Disordered systems / Systèmes désordonnés

A review of Monte Carlo simulations for the Bose–Hubbard model with diagonal disorder

Une revue des simulations Monte Carlo pour le modèle de Bose–Hubbard avec désordre diagonal

Lode Pollet

Department of Physics, Arnold Sommerfeld Center for Theoretical Physics and Center for NanoScience, University of Munich, Theresienstraße 37, 80333 Munich, Germany

A R T I C L E   I N F O

A B S T R A C T

We review the physics of the Bose–Hubbard model with disorder in the chemical potential focusing on recently published analytical arguments in combination with quantum Monte Carlo simulations. Apart from the superfluid and Mott insulator phases that can occur in this system without disorder, disorder allows for an additional phase, called the Bose glass phase. The topology of the phase diagram is subject to strong theorems proving that the Bose Glass phase must intervene between the superfluid and the Mott insulator and implying a Griffiths transition between the Mott insulator and the Bose glass. The full phase diagrams in 3d and 2d are discussed, and we zoom in on the insensitivity of the transition line between the superfluid and the Bose glass in the close vicinity of the tip of the Mott insulator lobe. We briefly comment on the established and remaining questions in the 1d case, and give a short overview of numerical work on related models.

R É S U M É

Nous passons en revue la physique du modèle de Bose–Hubbard en présence de désordre dans le potentiel chimique, en portant notre attention sur des arguments analytiques récemment publiés en les combinant avec des simulations Monte Carlo quantiques. Mis à part les cas du superfluide et des phases isolantes de Mott qui peuvent advenir dans ce système sans désordre, le désordre conduit à une phase additionnelle, appelée verre de Bose. La topologie du diagramme de phases est contrainte par un théorème prouvant que la phase du verre de Bose doit intervenir entre le superfluide et l'isolant de Mott, impliquant une transition de Griffiths entre ce dernier et le verre de Bose. Les diagrammes de phase complets en 3d et 2d sont discutés, et nous insistons sur l'insensibilité de la ligne de transition entre le superfluide et le verre de Bose au voisinage proche de l'extrémité du lobe de l'isolant de Mott. Nous commentons brièvement les connaissances établies et les questions qui demeurent dans le cas 1d, et présentons un rapide survol des travaux numériques sur les modèles liés.

E-mail address: lode.pollet@physik.uni-muenchen.de.

1. Introduction

The interplay between disorder and interactions is a subtle and long-standing problem in condensed matter physics, especially for bosons on a lattice. In the absence of interactions, the smallest amount of disorder will lead to Anderson localization of the atoms when all bosons occupy the same lowest-energy localized state. This limit is clearly pathological. In the absence of disorder, a sufficiently strong interaction will drive the system towards a gaped and incompressible Mott insulator provided the density is commensurate. Although both disorder and interaction tend to localize the bosons, the superfluid phase is quite robust against disorder. This compressible and gapless superfluid is found over a large part in the phase diagram: it exists for weak interactions, but also for interactions where the pure system is an insulator, so that disorder-induced superfluidity is possible. For even stronger disorder and/or interactions, it will ultimately go over in a new phase, the Bose glass phase [1,2]. This is a gapless and compressible but insulating phase. Its properties thus defy intuitive notions about conductivity based on Fermi liquid theory. Fisher et al., building on the one-dimensional work by Giamarchi and Schulz [1,2], argued the existence of the Bose glass phase in any dimension and gave an extensive description of its properties [3].

For a long time, there was controversy whether a direct transition between a superfluid (SF) and a Mott insulator (MI) in the presence of disorder was possible [4–25]. Fisher et al. argued that this was unlikely, though not fundamentally impossible [3]. Curiously, a large number of direct numerical simulations and some approximate approaches observed the unlikely scenario. In the past five years, a number of theorems have been published that allow us to answer these issues and prove the conjectures by Fisher et al. [3]. Here, we review these theorems, and discuss their influence on the numerically obtained phase diagrams of the disordered Bose–Hubbard model with disorder in the chemical potential by large-scale path–integral Monte Carlo simulations with worm-type updates.

