PHASE SEPARATION IN THE 2D HUBBARD MODEL A CHALLENGING APPLICATION OF FIXED-NODE QMC

GIOVANNI B. BACHELET AND ANDREA C. COSENTINI
Dipartimento di Fisica and INFM, Università La Sapienza, Piazzale Aldo Moro 2, I-00185 Roma, Italy

1. Introduction

A quantitative first-principles theory of atoms, molecules, and solids requires an accurate description of the electron-electron correlation. Fahy's [1], Mitáš' [2], and Umrigar's [3] contributions to this book suggest that Variational or Diffusion Monte Carlo simulations are rapidly approaching the required physical or chemical accuracy. However, at least at their present pioneering stage, such simulations are still limited to relatively simple systems, whose correlation energy is a small fraction of the total energy, and, more significantly, whose main physics and chemistry can be qualitatively understood without any reference to many-body theory. By contrast (and by definition), strongly correlated electron systems are those for which even a qualitative understanding of the main physics and chemistry is beyond the independent-electron picture; in this case, as also pointed out by Koch in his contribution to this book [4], rather crude models of real interacting electrons, like the Hubbard model, can be on one hand the only feasible target, and, on the other, a clever way to capture the essential physics of a whole class of materials without going into too many realistic details [5]. However, apart from very special limiting cases, even simple models of *interacting* electrons cannot be solved exactly, because, in spite of many severe approximations, they still give rise to an astronomically large Hilbert space; that's why quantum Monte Carlo methods, based on the idea of stochastic sampling, remain a natural option. In this new context, however, the absence of a preliminary qualitative understanding of the electronic correlations makes the choice of the trial wavefunction (Sect. 2), and possibly even of the size of the simulation box (Sect. 4), a much more crucial and less straightforward part of the investigation. A recent work by Hood *et al.* [9] gives the occasion of an instructive glance at the state of the art, illustrating the different theoretical challenges involved in today's QMC studies of "realistic" and "model" systems. Hood's variational Monte Carlo investigation of bulk silicon, based on Jastrow-Slater wavefunctions and a 54-atom 3D lattice, is more than adequate to obtain an essentially exact ground state, thereby explaining the great success of approximate local-density functionals for this kind of materials (similar considerations apply e.g. to *s-p* bonded materials, like boron nitride [1, 10]). By contrast, when investigating the Hubbard model, the extensive Fixed-Node Monte Carlo simulations by Cosentini *et al.* [11], based on the nodal structure of a correlated Gutzwiller-Slater wavefunction and a (proportionally much larger) 256-atom 2D lattice, may in fact approach the true ground state of the infinite 2D Hubbard lattice with a much better accuracy than most previous attempts, and yet may not be sufficient to proclaim the ultimate word on phase separation. Rather than giving a final answer to this really tough problem, the study presented in the second part of this chapter (Sects. 3 and 4) thus serves the more modest purpose of illustrating a particularly challenging application of the Fixed-Node quantum Monte Carlo method, recalled in the next Sect. 2.

2. Fixed-Node QMC for lattice fermions

In what follows we briefly illustrate some aspects of the Fixed-Node QMC scheme actually followed in the study of the one-band Hubbard model of the subsequent Sects. 3 and 4. We omit all demonstrations, which can be found either in other chapters of this book or in the original papers (in particular Ref. [12] for the power method and Ref. [13] for the fixed-node approximation).

2.1. POWER METHOD ON A LATTICE

Lattice models of quantum many-body systems (see e.g. Ref. [14]) are normally based on the preliminary choice of a tiny single-body basis set attached to each site of a lattice (for example, just one orbital per site in the simplest Hubbard model, or one spin per site in the Heisenberg model). On a finite lattice of N_s sites, the corresponding many-body Hilbert space is also discrete and finite, but its size is a very rapidly growing function of N_s (for example, it's exponential for Heisenberg, combinatorial for Hubbard). Since the purpose of lattice hamiltonians is to model infinite solids, small-N_s lattices (which can be either analytically solved or exactly diagonalized on the computer) may be interesting as a benchmark for new theories and numerical methods (see e.g. Table 1), but the real challenge is to reliably extrapolate the behavior of the infinite model solid from lattices of reasonably large N_s. As mentioned, this amounts to dealing with finite but huge

Hilbert spaces, where quantum Monte Carlo methods are often the only available alternative.

2.1.1. *A sparse matrix whose elements are all positive*

In this context it is possible to project the ground state out of an arbitrary initial state (not orthogonal to it) by the "power method" – the repeated application of some operator $\mathcal{G}$ simply related [15, 16] to the hamiltonian $\mathcal{H}$ (often conventionally called Green's function) – provided that, in the (huge) many-body Hilbert space of the lattice model, one can choose a representation such that (a) each of the basis states $|X\rangle$ is connected only to a few other basis states $|X'\rangle$ by the operator $\mathcal{G}$ (or, in other words, most off-diagonal matrix elements of $\mathcal{G}$ are zero and the matrix $\langle X'|\mathcal{G}|X\rangle$ is sparse), and (b) all the nonzero matrix elements of $\mathcal{G}$ are positive $\langle X'|\mathcal{G}|X\rangle \geq 0$. As described by Nightingale's contribution [17], the condition (a) is needed for the action of the hamiltonian $\mathcal{H}$ (or related [15] operators $\mathcal{G}$) on some arbitrary state to be computationally feasible – the number of operations involved must be much smaller than the size of the Hilbert space [18]; the condition (b) implies that, in the chosen representation, $\langle X'|\mathcal{G}|X\rangle$ can be split into a Markov matrix and a positive weight, and stochastic approaches become possible (see e.g. Ref. [17]); it also implies that the ground-state wavefunction [16], in that particular representation, is positive everywhere $\langle X|\Psi_{\rm o}\rangle \geq 0$, which rules out systems with three or more fermions (*sign problem*). In other words, the power method presented here does not apply to the ground state of lattice fermions unless some *approximation* (recalled in Subs. 2.2) is adopted; it instead directly applies to those lattice models (like e.g. the Heiseberg model on a square lattice [12]) for which both conditions (a,b) are satisfied.

2.1.2. *Markov chain (no weights)*

Let us first suppose for purely pedagogical reasons that, besides (a) being sparse and (b) having all elements greater or equal to zero, the relevant matrix also (c) has the additional, remarkable property $\sum_{X'}\langle X'|\mathcal{G}|X\rangle = 1$. If (c) holds, then $\langle X'|\mathcal{G}|X\rangle$ can be directly identified with a probability (Markov) matrix $P_{X'X}$, whose maximum right eigenvector $|\phi_{\rm max}\rangle$ (i.e. the right eigenvector corresponding to its largest eigenvalue $\lambda_{\rm max} = 1$) can be first approached, and then sampled, by a (sufficiently long) single random walk, which can be generated on the computer by e.g. the following simple rule:

1) take an initial state $|X_{\rm o}\rangle$
2) find all the $N_{\rm o}$ states $|X'\rangle$ for which $P_{X'X_{\rm o}}{=}\langle X'|\mathcal{G}|X_{\rm o}\rangle \neq 0$ (including $X'{=}X_{\rm o}$)

3) evaluate $P_{X'X_o} = \langle X'|\mathcal{G}|X_o\rangle > 0$ for each of the above states, obtaining a list of N_o probabilities which add up to 1
4) align them as N_o adjacent segments adding up to form a unit segment
5) take a random number between 0 and 1 and, depending where it falls on that unit segment, choose the corresponding state $|X'\rangle$ as the new initial state $|X_1\rangle$
6) go back to the first item of this rule and iterate K times this process (with very large K)

The above rule amounts to starting from an initial state $|X_o\rangle$ and to repeatedly (K times) multiplying from the left the matrix $P_{X'X}$. The matrix-vector multiplication is implemented each time in a statistical sense: for example in the first step, rather than yielding a linear combination of all the N_o states $|X'\rangle$ with coefficients $P_{X'X_o}$, the statistical result of the multiplication is just one of them, chosen at random among all the possible N_o states $|X'\rangle$ according to the probabilities $P_{X'X_o}$. After a good number n of such stochastic matrix multiplications, the random sequence $|X_{n+1}\rangle, |X_{n+2}\rangle, \dots |X_K\rangle$ starts to be distributed according to $|\phi_{\max}\rangle$. If the condition (c) holds, then the repeated multiplication of $P_{X'X}$ and the repeated application of the operator $\mathcal{G}$ are the same thing, $|\phi_{\max}\rangle$ is proportional to $|\Psi_o\rangle$, and the random walk $|X_{n+1}\rangle, |X_{n+2}\rangle, \dots |X_K\rangle$ can be directly used for statistical averages over the ground state.

