observables on the plateaus can be viewed as the asymptotic values in the $T = 0$ limit.

In Fig. 4, we plot those asymptotic values versus the inverse of the lattice size at $U = 6$. Spin wave theory predicts that the fluctuations giving rise to spin–spin correlations decay as $1/L$ [12]. Our results support this prediction for $L > 4$ giving a value for M_{stag} in the thermodynamic limit of 0.15(2). Including $L = 4$ the best
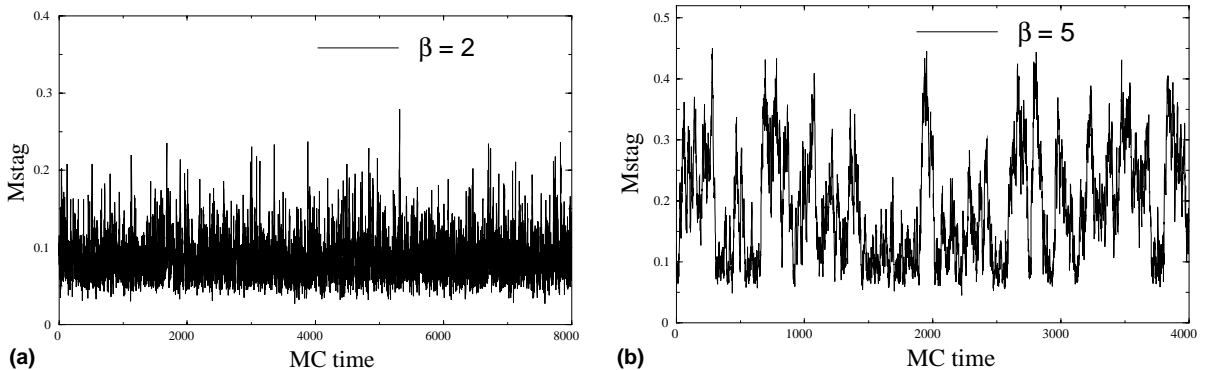

Fig. 2. Monte Carlo evolution of M_{stag} at $\beta = 2$ (a) and $\beta = 5$ (b) for a $L = 6$ lattice at $U = 6$.
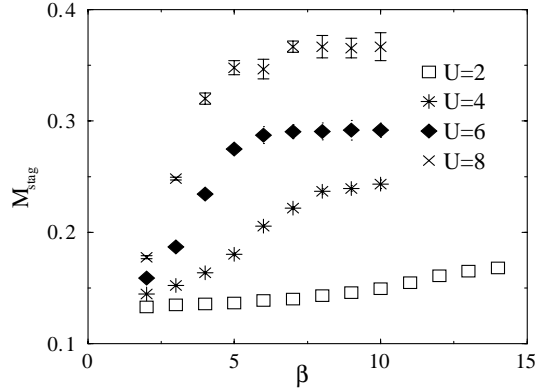
Fig. 3. M_{stag} versus β for different values of U in $L = 4$ at $\Delta\tau \sim 0.125$.
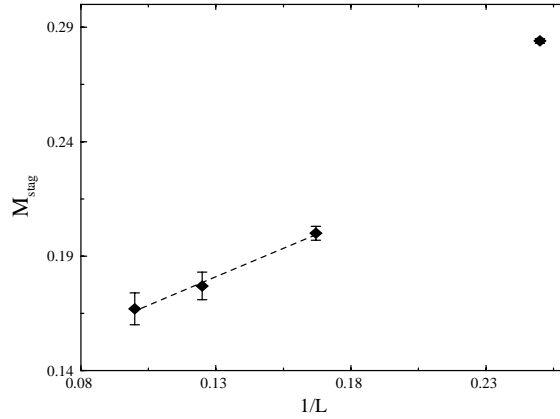

Fig. 4. Asymptotic value of M_{stag} versus $1/L$ at $U = 6$ ($\Delta\tau \sim 0.125$).

fit is a $1/L^2$ extrapolation which leads to a value of M_{stag} compatible with the previous one. We can therefore not give a conclusive answer to this issue but we are inclined to prefer the fit without the small lattice because finite size effects are likely to be uncontrolled for $L = 4$. Indeed, the asymptotic value of $(M_{\text{stag}})^2$ starts to stabilize only from $L = 6$ on.