The main results and the contents of the paper are as follows. To set the ideas, we specify our Hamiltonian and the disorder model in Section 2 with the definition of the phases of interest (superfluid (SF), Mott insulator (MI), and Bose glass (BG)). We proceed in Section 3 with a brief overview of what is known analytically about such systems. We will introduce Lifshitz regions, analyze the regime of weak interactions, review a sufficient condition by Fisher et al. for the presence of a BG phase in the vicinity of the MI–BG transition, review the theorem of inclusions stating that a glassy phase must intervene between a gaped and gapless phase in the presence of disorder, which in turn shows that the sufficient condition by Fisher is also necessary and leads to a Griffiths-type transition between the BG and the MI. We then move onto the numerical section (see Section 4) containing a brief discussion of path–integral Monte Carlo with worm-type updates, and a presentation of the phase diagrams of the disordered Bose–Hubbard Hamiltonian in $d = 3, 2$ and 1 dimensions. We briefly mention a number of experimental (see Section 5) and other recent numerical studies in Section 6, before concluding in Section 7.

2. Bose–Hubbard model

We consider the Bose–Hubbard model [3] describing scalar bosons on a lattice,

$$H = -t \sum_{\langle i,j \rangle} b_i^\dagger b_j + \frac{U}{2} \sum_i n_i(n_i - 1) - \mu \sum_i n_i \tag{1}$$

The operators b_j annihilate a scalar boson on site j while $b_i^\dagger$ creates such a boson on site i, and is the Hermitean conjugate operator of b_i. These operators satisfy the bosonic commutation relations, $[b_i, b_j^\dagger] = \delta_{i,j}$, and this yields zero for commutators with two creation or two annihilation operators. The operator $n_i = b_i^\dagger b_i$ counts the number of bosons on site i. The first term describes the hopping of bosons between neighboring sites with tunneling amplitude t, which we set as our unit $t = 1$. The second term describes the on-site repulsion with strength U, while the last term is proportional to the chemical potential μ. The lattice spacing is set to unity. Unless otherwise specified, we have a lattice in mind of linear size L and volume L^d where d is the dimension, and we use periodic boundary conditions. The Bose–Hubbard model is the simplest model that describes a conductor–insulator transition for bosons. It has three phases, which can easily be identified in the limiting cases [3,26]: first, in the limit of high temperature, the system is a normal liquid. At zero temperature and $t = 0$, the system is a Mott insulator with fixed, integer density, a gap, and zero compressibility. For finite hopping, stable Mott lobes around the $t = 0$ insulators are found, which are surrounded by a gapless, compressible superfluid. The superfluid phase also exists at finite temperature. The zero-temperature phase diagram in the mean-field approximation (also known as the decoupling approximation) is shown in Fig. 1. The true phase diagrams in 3d and 2d look very similar with only quantitative changes, especially in the vicinity of the tip of the Mott lobe, which are shifted to lower interaction strengths by about 20% and 30%, respectively. In 1d, the tip of the lobe has a cusp reflecting the Kosterlitz–Thouless transition where the shape of the Mott insulator also bends down. There is thus reentrant behavior in 1d when following a line of constant μ close to the tip of the lobe.