2.1.3. *Single random walk (weights)*

Unfortunately general Markov matrices $P_{X'X}$ are not symmetric (that's why one has to distinguish between left and right eigenvectors), and $\mathcal{H}$ or related hermitian operators $\mathcal{G}$ do not satisfy the condition (c) of previous Subs. 2.1.2: one has, instead, $\sum_{X'}\langle X'|\mathcal{G}|X\rangle = b_X \neq 1$. However, if the property (b) does hold, then $b_X > 0$, and we can still decompose $\langle X'|\mathcal{G}|X\rangle$ into the product of a normalized probability matrix $P_{X'X}$ and a positive weight b_X. Then, by the simple rule illustrated above (steps 1-6) we can still generate a random walk which samples the maximum right eigenvector of $P_{X'X} = \langle X'|\mathcal{G}|X\rangle / b_X$, except that now the eigenvector $|\phi_{\max}\rangle$ is no longer proportional to $|\Psi_o\rangle$, the ground-state eigenvector of $\mathcal{H}$. However, if we keep record not only of the state $|X_j\rangle$, but also of the weight $b_j = b_{X_j} = \sum_{X'}\langle X'|\mathcal{G}|X_j\rangle$ (which must be calculated at each step because it's needed to obtain a Markov matrix P from the original matrix $\mathcal{G}$), then the random walk $\{|X_i\rangle\}$ $(n<i\leq K)$ which samples $|\phi_{\max}\rangle$ can also be used to sample $\mathcal{G}^L|\phi_{\max}\rangle$: one only needs to associate an L-th order cumulative weight $w_i(L) = b_{i-L}b_{i-L+1}b_{i-L+2}\dots b_{i-1}$ to each state $|X_i\rangle$ of the sequence, and then consistently include these weights in the averages [12]. Since the repeated application of $\mathcal{G}$ tends to project the ground-state component out of $|\phi_{\max}\rangle$, a sufficiently large L could allow ground-state averages from a

single random walk [19]. Unfortunately this is not possible because, when L is large enough for $\mathcal{G}^L$ to project the ground state out of $|\phi_{\max}\rangle$, then the weights $w_i(L)$ – as suggested by their product form and argued in better detail elsewhere [12] – undergo wild fluctuations, which grow exponentially with L. So statistical averages are accompanied by a diverging variance and become pointless.

2.1.4. *Statistical averages, many walks, reconfiguration*

The standard solution to the problem of a diverging variance, which occurs with a single random walk, is to propagate many simultaneous random walks. These walks propagate independently most of the time, but not all the time (else they would be equivalent to a single, long walk): to prevent individual weights from getting too small or too large and the variance from blowing up, the different walks must periodically undergo some sort of global reshoveling of states and weights. As recently emphasized by Calandra and Sorella [12], it's this *reconfiguration* process which makes the difference. The results of Sects. 3 and 4 are actually based on their new reconfiguration scheme. A thorough discussion of the underlying theory, and even a complete presentation of their recipe, is beyond the scope of this chapter; it can be found by the interested reader in their recent, original paper [12]. But our numerical experience may be worth mentioning: in practice the Calandra-Sorella reconfiguration scheme allowed us to reproduce with a relatively small fixed number ($100 \sim 200$) of walkers equally accurate results as those obtained by means of standard branching schemes [20] with more than 2000 walkers.

2.2. THE SIGN PROBLEM

As mentioned in the previous Subs. 2.1, the power method can work as a ground-state [16] projector only if all the matrix elements of the operator $\mathcal{G}$ are positive [15], or, equivalently, if the off-diagonal matrix element of the hamiltonian are all negative [15, 16].

2.2.1. *Fermion nodes and other possible sources*

The above condition on matrix elements can be violated for various reasons: (i) particle statistics: the ground-state wavefunction of identical fermions changes sign under their exchange, and must be negative somewhere; (ii) one-particle effects: in certain lattice models, some hopping terms have the “wrong” sign (on continuum “realistic” models a similar problem may arise from nonlocal pseudopotentials [21] as illustrated by Mitáš in this book [2]); (iii) other effects, like the frustration of a quantum spin system on a triangular lattice [22]; (iv) combinations of the above. From this point of view our simplest 2D one-band Hubbard model on a square lattice falls in the first category (i), since its sign problem is entirely due to the Fermi statistics of the electrons [23].

2.2.2. *Importance sampling and Fixed-Node approximation*
As clearly explained in at least one of the other chapters of this book [17], the use of importance sampling is of crucial importance to improve the sampling efficiency; a trial function may also be the simplest way to incorporate the Fermi statistics in a basis of "labeled fermion" lattice configurations [13, 23]. So in actual fermion calculations a good trial function Ψ_T (typically available from preliminary Variational Monte Carlo calculations) is always adopted. In the power method, importance sampling amounts to replacing $\mathcal{G}$ by $\tilde{\mathcal{G}}$, whose matrix elements are $\tilde{\mathcal{G}}_{X'X} = \Psi_T(X')\langle X'|\mathcal{G}|X\rangle\Psi_T^{-1}(X)$, and then to consistently deal with $\tilde{\mathcal{G}}$ rather than $\mathcal{G}$ [17]. The sign problem amounts then to the occurrence of some unwanted $\tilde{\mathcal{G}}_{X'X} < 0$, and, apart from transient estimates which are rather hazardous for large systems, the only cure to this problem seems to be the fixed-node approximation recently proposed by An, Bemmel, Ceperley, Haaf, Leeuwen, and Saarloos [13, 24]. In this approximation the true hamiltonian $\mathcal{H}$ is replaced by an effective hamiltonian $\mathcal{H}_{\rm eff}$, defined as follows:

$$\begin{aligned} \langle X'|\mathcal{H}_{\rm eff}|X\rangle &= 0 && \text{if } \tilde{\mathcal{G}}_{X'X} < 0 \\ &= \langle X'|\mathcal{H}|X\rangle && \text{(otherwise)} \end{aligned} \tag{1}$$

for the off-diagonal terms $X \neq X'$, so that hops from the state $|X\rangle$ towards "sign-flipping states" $|X'\rangle$ (i.e. such that $\tilde{\mathcal{G}} < 0$) are now forbidden, and, for the diagonal terms,

$$\langle X|\mathcal{H}_{\rm eff}|X\rangle = \langle X|\mathcal{H}|X\rangle + \sum_{X'}^{sf}\langle X'|\mathcal{H}|X\rangle\frac{\Psi_T(X')}{\Psi_T(X)}, \tag{2}$$

the summation being over all the "sign-flipping states" $|X'\rangle$ such that $\langle X'|\mathcal{H}|X\rangle \neq 0$ but $\tilde{\mathcal{G}}_{X'X} < 0$: whenever in the "neighborhood" of the state $|X\rangle$ there are sign-flipping states $|X'\rangle$, appropriate repulsive terms are added to the old diagonal element $\langle X|\mathcal{H}|X\rangle$. Such a diagonal repulsive term has the intuitive effect of repelling the random walk away from the regions of the Hilbert space where $\tilde{\mathcal{G}}$ would tend to revert its sign, and also the less intuitive, but more important effect of producing, in analogy with the continuum case [25, 26], an upper bound for the true ground-state energy and a variational approximation [13]. This formulation, contained in Ref.[13], extends the original prescription of Ref. [24] to include all the possible sources of a sign-flip in $\mathcal{G}$ (Subs. 2.2.1) [27]. However for our simple Hubbard hamiltonian (Eq. 3) the original and new prescription are coincident, because in this case the only source of sign flips are the Fermi nodes [23]. Although the original papers by An, Bemmel, Ceperley, Haaf, Leeuwen, and Saarloos [13, 24] are only 3-4 years old, their "fixed-node Monte Carlo for lattice fermions" has been already used by other authors, like Boninsegni for frustrated Heisenberg systems [22], or Koch *et al.* for orbitally-degenerate Hubbard models [4, 28]. We (Cosentini *et al.* [11]) have chosen it for the extensive study of the 2D Hubbard model presented in the coming Sections 3,4, because it allows a reliable investigation of (previously unfeasible) very large lattice-fermion systems.