Next, we focus on the critical behavior of the system. We concentrate our largest statistics in this particular aspect. Our aim is studying the scaling of the order parameter and the susceptibility close to the phase transition temperature, and finally extracting the magnetic critical exponents. A first question arising is the actual value of the transition temperature. The most recent work the authors are aware of [5] quotes a value $T_c \sim 0.3$ for the Néel transition at $U = 6$. This temperature is certainly much lower than the ones reported in the pioneer works of the 1980s (see e.g., [4]). Our purpose is to give an estimation of the critical temperature based on the measurements of the order parameter.

In Fig. 5, we plot the histograms corresponding to the time evolution of M_{stag} at $U = 6$ for increasing lattice sizes at $\beta = 4$ (b) and $\beta = 5$ (a). The asymmetry of the distributions at $\beta = 4$ indicates that we are close to a phase transition. However the system is still clearly in the paramagnetic side because when increasing the lattice size the peak of the magnetization in the paramagnetic region tends to dominate the
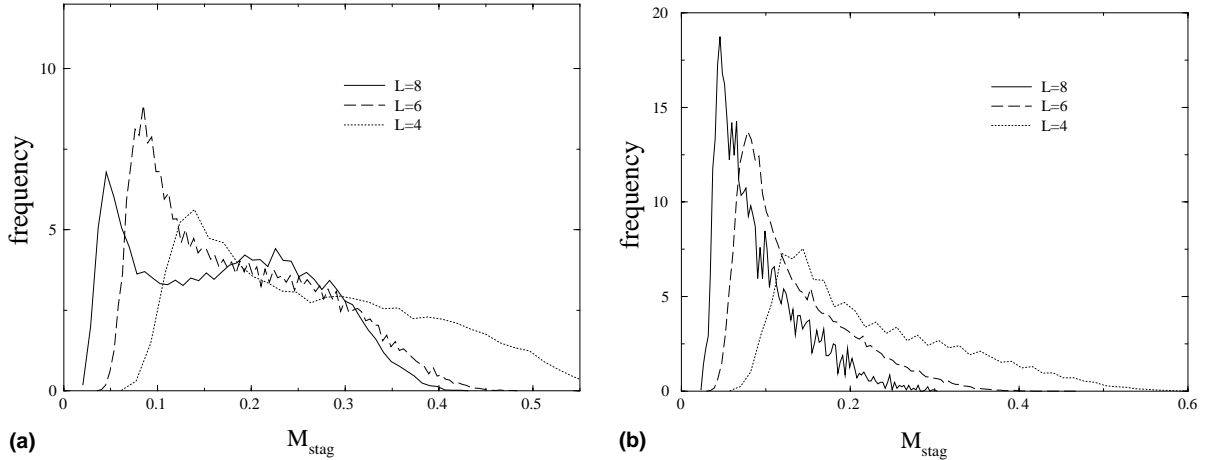
Fig. 5. Normalized distribution of M_{stag} at $U = 6$ at $\beta = 4$ (b) and $\beta = 5$ (a). The Néel phase transition takes place between this two values.

distribution while the other runs away. The behavior is radically different at $\beta = 5$. The peak corresponding to the paramagnetic phase decreases when the lattice size gets bigger, corresponding therefore to a finite size effect. Conversely, the peak in the symmetry broken phase tends to grow. Summarizing, the system is already in the antiferromagnetic Néel phase at $\beta = 5$ because for increasing lattice size the system stabilizes in the staggered phase. This somehow ad-hoc procedure to locate the phase transition has been very useful for us, since the measurement of cumulants of the magnetization turned out to be extremely noisy.

From the order parameter distributions it is clear that the magnetic transition takes place between $T = 0.25$ and $T = 0.20$. Our results therefore support the ones of Staudt and coworkers in the sense that the phase transition occurs at a value much lower than traditionally expected. The fact that we measure a value even slightly smaller might be related to the use of a different observable. In [5] the authors use cumulants of the energy to locate the critical point, while our results are based on order parameter measurements. In principle, different estimators give slightly different results in finite lattices.