In this review, disorder will be introduced at the level of the chemical potential by $\mu \to \mu_i = \mu - \epsilon_i$, where ϵ_i is the disordered on-site potential. This is known as diagonal disorder, unlike disorder introduced in the hopping, which is known as off-diagonal disorder. We can take ϵ_j as independent random variables distributed according to the probability density
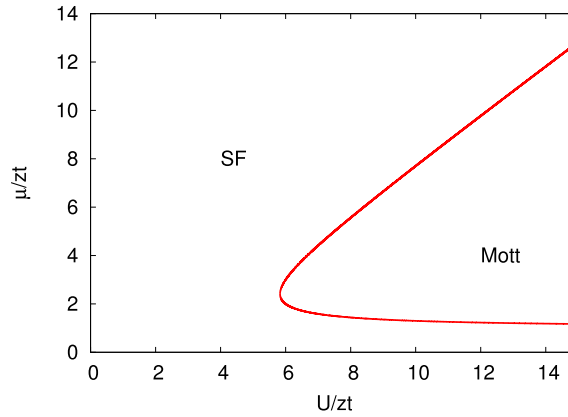
Fig. 1. (Color online.) Mean-field zero temperature phase diagram in the (U, μ) plane in units of zt (where z is the coordination number) of the Bose–Hubbard model without disorder. At commensurate densities, Mott insulators (MI) are formed for strong-enough interactions, while a superfluid (SF) is found elsewhere. In this approximation, the tip of the Mott lobe is found for $U/zt \approx 5.83$, where the Mott boundary has a vertical slope [3,26].

$p(\epsilon/\Delta)$, which satisfies the normalization condition $\int_{-1}^{1} \mathrm{d}u\, p(u) = 1$. The first moment can always be taken to be zero, $\int_{-1}^{1} \mathrm{d}u\, up(u) = 0$, by absorbing it in the definition of μ. The disorder distribution is taken to be bounded (otherwise the Mott insulator does not exist), that is $p(u) = 0$ if $|u| > 1$. Formally, the disorder bound Δ and the disorder distribution dispersion $\delta = \sqrt{\langle \epsilon^2 \rangle - \langle \epsilon \rangle^2}$ are independent parameters. For the most common choice of the uniform distribution $p(u) = $ const (also used in our numerical simulations), we have $\delta^2 = \Delta^2/3$. For the theoretical discussion, however, we allow more general distributions. The disorder can completely be characterized by specifying all its moments. One may also add parameters that control correlations between the potentials on different lattice sites. All these parameters are denoted by ξ, and we identify them with the definition of a particular model of disorder.

With this kind of disorder, a new phase is possible: the Bose glass phase [1,2], which is an insulator, though gapless and compressible. While the superfluid is self-averaging, the Bose glass phase is not; that is, most realizations of the superfluid have the same superfluid stiffness in the thermodynamic limit, while for a large though finite system size in the Bose glass phase, this property will be absent. With other types of disorder, other phases such as a Mott glass [27,19] may exist (the Mott glass is incompressible, and can occur for instance for hard-core bosons at half-filling with disorder in the hopping, featuring particle–hole symmetry), but we will not consider such cases in this paper.

3. Analytical aspects

In this section we review some old and some recent theoretical arguments that put very strong constraints on the phase boundaries of the disordered Bose–Hubbard model. We start with a standard textbook discussion of Lifshitz tails.

3.1. Lifshitz tails

In disordered systems, localized states with an energy close to the band edge $E_{\min} = -\Delta - 2dt$ play a crucial role. The probability P of finding a state with energy $\delta E = E - E_{\min}$ close to $E_{\min}$ extended over ℓ states is proportional to $(1/\Delta)^{\ell}$. The linear size of such a state is $R \sim \ell^{1/d} r_0$, with r_0 a typical length scale such as the lattice constant. Its energy is then $\delta E = t\ell^{-2/d} = t(R/r_0)^2$. From this, it follows that the probability P is given by:

$$P \sim e^{-(t/\delta E)^{d/2} c} \tag{2}$$

with c a constant. Such Lifshitz regions are thus exponentially rare, but can be arbitrarily large (such that the quantization energy can be very low). They pose a challenge for theoretical treatments of disordered systems and will show up several times in the discussions below.