3. 2D Hubbard model and phase separation

3.1. MOTIVATIONS

3.1.1. *High-T_c, phase separation, ICDW*

Strongly correlated electrons and holes are expected to play a key role in the high-T_c superconductors. Their possible instability towards phase separation, initially believed to inhibit superconductivity, is attracting a lot of interest since a few different authors [29, 30, 31] have pointed out that such a tendency may in fact be intimately related to the high-T_c superconductivity. Long-range repulsive interactions may turn the phase-separation instability into an incommensurate charge-density-wave (ICDW) instability, and the very existence of a quantum critical point associated with it may be a crucial ingredient of the superconducting transition [32]. Phase-separation and/or ICDW instabilities are related to a substantial reduction of the kinetic energy, which otherwise tends to stabilize uniformly distributed states; such a reduction is typical of strongly correlated electrons, both in real and model systems.

3.1.2. *What is phase separation?*

Phase separation is a thermodynamic instability associated with the violation of the stability condition $\chi^{-1} = \partial^2\mathcal{E}/\partial n^2 > 0$, which requires the energy density $\mathcal{E}$ of an infinite electronic system to be a convex function of the electron density n. If such a violation occurs in the density range $n_1 < n < n_2$, then the system will separate into two subsystems with electron densities n_1 and n_2. For the two-dimensional t–J and Hubbard models a phase separation, if any, is expected to occur in a density range close to half filling ($n \simeq 1$), and to yield a hole-rich phase with electron density $n_1 < 1$ and a hole-free phase with electron density $n_2 = 1$ [33]. In view of our further developments, the same physics can be conveniently rephrased in terms of the hole density $x = 1 - n$ and energy density $e_h(x) = \mathcal{E}(n(x))$, as done in the upper right panel of Fig. 1.

3.1.3. *Maxwell's construction*

In a truly infinite system a phase separation of the type just discussed would be associated to a vanishing inverse compressiblity χ^{-1} in the whole density range $n_1 < n < n_2$; in a finite system χ^{-1} may even become negative, because of surface effects; the same applies of course to the the second derivative of the hole energy density $e_h(x)$ with respect to the hole density x. So for finite systems it's preferable to pinpoint the phase separation using Maxwell's construction (Fig. 1, upper left panel). It has been shown in Ref.[33] that Maxwell's construction is equivalent to study, as a function of the hole density $x=1-n$, the quantity $e(x)=[e_h(x)-e_H]/x$, i.e. the energy per hole $e_h(x)$ measured relative to its value at half filling $e_H = e_h(x=0)$ (Fig. 1, lower panels). For an infinite system, if the inverse compressibility χ^{-1} vanishes between some critical density x_c and half filling $x = 0$ (upper right), then the function $e(x)$ is a constant for $0 \leq x \leq x_c$, and the fingerprint of a phase separation is thus a horizontal plot of $e(x)$ below x_c (lower right). For a finite system, instead, the inverse compressibility χ^{-1} may become negative (upper left) and the fingerprint of a phase separation is a minimum of $e(x)$ at $x = x_c$ [33] (lower left). On finite lattices, where only a discrete set of densities is available and a statistical error bar usually accompanies the estimated QMC energies, the plot of the function $e(x)$ (lower panels of Fig. 1) is supposed to allow an easy visual judgement of the presence or absence of such an instability. The attempt of telling whether near half filling a few electron or hole energies are better interpolated by a straight or by a somewhat curved line (upper panels of Fig. 1) is definitely more difficult than looking at $e(x)$. However even $e(x)$ can give a reliable criterion only for medium-large finite systems; really small systems (for which most numerical results have been up to now available) can attain so few and coarse densities, and suffer from so large finite-size errors, that their predictions of the relevant trends remain largely inconclusive, no matter which energy function we look at. More remarks on $e(x)$ vs. phase separation will be presented later on, when examining the numerical results.

3.1.4. *Limited experimental evidence*

Phase separation has been experimentally observed in $La_2CuO_{4+\delta}$[34, 35], where the oxygen ions can move: in the doping interval $0.01 \leq \delta \leq 0.06$ the compound separates into a nearly stoichiometric antiferromagnetic phase and a superconducting oxygen-rich phase. In generic compounds, where charged ions cannot move, the possibility of a macroscopic phase separation is spoiled by the long-range Coulomb repulsion, and should lead to an ICDW instability [36]; here, however, the identification of charge inhomogeneities with spoiled phase separation is less obvious [37].

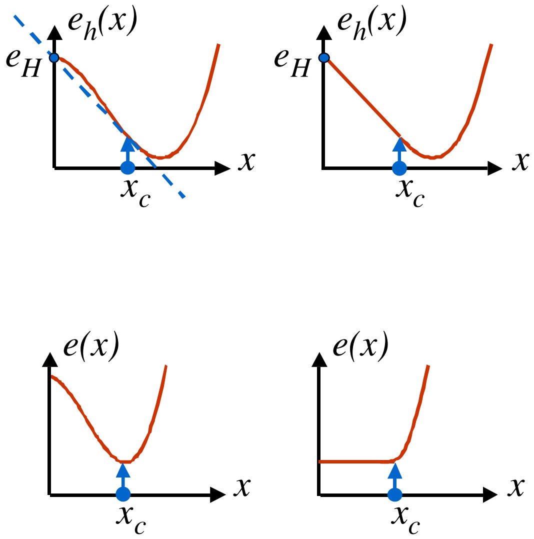

Figure 1. Sketch of a hypothetic system of holes which, when its density falls between zero and a critical value x_c, tends to phase-separate between a hole-free phase (zero density) and another hole-rich phase (density equal to x_c). Upper panel: behavior of the hole energy density e_h vs. the hole density x in the case of a finite (left) and an infinite (right) system. For the finite system (left), between $x = 0$ and $x = x_c$, the actual energy of the phase-separated system (dashed straight line, Maxwell's construction) is lower than the energy of the homogeneous system; for the infinite system, in the same density range, the energy density is a straight line, and the inverse compressibility vanishes. Lower panel: behavior of the function $e(x) = [e_h(x) - e_H]/x$ for a finite (left) and an infinite (right) system. As shown by Emery *et al.*, in a finite system the study of $e(x)$ is completely equivalent to Maxwell's construction (see text).