The behavior of the staggered susceptibility (Eq. (22)) at $U = 6$ across the parameter space for different lattice sizes is presented in Fig. 6. We observe the susceptibility growing monotonically until reaching a
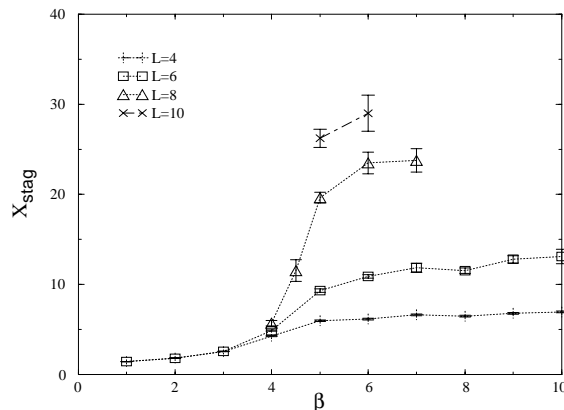

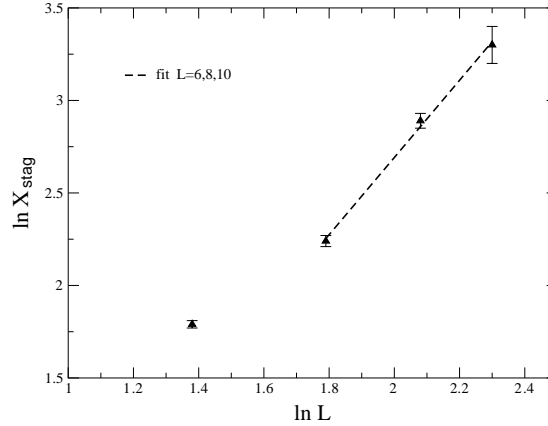Fig. 6. χ_{stag} versus β at $U = 6$ and lattice sizes $L = 4, 6, 8$ and 10.

Fig. 7. Log–log plot of the magnetic susceptibility at $\beta = 5$, $U = 6$ versus lattice size. The linear fit gives a critical exponent $\gamma/v = 2.08(9)$.

plateau reflecting the behavior of the magnetization itself. The saturation of the magnetization, and the subsequent plateau in the susceptibility raises doubts on the interpretation of the peak in the susceptibility as a good observable to locate the phase transition. In principle, there is no reason to believe that the β value at which the susceptibility reaches the plateau has anything to do with the critical temperature of the Néel transition.

We observe that the onset of the plateaus comes close to but clearly after $\beta = 5$ in all cases. Taking $\beta = 5$ as our best estimate for the critical temperature, and using the scaling law

$$\chi_{stag}(T_c) \propto L^{\gamma/v}. \tag{23}$$

Our result for the magnetic critical exponent should agree with the one of the three dimensional Heisenberg model [6], that is, $\gamma/v \approx 1.98$. Our estimation for the magnetic critical exponent is in good agreement with this expectation. In Fig. 7 we plot in log–log scale the values of the susceptibility versus the lattice size. The result of the linear fit gives a value of $\gamma/v = 2.08(9)$.

5. Anisotropic Hubbard model

The introduction of an anisotropic hopping parameter tz allows us to interpolate between the purely two dimensional behavior ($tz = 0$) and the perfectly isotropic three dimensional lattice ($tz = t$). We have done some exploratory studies at intermediate values of tz to get some insight on the crossover behavior of the model.

In Fig. 8 we plot the result for the kinetic energy in $D = 2, 3$ and at several intermediate values of the hopping parameter. The interpolation is smooth in all cases. As a function of tz, the interpolation is linear for the high temperature case, and faster than linear when the temperature is lowered. This can be understood easily taking into account that at low temperatures the correlation between planes is higher and compensates the smaller value of the interplanar tz hopping coupling.