3.2. Weak interactions

This regime was analyzed in detail by Falco et al. [28] and Lugan et al. [29], and we reproduce here the dependence of the critical line on U/t and Δ/t by a derivation based on the central limit theorem. When $U \ll 1$ and $\Delta \ll 1$, the transition between the BG and the SF happens through percolation between localized states (with large localization length) at high energy E. The precise form of the disorder distribution is then irrelevant, and we can use a delta-correlated, Gaussian distribution V, with mean 0 and variance Δ^2 satisfying $\langle V(x)V(y)\rangle = \Delta^2 a^d \delta(x - y)$ with a the lattice constant. The requirement of commensurability of the density n is likewise unimportant since Mott physics does not enter into the picture. The density
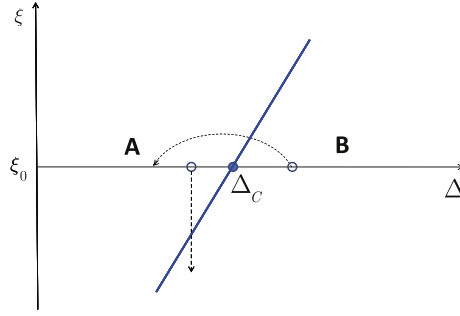
Fig. 2. (Color online.) A sketch of the generic phase transition line between some phases A and B in the plane of disorder distribution parameters Δ and ξ, where Δ is the bound and ξ is one of the infinite number of parameters characterizing the disorder distribution function, *e.g.*, the variance and its spatial correlations. Dashed lines with arrows originate from points which determine the disorder properties in the macroscopic (thermodynamic limit) system, and end at points which characterize disorder parameters in an arbitrarily large, but finite, domain as a result of rare statistical fluctuations in the same system. Reprinted figure with permission from Ref. [31]. Copyright (2009) by the American Physical Society.

of (localized) states with a high energy E is the probability that the average potential energy over a large region of linear size R is E. This leads to the following estimate for the density of states, using the central limit theorem,

$$\rho(E) = \left\langle \delta\left(\int \frac{d^d x}{R^d} V(x) - E \right) \right\rangle = \int d\phi \, e^{-i\phi E} \left\langle e^{i\phi \int \frac{d^d x}{R^d} V(x)} \right\rangle$$
$$= \int d\phi \, e^{-i\phi E} e^{-\frac{\phi^2}{2} \iint \frac{d^d x}{R^d} \frac{d^d y}{R^d} \langle V(x) V(y) \rangle} \propto e^{-\frac{E^2 R^d}{2\Delta^2 r_0^d}} \tag{3}$$

Since $E \propto t \frac{r_0^2}{R^2}$ it follows that $\rho(E) \propto e^{-\sqrt{\frac{E^{4-d} t^d}{2\Delta^4}}}$. The chemical potential parameter can be estimated as nU in this regime, setting the scale for the critical E when the localized states start overlapping,

$$\Delta_c \propto U^{1-d/4} \tag{4}$$

3.3. A necessary condition for a gaped system

Fisher et al. proved rigorously that if the bound Δ on the disorder strength is larger than the half-width of the energy gap $E_{g/2}$ in the Mott insulator of the disorder free system, then the system is compressible. This implies that the transition, characterized by a vanishing superfluid density, goes from a SF to a BG and not to a MI whenever:

$$\Delta_c > E_{g/2} \tag{5}$$

The proof of this theorem is a reasoning typical for disordered systems and based on Lifshitz regions. Consider an arbitrarily large region (however rare) where the chemical potential is roughly constant and lower than $-E_{g/2}$. Locally, this system looks like a superfluid. Provided that the finite-size quantization of the energy due to the large volume is small enough, such regions can be doped (take for instance two such regions, then particles can be transferred from the one to the other at no cost). So the whole system can have no gap in the spectrum and is compressible.

Although this proves that the system is certainly not an insulator whenever $\Delta > E_{g/2}$, it is not a necessary and sufficient condition fixing the location of the MI boundary since the argument does not prove that systems with $\Delta < E_{g/2}$ (and $U > U_c \, (\Delta = 0)$) are gaped. It also does not rule out a direct transition between a SF and the MI. Although the latter was conjectured to be very unlikely by Fisher et al., it was only rigorously ruled out in 1d by Svistunov [13] before the theorem of inclusions (valid in any dimension) was proven, which is the topic of the next section.