3.1.5. *Theoretical controversy*

On the theoretical side, evidence for phase separation has been suggested for various models of strongly correlated electrons, such as the $t-J$ model [33], the three-band Hubbard model, the Hubbard-Holstein model and the Kondo model (see e.g. Ref. [32] and references therein). Despite intensive studies, even for simple models there is no general agreement on the phase separation boundary: for the very popular $t-J$ model, phase separation is fully established only at large J, but at small J (which unfortunately happens to be the physically relevant case) theoretical and numerical results are quite controversial. Emery *et al.*'s [33] theory that phase separation occurs at *any* value of J in the t–J model is confirmed by a recent numerical study by Hellberg and Manousakis [38], but is in contrast with Dagotto *et al.*'s [31] exact numerical results on small clusters, suggesting no tendency toward phase separation for both the Hubbard model and the $t-J$ model below a critical value $J < J_c \sim t$, and with Shih *et al.*'s [39] numerical results. We also mention the recent suggestion by Gang Su [40], according to which the Hubbard model does not show phase separation for any value of U/t at any temperature.

3.1.6. *A challenging application of fixed-node quantum Monte Carlo*
Given the above theoretical and experimental situation, the availability of the fixed-node quantum Monte Carlo, described in the above Subs. 2.1, has recently provided us (Cosentini *et al.* [11]) with a strong motivation to further investigate the Hubbard model. Whether the 2D Hubbard hamiltonian, a prototype for interacting electrons with no long-range repulsion, shows any instability towards phase separation, is by itself an interesting open question. Besides the intrinsic interest in the model's behavior, a numerical investigation of this question may also shed some indirect light on two related issues: (i) whether phase separation is likely to occur in the $t-J$ model at small J (the "physical region", where the ground-state energy difference with the Hubbard model should become negligible); and (ii) whether the measured phase separation of some real high-T_c superconductors (Subs. 3.1.4) is such a fundamental and elementary fact as to be explained by the simplest one-band Hubbard hamiltonian.

3.2. NUMERICAL STRATEGY AND TESTS

3.2.1. *Hamiltonian*
To obtain numerical estimates of functions like those sketched in Fig. 1 we need to to evaluate the ground-state energy of the Hubbard hamiltonian

$$\mathcal{H} = -t \sum_{\langle i,j \rangle \sigma} (c^\dagger_{i\sigma} c_{j\sigma} + h.c.) + U \sum_i n_{i\uparrow} n_{i\downarrow} \tag{3}$$

for many large finite lattices and many different electron densities. As already mentioned (and easily double-checked by the reader), for fermions this hamiltonian $\mathcal{H}$ and the corresponding importance-sampled $\tilde{\mathcal{G}}$ satisfies the requirement of sparseness (Subs. 2.1.1 and Ref. [18]) but not the one on sign (Subs. 2.2 and Ref. [23]). The choice of a good variational wavefunction is thus needed not only to improve the sampling efficiency, but also for the power method to become feasible through the fixed-node approximation (Subs. 2.2.2).

3.2.2. *Choice of variational wavefunction*
The variational wavefunctions we use to guide the random walks and to fix the nodes are the product of a Gutzwiller factor and two Slater determinants of single-particle, mean-field wavefunctions (with uniform electron

density n and uniform staggered magnetization density m) for an equal number of up- and down-spin electrons [41]. Preliminary estimates of the optimal Gutzwiller parameter $0 < g < 1$ and mean-field wavefunctions (parametrized by the staggered magnetization $0 < m < 1$) were obtained, for each U/t and density, by variational Monte Carlo (VMC) runs; they are shown in Fig. 2 as a function of the electron density n for three choices of the coupling strength U (in units of the hopping parameter t). As U/t increases, the optimal g decreases, reflecting a stronger depression of doubly-occupied configurations. As far as the density dependence is concerned, the staggered magnetization m is close to zero for all U/t and low-medium electron densities, but sharply raises above $n \sim 0.75$, thus signaling, at the variational level, the system's tendency towards an antiferromagnetic ordering as half filling approaches. The optimal Gutzwiller factor steadily decreases as a function of the electron density n, except for very high electron densities ($n \simeq 0.9$ and above), where it apparently goes up again, as also noticed elsewhere [42]. In this narrow density range, however, the energy minimum appears to be flat: slightly different g values give essentially the same VMC energy, and the corresponding trial wavefunctions also yield the same energy at the fixed-node level.

3.2.3. *Comparison with exact results*

A few representative variational (VMC) and fixed-node Monte Carlo (FN-MC) energies are shown in Table 1 for the 4×4 Hubbbard lattice, for which exact results [43] are available. As expected, the VMC energy is always above the fixed-node energy, which in turn, for these coupling strengths, is slightly ($\sim 3\%$) above the exact energy. For comparison we show the Constrained-Path Monte Carlo (CPMC) energies of Zhang *et al.* [44, 45], which also include a larger 16×16 lattice (last row). Especially at large U/t our results appear to be of comparable quality to theirs. As far as the 4×4 results are concerned, we notice that for $N_e = 10$, which corresponds to a closed-shell configuration, both fixed-node and CPMC are much closer to the exact energy than for $N_e = 14$, which corresponds to an open-shell configuration. This could be a serious problem when numerically studying the behavior of the energy as a function of the density; the results presented here suggest that, for lattices larger than 12×12, the shell effects become almost irrelevant [47].

3.2.4. *How to vary the density?*

To study the energy as a function of the electronic density we have first tried out the less usual way of varying the density suggested in Ref. [38] to "avoid spurious Fermi-surface shape effects": keep the number of electrons N_e fixed while the size of the underlying lattice is varied. But we discovered

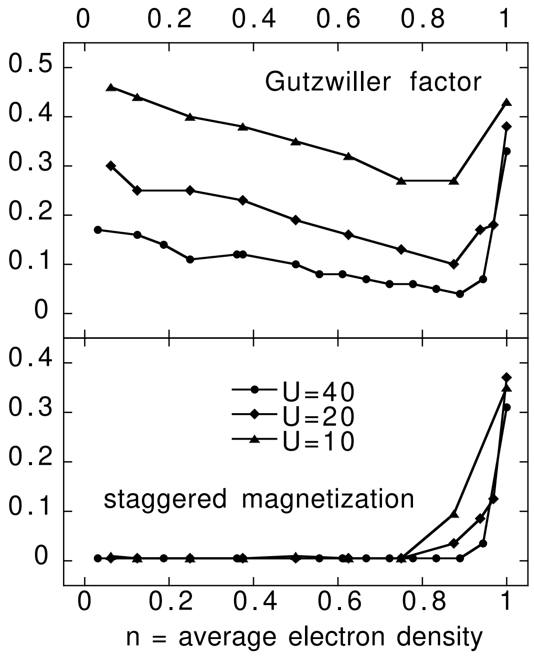

Figure 2. A collection of optimal Gutzwiller factors $0 < g < 1$ (upper panel) and staggered magnetizations $0 < m < 1$ (lower panel) for many different electron densities and three values of the coupling strength U/t: 10 (triangles), 20 (diamonds), and 40 (dots). Most results correspond to 16×16 lattices, but a subset corresponds to 12×12 and to other large but different lattices. When the size is large enough, the optimal g, m appear thus to be almost independent of the size.

TABLE 1. Ground-state energy per site (in units of the hopping parameter t) for a 4×4 Hubbard lattice and various values of U/t. N_e is the number of electrons and n is the corresponding average density. VMC: variational Monte Carlo, Ref. [11]; FNMC: Fixed-Node Green's function Monte Carlo, Ref. [11]; CPMC: Constrained-Path Monte Carlo, Ref. [44]; EXACT: exact diagonalization results, Ref. [43].(see text).

size	N_e	n	U/t	VMC	FNMC	CPMC	EXACT
4×4	10	0.625	4	$-1.211(2)$	$-1.220(2)$	$-1.2238(6)$	-1.2238
4×4	10	0.625	8	$-1.066(2)$	$-1.086(2)$	$-1.0925(7)$	-1.0944
4×4	14	0.875	8	$-0.681(2)$	$-0.720(2)$	$-0.728(3)$	-0.742
4×4	14	0.875	12	$-0.546(2)$	$-0.603(2)$	$-0.606(5)$	-0.628
16×16	202	0.789	4	$-1.096(2)$	$-1.107(5)$	$-1.1193(3)$	$-$

that by this prescription, if the number of electrons is really small (e.g. $N_e = 16$), then artificial changes in the convexity of the curve may occur. If, instead, the system is large enough (e.g. 12×12 lattices or larger), then it doesn't matter how the density is varied, as this prescription and the usual prescription give the same results. So for our systematic study (many densities and three U/t values) we finally adopted a large 16×16 lattice ($N_s = 256$ sites), and varied the number of electrons N_e to yield electronic densities $n = N_e/N_s$ ranging from empty $n=0$ to half filling $n=1$.