In two dimensions a phase transition to an antiferromagnetic ordered phase is expected at $T = 0$. In principle we can ask for the dependence of the phase transition temperature on the anisotropic hopping parameter tz. Evidently decreasing tz will result in a decrease of the transition temperature. The actual dependence $T_c(tz)$ is an interesting monitor of the form of the quantum fluctuations which disorder the ground state at $T = 0$ in $D = 2$.
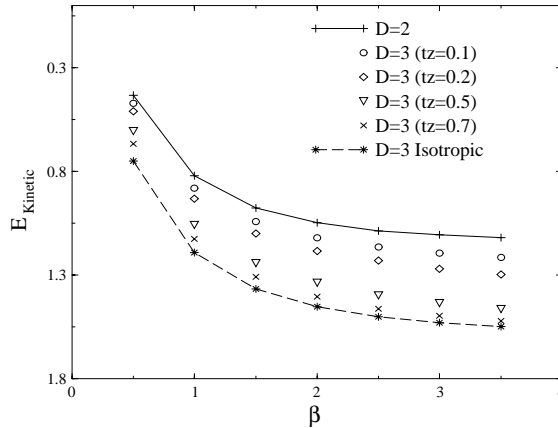
Fig. 8. Kinetic energy for different values of tz at $U = 4$. The values of the energies interpolate smoothly between the $D = 3$ and $D = 2$ limits.

As a first step in that direction we have measured the staggered magnetization for a small lattice, $L = 4$, for different values of tz. The order parameter flips between the disordered phase and the ordered one at $\beta = 6$ producing these distributions plotted in Fig. 9. We can observe that as tz is lowered, the system is the
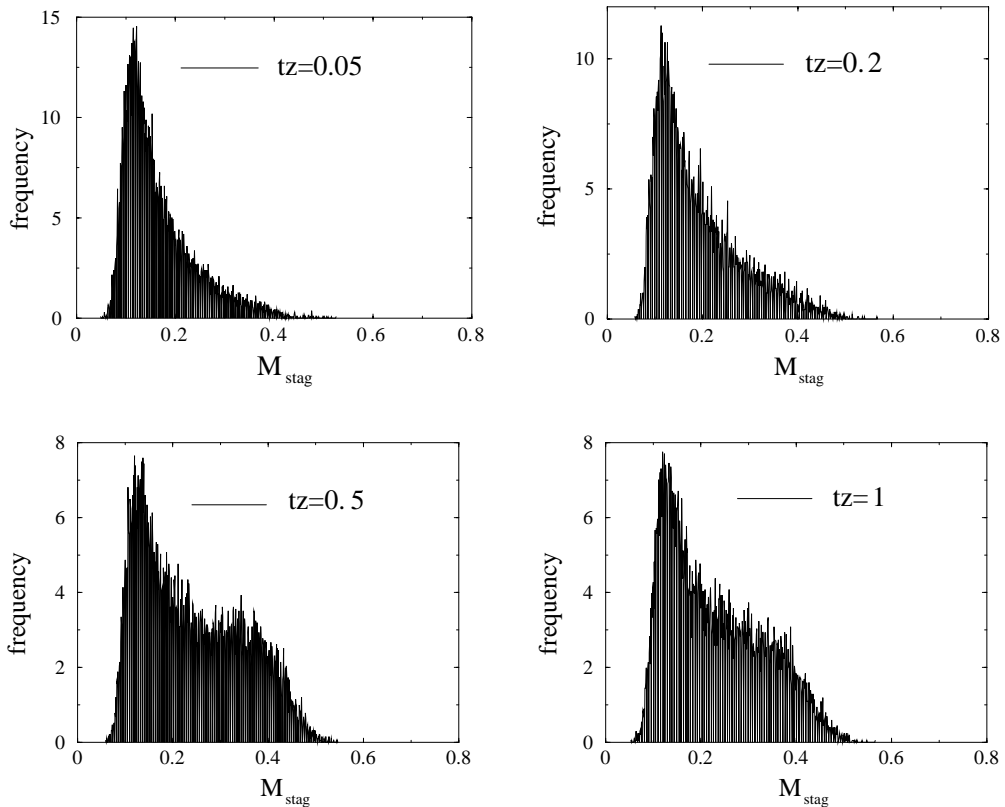

Fig. 9. Distribution of M_{stag} at $\beta = 6$ in a $L = 4$ lattice at $U = 6$ for several values of tz.

more and more in the disordered phase, meaning that the transition temperature in fact decreases as *tz* goes to zero. Eventually, the value of the critical temperature should go to zero, where the quantum critical point of the $D = 2$ Hubbard model is expected to be.