3.4. Theorem of inclusions

This and the next section follow the arguments presented in Refs. [30,31]. The theorem of inclusions is explained schematically in Fig. 2. Consider a phase transition driven by disorder between two arbitrary phases A and B. The curved dashed line between A and B means that inside phase B there exist arbitrarily large lakes that locally look like phase A, since B occurs for stronger disorder than A. This amounts to the same argument as in the previous section. How can we prove the converse, namely that inside A there are arbitrarily large regions that locally look like B? For a generic disorder model, the transition between A and B will be a function of the disorder bound Δ as well as of the properties of the disorder distribution specified by ξ (see Section 2). Close to the transition line, there exist then arbitrarily large domains in phase A that look like phase B but for a model with a slightly different disorder distribution (and it is always possible to

find such a distribution since there are an infinite number of continuous parameters determining the disorder properties), as is explained in the figure by the vertical dashed line with an arrowhead.

The consequence of the above argument is that it is not possible for phase A to be gaped if phase B is gapless since the domains of arbitrarily large size that look like phase B inside phase A make phase A gapless. As a consequence, no direct transition between the gaped Mott insulator phase and the gapless superfluid phase is possible, but a gapless Bose glass must always intervene in any dimension. The compressibility of the BG phase follows immediately from the existence of rare but arbitrarily large regions where the energy quantization is low enough such that the chemical potential can be slightly varied away from its equilibrium value (essentially the disorder average over the region) without violating the disorder bound. This fails for a Mott glass, however.

3.5. Griffiths transitions

Another consequence of the theorem of inclusions is that it seems to rule out the transition between the MI and the BG at first sight. The paradox is resolved by considering the exception implied by the rule: the argument of the previous section does not work when the critical point Δ_c cannot be identified with a generic case, i.e., that Δ_c is not dependent on ξ. A transition which depends only on the bound Δ cannot be linked to any local physics because when the variance of the disorder distribution goes to zero, the system becomes indistinguishable from a disorder-free system in the thermodynamic limit. The transition mechanism must hence be based on statistical fluctuations that explore the possibility of reaching the disorder bound at all sites on larger and larger scales which mimic a regular pure system in an external field. For the MI, the external field is the chemical potential which drives the MI to a SF whenever it exceeds the gap for particle or hole creation. In the general case, it can be any regular external field whose amplitude scales with Δ. This mechanism is nothing but the conjectured Griffiths-type MI–BG transition when the vanishing of the gap at the critical point is due to an infinitesimal concentration of rare regions in which the fluctuation of disorder mimics a *homogeneous* chemical potential shift [3]. So, the boundary of the MI is given by $E_{g/2}$, and there is no hope to observe this transition numerically.

4. Quantum Monte Carlo simulations

Before presenting the phase diagrams of the disordered Bose–Hubbard model in $d = 3$, 2 and 1, let us first make a few remarks on how such systems can be studied numerically by path–integral Monte Carlo simulations with worm-type updates. We refer to Ref. [32] for a recent review of the method with applications for cold gases.

4.1. The worm algorithm

The worm algorithm [33] is a path–integral Monte Carlo method (PIMC) ideally suited for studying systems with superfluidity. Let us first see how the algorithm works in the absence of disorder. The starting point is the following decomposition for the partition function,

$$Z = \mathrm{Tr}\, \mathrm{e}^{-\beta H} = \mathrm{Tr}\,\mathcal{T} \mathrm{e}^{-\beta H_0} \exp\left[-\int_0^\beta \mathrm{d}\tau\, H_1(\tau) \right] \tag{6}$$

where $H_1(\tau) = \mathrm{e}^{\tau H_0} H_1 \mathrm{e}^{-\tau H_0}$. As a basis we use the Fock basis of occupation numbers, $|\{n_i\}\rangle$, in which case the above formula refers to a strong coupling expansion. The potential energy and the chemical potential term are diagonal in this basis, and are combined into H_0. The kinetic term (H_1), which is a one-body operator, is not diagonal in this basis but will lead to transitions between Fock states with matrix elements $\langle \ldots n_i - 1, n_j + 1, \ldots | -t b_j^\dagger b_i | \ldots n_i, n_j \ldots \rangle = -t\sqrt{n_i(n_j + 1)}$. The trace is taken over all Fock basis states. The exponential is expanded into a time-ordered product [34–36],