3.3. RESULTS

3.3.1. *Energy vs. density for large lattices*

In Figs. 3-6 we show the electronic ground-state energy per site, obtained by FNMC runs as a function of the electron density [46]. Energies are in units of the hopping parameter t throughout this paper; the statistical errors are smaller than the marker size, and thus are not visible. The calculated points are shown as full markers for closed shells, and as empty markers for open shells. From the comparison of Fig. 3 and Table 1 it appears that the open-shell error, which we found to be significant for a small 4×4 lattice, becomes of the order of the statistical error (and thus negligible [47]) for our large lattices [48].

3.3.2. *Accuracy*

At all densities our three sets of data for $U/t = 10$ (Fig. 3), 20 (Fig. 4), and 40 (Fig. 5) are bracketed, as seen in the cumulative Fig. 6, by the noninteracting unpolarized energy and the fully spin-polarized energy (both dashed in Fig. 6), and display a smooth and reasonable behavior. To evaluate the absolute accuracy of our results, we can rely on two exact limits: the low-density ($n \simeq 0$) regime, where we expect $\mathcal{E} = -4n$, and the half filled case ($n=1$), for which the strong-coupling expansion provides the correct large-U/t behavior: to leading order in t/U, the model maps onto an Heisenberg model, whose ground-state energy has been evaluated with great accuracy [12, 49]. We can also consider the next correction term $34.6t^4/U^3$ [50]. At low density our results are essentially exact; at half filling our error is small ($\sim 3\%$) for $U/t = 10$ but (as already noticed in Table 1) it tends to grow with U/t: for $U/t = 20$ it's $\sim 9\%$, and for $U/t = 40$ it's $\sim 11\%$. We have made sure (see markers other than dots in Figs. 3,5 and Refs. [46, 48]) that such an energy discrepancy is not due to shape, open-shell, finite-size [47], and boundary condition effects; as far as systematic errors are concerned, we are thus left with the fixed-node approximation: as U/t grows, more flexible trial wavefunctions may be required to obtain more accurate energies [51].

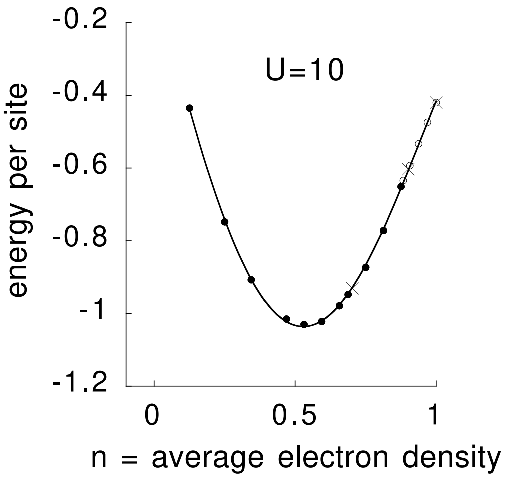

Figure 3. Ground-state energy per site (in units of the hopping parameter t) as a function of the electronic density, for a 2D Hubbard lattice of $N_s = 16 \times 16 = 256$ sites with $U/t = 10$. Errors are smaller than the dot size. Full dots correspond to closed shells and empty dots correspond to open shells. The three crosses correspond to closed-shell densities of a $11\sqrt{2} \times 11\sqrt{2}$ lattice, and are shown for comparison (see text and Refs.[46, 48]).

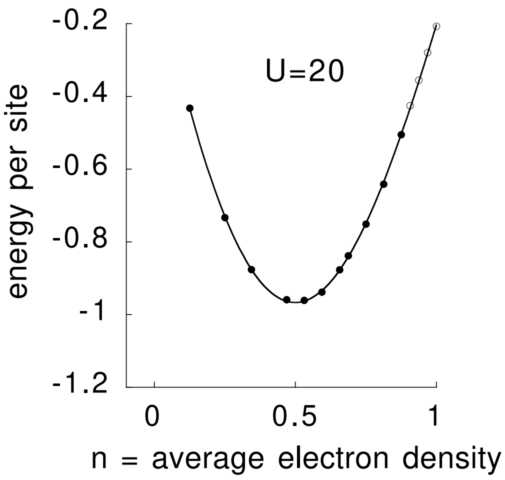

Figure 4. Ground-state energy per site (in units of the hopping parameter t) as a function of the electronic density, for a 2D Hubbard lattice of $N_s = 16 \times 16 = 256$ sites with $U/t = 20$. Errors are smaller than the dot size. Full dots correspond to closed shells and empty dots correspond to open shells.

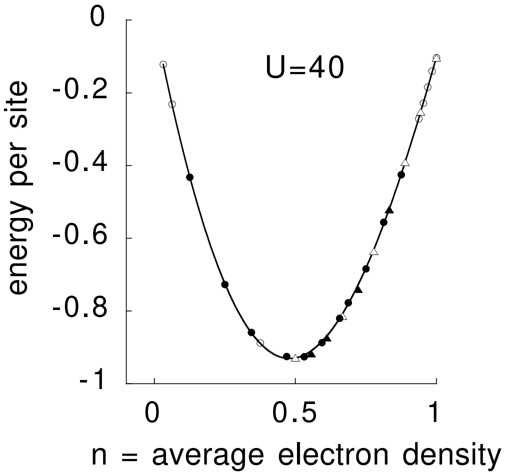

Figure 5. Ground-state energy per site (in units of the hopping parameter t) as a function of the electronic density, for a 2D Hubbard lattice of $N_s = 16 \times 16 = 256$ sites with $U/t = 40$. Errors are smaller than the dot size. Full dots correspond to closed shells and empty dots correspond to open shells. Full and empty triangles, corresponding to closed and open shells of a smaller 12×12 lattice, are shown for comparison, and give a measure of the size effects (see text and Ref. [47]).

3.3.3. *Phase separation*

Keeping in mind the virtues and limitations of our numerical study, we can now turn to the question of phase separation in the Hubbard model. In some sense the function $e(x)$ works like a magnifying lens of the phase separation (see Fig. 1). It should be stressed that in a consistent definition of $e(x)$ the half-filling energy e_H must be obtained as $e_h(x=0)$ from the same calculation as any other $e_h(x \neq 0)$ (here, from the FNMC at half filling). If that's not the case (for example, if the Heisenberg value is used instead), then $e(x)$ will spuriously diverge near $x=0$, with a good chance of artificially creating, rather than magnifying, the occurrence of phase separation. In Figs. 7-9 we find plots of $e(x)$ for $U/t = 10$, 20, and 40; these values, as well as the associated error bars, are directly obtained from those of Figs. 3-5 (i.e. from the original FNMC energies and tiny error bars). Despite the error bars, a common trend is evident for all the calculated coupling strenghts: $e(x)$ has a positive slope for large hole densities, far from half-filling, but then it clearly changes slope below some small critical density x_c. Such a minimum in $e(x)$ implies that, at least for the FNMC effective hamiltonian determined by our choice of wavefunction and cell size [47], phase separation occurs below $x = x_c$. Although a finer grid of