6. Summary and conclusions

We have investigated the properties of the magnetic phase transition in the three dimensional Hubbard model. The measurement of the order parameter allows us to give an estimation of the critical temperature at $U = 6$. From the scaling of the magnetic susceptibility we compute the magnetic critical exponent γ/ν which is in agreement with the magnetic exponent of the three dimensional Heisenberg model.

The dependence of the phase transition temperature on the anisotropic hopping parameter is a very interesting project from the numerical point of view. In this area we are aware of results based on Dynamical Mean Field Theory and the Two Particle Self Consistent approach [13]. It would certainly be of interest to probe such results in a Monte Carlo simulation.

We have shown that computationally costly thermodynamic quantities such as distributions of the order parameter and critical exponents can be computed with moderate-large computing resources using a well known algorithm. This fact should not dismiss a very important issue, which is the development of better core algorithms. A consequence of the impressive development of computer technologies is that we are now able to produce results that we could not have dreamed of only 5 years ago. Such technical developments should go hand by hand with work in algorithm improvement.

Acknowledgements

Appendix A. The Hubbard model on a PC cluster

The question we try to answer here is how to make optimal use of the PC cluster to speed up our numerical investigation of the Hubbard model using the determinantal method. As we pointed out in the introduction the farm method might not been suitable to simulate large lattices due to thermalization issues. The answer to our question is probably that a combination of both, the *farm method*, and a parallel version of the algorithm would do best.

The most serious problem regarding parallelization techniques applied to this algorithm is the extreme non-locality of the update mechanism. Let us suppose that we distribute the Ising spin variables among *np* processors in such a way that each processor takes care of the update of a piece of the spatial lattice. It is easy to show that such strategy could not work. Consider for instance the update of the spin σ_s (we will omit the time index l throughout this section for notation clarity) The update probability depends, among other things, on the diagonal element of the Green's function G_{ss}^{α}. The crucial point is that if the spin σ_s is flipped all the elements of the Green's function change [2]

$$
\begin{aligned}
G_{ij}^{\alpha} &\to f G_{ij}^{\alpha} + G_{is}^{\alpha} \cdot G_{sj}^{\alpha} && \text{if } j \neq s, \\
G_{is}^{\alpha} &\to G_{is}^{\alpha} + f \left(G_{is}^{\alpha} \cdot G_{ss}^{\alpha} - G_{is}^{\alpha} \right) && \text{if } j = s,
\end{aligned}
\tag{A.1}
$$

where f is a constant factor independent of the spatial coordinates. In particular the change affects all diagonal elements which enter in the update probability of all the other σs on the the processors. This implies that every time a spin is flipped the recomputed Green's function should be communicated to all processors. Synchronizing such communication is likely to be impossible. In any case, from a strictly performance point of view it is clear such an strategy could not pay off. The update of the Ising fields within a same time slice must therefore be sequential.

A look at the definition of the Green's function (Eq. (10)) tells us that the update of the different time slices cannot be a distributed task either. The value of the Green's function at a particular time slice l, depends on the state of the Ising spins on all the other time slices. From the previous analysis we conclude that the update process is inherently sequential in all dimensions.

A possibility to still make use of computing cooperation among several processors is to parallelize the matrix operations [14]. We observe in Eq. (A.1) that despite all the matrix elements G_{ij}^α of the Green's function change after a spin flip, such change can be computed from the original G_{ij}^α plus a factor which only depends on the elements of G^α belonging to the row and column of the particular site s being updated

$$
\hat{G}^\alpha = \begin{pmatrix}
\cdots & \cdots & G_{1s}^\alpha & \cdots & \cdots \\
\cdots & \cdots & \cdots & \cdots & \cdots \\
\cdots & \cdots & G_{s-1s}^\alpha & \cdots & \cdots \\
G_{s1}^\alpha & \cdots & G_{ss}^\alpha & \cdots & G_{sV}^\alpha \\
\cdots & \cdots & G_{s+1s}^\alpha & \cdots & \cdots \\
\cdots & \cdots & \cdots & \cdots & \cdots \\
\cdots & \cdots & G_{Vs}^\alpha & \cdots & \cdots
\end{pmatrix}.
$$