$$Z = \mathrm{Tr}\,\mathcal{T} \mathrm{e}^{-\beta H_0} \left[1 - \int_0^\beta \mathrm{d}\tau\, H_1(\tau) + \int_0^\beta \mathrm{d}\tau_1 \int_{\tau_1}^\beta \mathrm{d}\tau_2\, H_1(\tau_1) H_1(\tau_2) + \cdots \right] \tag{7}$$

The inverse temperature $\beta = 1/T$ is understood as an imaginary time, where the matrix elements of the operators $\exp[-\Delta\tau H_0]$ act as propagators between the different states at $\tau_1, \tau_2, \ldots$, with $0 < \tau_1 < \tau_2 < \cdots < \tau_j < \cdots < \tau_n < \beta$. There is a pictorial representation of this expansion formula, in which each H_1 term is shown as a kink where the occupation numbers change. The legs of the kinks are connected to other kinks by lines whose properties (it can be thickness or multiple lines) represent the occupation number of the propagator connecting the two kinks. At time $\tau = 0$, one can now select a particle and follow its motion in space and imaginary time. One sees that all permutations of the particles can be represented this way, since at time $\tau = \beta$ the propagated line either closes on the same particle or on another particle, which is allowed by indistinguishability of the particles. One calls a worldline the trajectory of a particle propagating in space and imaginary time.

The expansion in path–integral Monte Carlo is always understood over a finite volume and finite maximum imaginary time. There is no singularity caused by a phase transition, which can only be studied in a finite-size scaling analysis. The expansion is written in terms of an entire function, so that there are no non-physical divergencies in the PIMC formulation.

According to statistical mechanics, the expectation value of an observable A is given by $\langle A \rangle = \frac{1}{Z} \operatorname{Tr} A e^{-\beta H}$. In PIMC, a statistical interpretation is given to Eq. (7) by introducing weights w through $Z = \operatorname{Tr}_{|\{n_i\}\rangle} w(|\{n_i\}\rangle)$. We now have to statistically generate all possible configurations according to their weights by generating all possible expansion orders and matrix elements, assign an (unnormalized) weight to each of them, evaluate the observable A in every configuration, and sum all these contributions. For instance, one can perform local updates by inserting pairs of hopping elements, in which a particle hops from a site to its neighbor, and back at a later time.

At high temperatures or deep in the Mott insulating phase, the kinetic energy is small compared to U and/or T, and few perturbation orders are needed in Eq. (7). Such is not the case when the contributions to the free energy coming from the hopping terms are large: there is no reason to expect that the local updates would be efficient. Even worse, they are not ergodic: the low-energy states in a superfluid are given by states with a different winding number. These are paths in which a particle winds around the full linear length of the system before closing on itself again. Configurations with different winding numbers are topologically distinct and cannot be transformed into each other by local updates alone. The winding number W is directly related to the superfluid density through $\rho_s = \frac{\langle W^2 \rangle L^{2-d}}{d\beta}$ [37], with d the dimension of the system.

The worm algorithm has completely solved those ergodicity problems [33]. Instead of working with the partition function Z alone one also works in the Green function sector Z_G,

$$Z_G = \operatorname{Tr} \mathcal{T}\left\{ b_i(\tau_0) b_j^\dagger(\tau) e^{-\beta H} \right\} \tag{8}$$