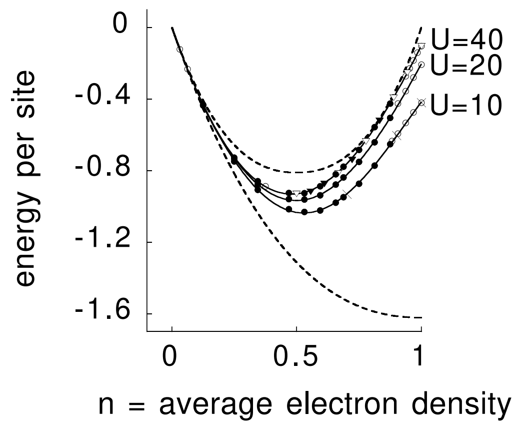

Figure 6. This picture summarizes the results of the previous Figs. 3-5, showing the ground-state energy per site (in units of the hopping parameter t) as a function of the electronic density, for a 2D Hubbard lattice of $N_s = 16 \times 16 = 256$ sites with $U/t = 10$ (lower), 20 (middle), and 40 (upper data). Errors are smaller than the dot size and symbols are the same as in Figs. 3-5. The dashed curves correspond to two known results for the infinite lattice at $U = 0$: the fully spin-polarized case (upper curve), whose total energy per site is symmetric with respect to quarter filling, and the unpolarized case (lower curve), whose total energy per site is symmetric with respect to half filling. (see text).

hole densities would be required to locate with high precision the critical density x_c as a function of U/t, we already see that x_c decreases as U/t is increased; this qualitatively agrees with the original predictions [33] and with some previous calculations on the $t-J$ model at corresponding values of $J = 4t^2/U$ [38].

4. Summary, open problems, conclusions and perspectives

4.1. SUMMARY

In summary, the extensive fixed-node Monte Carlo simulations of the Hubbard model for 16×16 two-dimensional lattices by Cosentini *et al.* [11], just recalled in the previous Section, suggest the presence of a phase separation for $U \gg t$. If confirmed, this result would imply that the $t-J$ model is also likely to show a phase separation in the physically relevant regime $J < 0.4$, and that even a single-band Hubbard model is sufficient to reproduce this physical tendency of some high-T_c superconductors (Subs. 3.1). Before these statements are confirmed beyond any conceivable doubt, further tests may be required because of the following two problems.

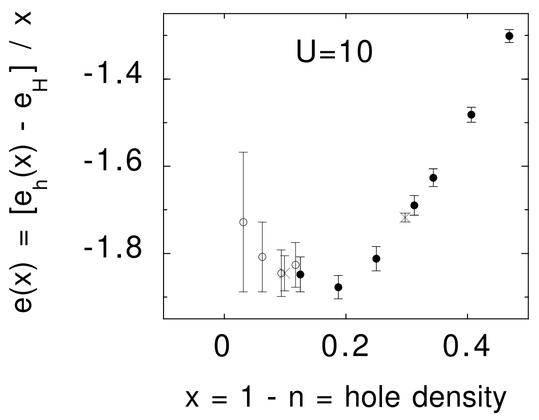

Figure 7. Plot of $e(x)$ vs. x for $U/t = 10$. Full dots correspond to closed shells and empty dots correspond to open shells of a 16×16 lattice. Crosses (corresponding to a $11\sqrt{2}\times 11\sqrt{2}$ lattice) are shown for comparison. Obviously at small x the error bar associated to $e(x)$, $\Delta e(x) = [\Delta e_h(x) + \Delta e_h(x=0)]/x$, becomes significant even if the statistical FNMC error $\Delta e_h(x)$ is tiny.

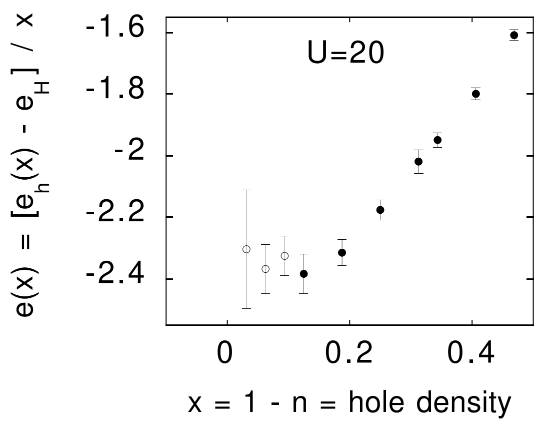

Figure 8. Plot of $e(x)$ vs. x for $U/t = 20$. Full dots correspond to closed shells and empty dots correspond to open shells of a 16×16 lattice. Other comments like in the previous figure.

4.2. OPEN PROBLEMS

4.2.1. *Fixed-node error at large U/t*

First of all, the fixed-node approximation is a variational approximation. Especially for the two larger U/t values, the energy discrepancy at half

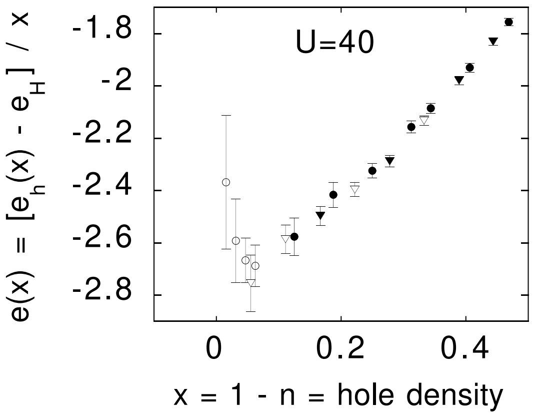

Figure 9. Plot of $e(x)$ vs. x (see text) for $U/t = 40$. Full dots correspond to closed shells and empty dots correspond to open shells of a 16×16 lattice. Full and open triangles correspond to closed and open shells of a smaller 12×12 lattice and are shown for comparison (see text). Other comments like in the previous figure.

filling (as discussed in Subs. 3.3.2) is unfortunately significant. Not knowing in advance the density dependence of such an energy discrepancy, even the conclusions on phase separation, based on the function $e(x)$, could be at risk. Richer variational wavefunctions (yielding different nodal topologies) [51], or possibly the stochastic reconfiguration scheme recently proposed by Sorella [53], must be employed to go beyond our simple Gutzwiller-Slater nodes, and to settle this point completely.

4.2.2. *Finite-size and "phase separation" at $U=0$*

A second motivation for further tests is illustrated in Fig. 10. Here we show the energy density of a finite 16×16 lattice with $U/t = 0$ (noninteracting electrons). Close to half filling, below $x_c = 0.125$, the function $e(x)$ becomes completely flat. Such a behavior, already pointed out by Lin a few years ago [54] but surprisingly forgotten in most subsequent numerical studies of small Hubbard lattices, is an artifact related to the high degeneracy of the Fermi level, and disappears in a truly infinite lattice (x_c vanishes as $2/L$ when L, the side of a square $L\times L$ lattice, goes to infinity). Comparing Fig. 10 with previous Figs. 7-9, one may even fear that the artificial flattening at $U = 0$ has something to do with the behavior found for $e(x)$ at $U > 0$. This observation suggests caution even if, as stated in various parts of Subs. 3.3, many severe tests of open-shell and finite-size effects have already been passed. The almost perfect convergence of energies $\mathcal{E}(n)$, shown in Figs. 3-

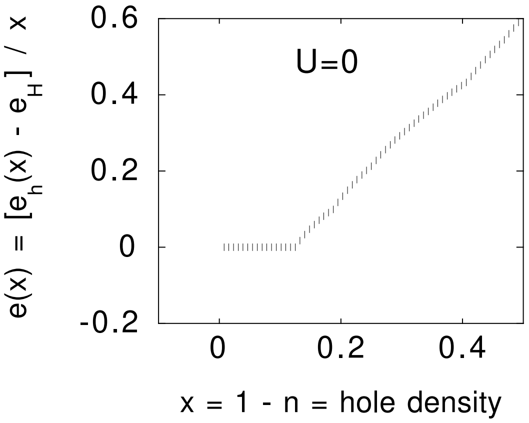

Figure 10. This figure contains plot of $e(x)$ vs. x for $U/t = 0$ (noninteracting electrons, analytic results) for a 16×16 lattice. As a marker we use here vertical bars; we display data for all the 128 densities which can be obtained at this lattice size with an even number of electrons. Note that below $x = 2/16$ $e(x)$ is completely flat.