The parallelization strategy takes profit of this regularity. The matrix elements of the different operators are distributed column-wise among the processors. When a matrix operation takes place, e.g., a matrix multiplication, each processor computes only the part corresponding to the column it is responsible for. The update of the field σ_s is done simultaneously on all the processors. In principle, since the random number sequence is the same for all of them, the result of the update is the same on all of them. Therefore this mechanism, although redundant, is harmless. In the program, before the update function is called, the processor containing the column G_{is}^α broadcast this column to all the others. As explained before this is the only information needed to recompute the Green's function during the update process. Note that the portion of the s-row needed by each processor is stored locally, and therefore needs not to be communicated.
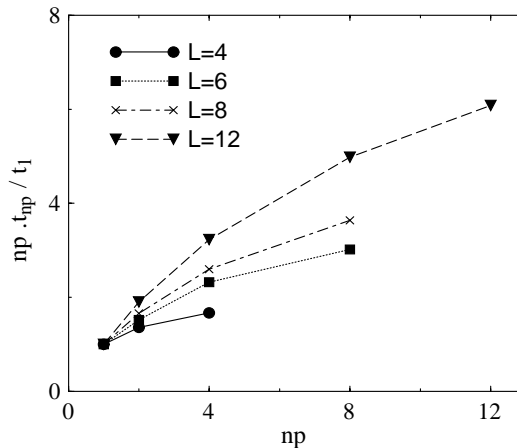

Fig. 10. Speed up of the simulation as a function of the number of processors using the parallelization method described in the text.

In Fig. 10 we show the speed up in the calculation of the Green's function when using the proposed strategy. For small and intermediate lattice sizes the *farm method* is still the best option. For big lattices ($L \geqslant 10$) the algorithm starts to scale reasonably well. It is clear that there is a bottleneck generated by the communication of matrix elements during the calculation. The jam tends to improve when the lattice size gets big because the processors have more operations to perform and therefore do not block the channels trying to submit and retrieve information constantly.

Summarizing, the lesson to extract from here is that the algorithm is parallelizable. Using this scheme instead of a *farm method* pays off for big lattices. It is also clear that we have a very modest switch (Fast Ethernet at 100 Mbits/s).The performance of the parallelization is probably boosted using a more advanced switch for instance a Mirinet^{TM} [15].

References

[1] J. Hubbard, Proc. R. Soc. Lond. A 276 (1963) 238.
[2] S.R. White et al., Phys. Rev. B 40 (1989) 506.
[3] J.E. Hirsch, Phys. Rev. B 35 (1987) 1851.
[4] R.T. Scalettar, D.J. Scalapino, R.L. Sugar, D. Toussaint, Phys. Rev. B 39 (1989) 4711.
[5] R. Staudt, M. Dzierzawa, A. Muramatsu, Eur. Phys. J. B 17 (2000) 411.
[6] A.W. Sandvik, Phys. Rev. Lett. 80 (1998) 5196.
[7] Martin Lüscher, Lattice QCD on PCs, Nucl. Phys. Proc. Suppl. 106 (2002) 21.
[8] R. Blanckenbecler et al., Phys. Rev. D 24 (1981) 2278.
[9] E.Y. Loh, J.E. Gubernatis, in: W. Hanke, Yu.V. Kopaev (Eds.), Electronic Phase Transitions, 1992, p. 177.
[10] J.E. Hirsch, Phys. Rev. B 28 (1983) 4059.
[11] A.D. Sokal, in: M. Creutz (Ed.), Bosonic Algorithms in Quantum Fields on the Computer, Advanced Series on Direction in High Energy Physics, vol. 11, World Scientific, Singapore, 1992.
[12] D.A. Huse, Phys. Rev. B 37 (1988) 2380.
[13] A.M. Daré, Y.M. Vilk, A.-M.S. Tremblay, Phys. Rev. B 53 (1996) 14236.
[14] I. Campos, J.W. Davenport, Wonho Oh (unpublished).
[15] Available from http://www.mirinet.com.