Pictorially, the operators $b_i(\tau_0)$ $b_j^\dagger(\tau)$ are open ends delimiting a segment of a worldline and are called the worm head and the worm tail. The worms can be on any site and at any time. A correct algorithm allows for the transition between the Green function sector and the partition function sector (where measurements of observables such as the energy and the superfluid stiffness are done), and allows to move the worms around in configuration space. Worms also have the ability to insert and remove hopping elements (kinks). Since the worm operators correspond to open ends on a worldline segment, they have no problem in exploring configurations with different winding numbers. All updates are local in the Green function sector, which means that all acceptance factors can be made of order unity. In phases with off-diagonal long-range order (either true long-range order such as seen in a Bose–Einstein condensate or quasi long-range order with correlation functions decaying as a power law), the worms will preferentially be far away from each other in configuration space, *i.e.*, we are efficient in describing the physics of those phases. In order to distinguish between the superfluid, Bose glass and Mott insulating phases, we need to compute the superfluid stiffness and the compressibility. The latter corresponds to the winding number fluctuations in imaginary time, *i.e.* by the fluctuation–dissipation theorem, the compressibility is $\kappa = \frac{\partial n}{\partial \mu} = \beta(\langle N^2 \rangle - \langle N \rangle^2)$. Because the worm algorithm is formulated in the grand-canonical ensemble, the latter quantity is readily accessible in simulations. There exist other algorithms that build on the same physical idea of worms such as the stochastic series expansion algorithm with directed loops [38,39], directed loops in path–integral representation [40], algorithms with multiple worms present (which may lead to slowdowns) [41,42], as well as canonical formulations [43,44]. Rather than emphasizing minor differences, the important issues are that sectors with different winding numbers must be sampled efficiently (this is the idea behind worms and ensures that autocorrelation times of order unity can be reached everywhere in the phase diagram, at least in theory) as well as an implementation with an efficient data structure (such that the autocorrelation times of order unity can also be reached in practice).

Second, we discuss how the worm algorithm works in the presence of chemical potential disorder. Because such disorder is diagonal in the Fock basis, there are essentially no changes to the above algorithm other than that the chemical potential is now site dependent. One has to compute thermodynamic observables for a fixed disorder realization, and then average those observables over all possible disorder realizations (by the replica trick). So although we see that the algorithm is essentially unaltered, simulations of disordered systems can still be notoriously difficult and contain a number of pitfalls where the inexperienced (and often also the experienced) practitioner needs to pay attention (most of these considerations not only apply to Monte Carlo simulations, but to any numerical approach):

1. There is the risk of a lack of ergodicity and inefficient sampling. For the disordered Bose–Hubbard model, the worm algorithm has effectively solved this problem, and this should no longer be a main issue if one checks that the simulations are well converged.
2. One underestimates the number of realizations that are needed. Especially when rare fluctuations (Lifshitz regions) have an impact, a too small number of realizations can easily lead to wrong conclusions. Cases are known where 10^5 realizations are needed in order to have converging answers; otherwise small drifts can be observed. The Bose glass also does not have the property of self-averaging, meaning that the distribution of the stiffness has to become very broad on accessible system sizes, so an accurate description of a broad distribution understandably requires many disorder samples. But also weak and marginal superfluid phases have broad distributions on mesoscopic system sizes.
3. System sizes may be too small. Griffiths transitions can never be observed on a finite system, and so numerical results should always be checked against the analytical theorems. The compressibility turns out to be a very small number
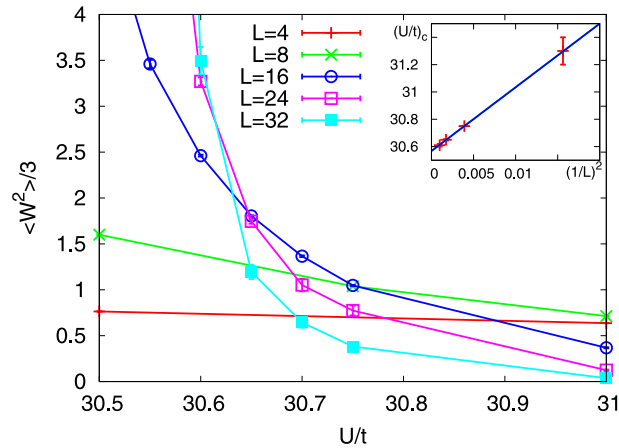
Fig. 3. (Color online.) Finite size scaling analysis for the winding number fluctuations for $\Delta/t = 5$ using the dynamical exponent $z = 2$, chosen for numerical convenience (a calculation with $z = 1$ produced the same critical point), and starting value $\beta t = 1$ for $L = 4$. The inset shows the extrapolation of the intersection points with a $(1/L)^2$ law yielding $(U/t)_c = 30.57(2)$. The figure is a reprint combination of figures from Ref. [30] with permission. Copyright (2009) by the American Physical Society.