5, and even the convergence (within the appropriate statistical error) of the $e(x)$ data of Figs. 7-9, could, incredibly, not be enough to rule out a very slow $2/L$ drift of the critical density towards zero. Perturbation theory or further numerical work at small U/t, as well as a thorough finite-size scaling study based on yet larger lattices may be required to convincingly rule out the possibility of such a spurious effect.

4.3. CONCLUSIONS AND PERSPECTIVES

Neither previous, quite crude numerical studies, nor the recent (and definitely much more accurate) work by Cosentini *et al.*, recalled in this chapter, seem to be sufficient for a 100% safe and general claim on phase separation. This is by itself information of the greatest importance to whoever is active either in the theory or in the numerical simulation of strongly correlated two-dimensional electron systems: phase separation is a very delicate effect, which can be masked by otherwise negligible variational or finite-size errors. The simplicity and ease of the lattice fixed-node Monte Carlo [24, 13], the efficiency of the new Calandra-Sorella reconfiguration scheme [12], and the experience gained in the preliminary study recalled here [11] provide a good starting point and very useful guidelines for a future, conclusive numerical study of this challenging problem.

5. Acknowledgements

Thanks are due to M. Capone and L. Guidoni for many stimulating discussions during the preparation of both the lectures and the manuscript, and to M. Calandra Buonaura, C. Lavalle and G. Senatore for patiently listening to GBB's rehearsals in Ithaca, and for their useful comments . Partial support from the Italian National Research Council (CNR, Comitato Scienza e Tecnologia dell'Informazione, grants no. 96.02045.CT12 and 97.05081.CT12), the Italian Ministry for University, Research and Technology (MURST grant no. 9702265437) and INFM Commissione Calcolo is gratefully acknowledged.

References

1. Fahy, S. (1998) – contribution to this book.
2. Mitáš, L. (1998) – contribution to this book.
3. Umrigar, C.J. (1998) – contribution to this book.
4. Koch, E., Gunnarsson, O., and Martin, R.M., (1998) – contribution to this book; the "downfolding" appropriate to C_{60} is discussed at the very beginning.
5. In this respect it may be worth mentioning that, although model hamiltonians can be intrinsically exciting to some theorists [6], their link to real bands and bonds, following a pioneering paper by Andersen and Jepsen [7], can in fact be reasonably worked out, at least as far as the one-electron matrix elements are concerned [4, 8].
6. The work presented in the second part of this chapter belongs in some sense to this line of thought: it's based on the simplest 2D Hubbard hamiltonian (just one orbital per site) and unrealistically large values of the electron-electron repulsion. However, as discussed in Sect. 3, the state of the controversy on phase separation is such that even an oversimplified model like this might give a significant contribution, and possibly shed some light even on real systems.
7. Andersen, O.K., and Jepsen, O. (1984) Explicit, first-principles tight-binding theory *Phys. Rev. Lett.* **53**, 2571.
8. Andersen, O.K., *et al.* (1998) Third-generation TB-LMTO in: *Tight-Binding Approach to Computational Materials Science, MRS Symposia Proceedings No. 491* (MRS, Pittsburgh) and *cond-mat/9804166*; the "downfolding" appropriate to some high-Tc superconductors is in one of its sections.
9. Hood, Randy Q., *et al.* (1997) Quantum Monte Carlo Investigation of Exchange and Correlation in Silicon *Phys. Rev. Lett.* **78**, 3350.
10. Malatesta, A., Fahy, S.B., and Bachelet, G.B. (1997) Variational quantum Monte Carlo calculation of the cohesive properties of cubic boron nitride *Phys. Rev. B* **56**, 12201.
11. Cosentini, A.C., Capone, M., Guidoni, L., and Bachelet, G.B. (1998) Phase separation in the 2D Hubbard model: a fixed-node quantum Monte Carlo study *cond-mat/9801299* and submitted to *Phys. Rev. B.*
12. Calandra Buonaura, M., and Sorella, S. (1998) Numerical study of the two-dimensional Heisenberg model using a Green function Monte Carlo technique with a fixed number of walkers *Phys. Rev. B* **57**, 11446; Hetherington, J.H. (1984) Observations on the statistical iteration of matrices *Phys. Rev. A* **30**, 2713.
13. ten Haaf, D.F.B., van Bemmel, H.J.M., van Leeuwen, J.M.J., van Saarloos, W. and Ceperley, D.M. (1995) Proof for an upper bound in fixed-node Monte Carlo for lattice fermions *Phys. Rev. B* **51**, 13039.
14. Fisher, Matthew P.A. (1998) Mott Insulators, Spin Liquids and Quantum Disor-

dered Superconductivity, *cond-mat/9806164*. The link between real solids and popular lattice models of correlated electrons is discussed in the first part.

15. A common option is $\mathcal{G} = 1 - \tau(\mathcal{H} - E_{ref})$, with a convenient choice of a small positive τ and of the energy reference E_{ref}. This operator looks like a short-time Green's function, but no short-time approximation is involved here: on a lattice this is just a way of constructing an iterative ground-state projector, because the spectrum of the hamiltonian is bounded both from above and from below.
16. Under the conditions of the previous Ref. [15], analogous projectors can be constructed for the highest excited state by e.g. reversing the sign of $\mathcal{H}$; in this chapter, however, we only deal with ground-state projectors.
17. Nightingale, M.P. – contribution to this book.
18. For example, in the Hubbard model of this chapter, the number of basis states $|X\rangle$ connected to a given state $|X'\rangle$ - and thus the number of operations needed to evaluate, say, $\langle X|\mathcal{H}|\Psi\rangle$ - grows linearly with the number of electrons N_e; which in turn, because of the Pauli principle, can be at most twice the number of sites N_s. This is fortunately much less than the size of the Hilbert space, which, for N_e identical particles on a lattice of N_s sites and an unspecified total spin S and projection S_z, is $(2N_s)!/(2N_s - N_e)!$ for "labeled fermion" configurations (see e.g. the appendix B of Ref. [13]), and $\binom{2N_s}{N_e}$ for (antisymmetric) true fermion states. In both cases, the overwhelming majority of hamiltonian matrix elements are zero (sparse matrix).
19. For example, one could estimate the largest eigenvalue of $\mathcal{G}$ as $\gamma = \sum_i w_i(L) b_i / \sum_i w_i(L)$; then, if the relation between $\mathcal{G}$ and $\mathcal{H}$ follows the choice of Ref. [15], the estimate for the ground-state energy would be $E_\circ = [(1-\gamma)/\tau] + E_{ref}$.
20. See e.g. Trivedi, N., and Ceperley, D.M. (1990) Ground-state correlations of quantum antiferromagnets: a Green-function Monte Carlo study *Phys. Rev. B* **41**, 4552.
21. Hamann, D.R., Schlüter, M., and Chiang, C. (1979) Norm-conserving pseudopotentials *Phys. Rev. Lett.* **43**, 1494; Bachelet, G.B., Hamann, D.R., and Schlüter, M. (1982) Pseudopotentials that work: from hydrogen to plutonium *Phys. Rev. B* **26**, 4199.
22. Boninsegni, M. (1995) Ground state of a triangular quantum antiferromagnet: Fixed-node Green-function Monte Carlo study *Phys. Rev. B* **52**, 15304; (1996) Ground state of a frustrated quantum antiferromagnet: fixed-node Green Function Monte Carlo study *Phys. Lett. A* **216**, 313.
23. The following digression could be useful to some reader. The Hubbard hamiltonian is usually written in terms of creation and annihilation operators, according to the formalism of second quantization (see e.g. Eq. 3); in this context each of the basis states $|X\rangle$ of the many-body Hilbert space is identified by an array of occupation numbers, which in turn corresponds to a Slater determinant of occupied single-particle spin-orbitals. Each individual basis state $|X\rangle$, and thus any linear combination of these basis states, is therefore an automatically antisymmetric fermion function. With this choice the "sign problem" is built into the basis set, and directly emerges, without the need of any fermionic trial function, from the off-diagonal matrix elements: because of the fermionic commutation rules of the creation and annihilation operators, some $\langle X'|\mathcal{H}|X\rangle$ turn out to be equal to $-t$, while others are equal to $+t$. If we instead choose a different (and larger [18]) basis $|X\rangle$ of "labeled fermion" configurations, as explained e.g. in appendix B of Ref. [13], then the sign problem separates into "good" off-diagonal hopping terms (all $\langle X'|\mathcal{H}|X\rangle$ are now negative and equal to $-t$) and a "bad" antisymmetry requirement, which is imposed through a fermionic trial function $\Psi_T(X)$, and causes sign flips in the importance-sampled Green's function matrix $\tilde{\mathcal{G}}_{X'X} = \Psi_T(X')\langle X'|\mathcal{G}|X\rangle\Psi_T^{-1}(X)$. The latter choice corresponds to an intuitive picture which parallels the continuum case and may be convenient for practical purposes, but it's inconsistent with the usual second quantization.
24. van Bemmel, H.J.M., ten Haaf, D.F.B., van Saarloos, W., van Leeuwen, J.M.J., and An, G. (1994) Fixed-node quantum Monte Carlo method for lattice fermions *Phys.*

Rev. Lett. **72**, 1442.
25. Anderson, J.B. (1975) A random-walk simulation of the Schrödinger equation: H_3^+ *J. Chem. Phys.* **63**, 1499.
26. Ceperley, D.M., and Alder, B.J. (1981) The Calculation of the Properties of Metallic Hydrogen using Monte Carlo *Physica B* **108**, 875.
27. It also gives a sound proof of the variational principle and withdraws an incorrect statement of Ref. [24] about "lever rules".
28. Gunnarsson, O., Koch, E., and Martin, R.M. (1996) Mott transition in degenerate Hubbard models: Application to doped fullerenes *Phys. Rev. B* **54**, R11026; (1997) Mott-Hubbard insulators for systems with orbital degeneracy *ibid.* **56**, 1146; Aryasetiawan, F., Gunnarsson, O., Koch, E., and Martin, R.M. (1997) Pauli susceptibility of A3C60 (A = K,Rb) *Phys. Rev. B* **55**, 10165; also see Ref. [4], and references therein.
29. Grilli, M., Raimondi, R., Castellani, C., Di Castro, C., and Kotliar, G. (1991) Superconductivity, phase separation, and charge-transfer instability in the U=∞ limit of the three-band model of the CuO2 planes *Phys. Rev. Lett.* **67**, 259; Di Castro, C., and Grilli, M. (1992) *Physica Scripta* T**45**, 81, and references therein.
30. Emery, V.J., and Kivelson, S.A. (1993) Frustrated electronic phase separation and high-temperature superconductors *Physica C* **209**, 597.
31. Dagotto, E., Moreo, A., Ortolani, F., Poilblanc, D., and Riera, J. (1992) Static and dynamical properties of doped Hubbard clusters *Phys. Rev. B* **45**, 10741.
32. Castellani, C., Di Castro, C., and Grilli, M. (1995) Singular quasiparticle scattering in the proximity of charge instabilities *Phys. Rev. Lett.* **75**, 4560.
33. Emery, V.J., Kivelson, S.A., and Lin, H.Q. (1990) Phase separation in the $t-J$ model *Phys. Rev. Lett.* **64**, 475.
34. Jorgensen, J.D.*et al.* (1988) Superconducting phase of La_2CuO_{4-y}: Superconducting composition resulting from phase *Phys. Rev. B* **38**, 11337.
35. Chou, F.C. *et al.* (1996) Phase separation and oxygen diffusion in electrochemically oxidized $La_2CuO_{4+\delta}$: A static magnetic susceptibility study *Phys. Rev. B* **54**, 572.
36. Tranquada, J.M., Sternleb, B.J., Axe, J.D., Nakamura, Y., and Uchida, S. (1995) Evidence for stripe correlations of spins and holes in copper oxide superconductors *Nature* **375**, 561.
37. Bianconi, A. *et al.* (1996) Determination of the local lattice distortions in the CuO_2 plane of $La_{1.85}Sr_{0.15}CuO_4$ *Phys. Rev. Lett.* **76**, 3412.
38. Hellberg, C.S. , and Manousakis, E. (1997) Phase Separation at all Interaction Strengths in the $t-J$ Model *Phys. Rev. Lett.* **78**, 4609.
39. Shih, C.T., Chen, Y.C., and Lee, T.K. (1998) Phase separation of the two-dimensional $t-J$ model *Phys. Rev. B* **57**, 627.
40. Gang Su (1996) Phase separation in the two-dimensional Hubbard model *Phys. Rev. B* **54**, R8281.
41. This choice corresponds to a projection of the total spin $S_z = 0$ and, for those electron densities which correspond to a closed-shell situation, also to a total spin $S = 0$.
42. Koch, E. (1998), *private communication.*
43. Parola, A., Sorella, S., Baroni, S., Car, R., Parrinello, M., and Tosatti, E. (1989) Recent numerical results on the two dimensional Hubbard model *Physica C* **162-164**, 771; Parola, A., Sorella, S., Parrinello, M., and Tosatti, E. (1991) d-wave, dimer, and chiral states in the two-dimensional Hubbard model *Phys. Rev. B* **43**, 6190.
44. Shiwei Zhang, Carlson, J., and Gubernatis, J.E. (1995) Constrained path quantum Monte Carlo method for fermion ground states *Phys. Rev. Lett.* **74**, 3652.
45. See also Shiwei Zhang (1998) – contribution to this book.
46. Different boundary conditions yield energies within an error bar for 16×16 lattices, and within at most two error bars for our smallest lattices; all results presented here refer to antiperiodic boundary conditions.
47. We will come back to this point while discussing our results in Subs. 4.2.2.
48. We have checked a few densities which correspond to open shells in our usual 16×16

lattice (empty markers in Figs. 3-9), but to closed shells in a 45^o-rotated $11\sqrt{2}\times 11\sqrt{2}$ lattice (crosses in Figs. 3,6 and 7); the energy difference turns out to be negligible [47].

49. Zhong, Q.F., and Sorella, S. (1993) Spin-wave theory on finite lattices: application to the J1-J2 Heisenberg Model *Europhys. Lett.* **21**, 629.
50. This term has been proposed long ago by Takahashi (1977) Half-filled Hubbard model at low temperature *J. Phys. C* **10**, 1289; the resulting large-U expansion agrees remarkably well, for $U/t \geq 10$, with the more recent results of Polatsek, G., and Becker, K.W. (1996) Ground-state energy of the Hubbard model at half filling *Phys. Rev. B* **54**, 1637.
51. The Gutzwiller-Slater wavefunction, a very popular choice for the Hubbard model, suffers from at least two drawbacks (which probably reflect into the "quality" of its nodes): for open shells it does not describe a spin singlet, and it only has on-site correlations; more general forms, such as those proposed in Ref. [52], could gain some of the missing energy at $U/t = 20, 40$.
52. Giamarchi, T., and Lhuillier, C. (1990) Variational Monte Carlo study of incommensurate antiferromagnetic phases in the two-dimensional Hubbard model *Phys. Rev. B* **42**, 10641.
53. Sorella, S. (1998) Green Function Monte Carlo with Stochastic Reconfiguration. *Phys. Rev. Lett.* **80**, 4558.
54. Lin, H.Q. (1991) Ground-state properties of the two-dimensional Hubbard model *Phys. Rev. B* **44**, 7151.