(10^{-5} and smaller was observed numerically) in the vicinity of the Mott lobe, also making a direct observation extremely challenging, especially in the vicinity of the tip of the Mott lobe.

4. In cases where the theory is not well established, numerical fluctuations may be hard to control and difficult to interpret. Simulations in one dimension for weak interactions provided a prominent example (see however Section 4.4).

In the next sections we discuss the phase diagrams in 3d, 2d and 1d as found by Monte Carlo simulations.

4.2. Phase diagram of the disordered Bose–Hubbard model in 3d

Critical points on the SF–BG transition line are determined using a finite size scaling analysis of the superfluid density. The scaling equation for the superfluid density for a dimensionless detuning δ reads [3]:

$$\rho_s(\delta, L) = \xi^{-(d+z-2)} f_s(\xi/L, \xi^z/\beta) \tag{9}$$

On finite system sizes, the correlation length $\xi \sim \delta^{-\nu}$ occurring in front of the universal scaling function f_s will be cut off by the system size. For a quantum phase transition, β must be scaled with L^z according to the dynamic exponent z. Hence, by plotting $\rho_s L^{-(d+z-2)}$ versus detuning for different system sizes L and determining the consecutive crossing points, the critical point can be determined. It will be a single crossing point to leading order by scale invariance if the correct scaling exponents are used. An equivalent viewpoint is noting that at the critical point the winding number fluctuations, which are integer numbers, must be scale invariant. For a disorder-free Bose–Hubbard model at the tip of the Mott lobe, $z = 1$. In the presence of disorder and when the compressibility κ is finite, it was argued that $z = d$ in Ref. [3]. This was however recently questioned in numerical [9,45] and analytical [46] studies. How the finite size scaling is done in practice can be seen in Fig. 3. Because $z = 3$ implies a too fast increase of β with L, we opted for a lower $z = 2$ out of convenience. Any $z > 0$ is fine to find the critical point. It is clear from Fig. 3 that the critical point on the SF–BG transition line can be determined reliably, which allows an accurate numerical determination of the entire superfluid boundary.

The resulting phase diagram of the 3d Bose–Hubbard model is shown in Fig. 4 for unit density. The BG phase always intervenes between the SF and MI, while the transition from the MI to the BG phase is of the Griffiths type, in agreement with the preceding theoretical discussion. The transition line is then determined by measuring the gaps in the disorder-free MI. At $U/t \to 0$, the data are consistent with Eq. (4). Reentrant behavior is also clearly seen in the phase diagram [6,47], as well as a large region where disorder induced superfluidity is possible (i.e., superfluidity is possible with disorder, but does not occur in the disorder-free system). The BG phase is connected to the $U = 0$ limit described by Anderson's localization. The transition temperature inside the superfluid finger is very low, e.g. $T_c/t = 0.37(5)$ for $\Delta/t = 65$ and $U/t = 60$ [30,31], and the corresponding condensate fraction at low temperature is equally low. What is also remarkable are the large scales for the SF–BG transition lines. For intermediate interactions, the superfluid phase extends to $\Delta/t \approx 300$, which at first sight seems to have little to do with the microscopic parameters of the Hubbard model. It was explained in Ref. [31] as follows: in a similar way as in Section 3.2, we expect that the transition between the SF and the BG is still given by percolation at moderate interaction strength and $\Delta \gg t$, but commensurability plays a role now and the localization length will be of the order of the lattice constant [48]. From the local Hamiltonian, the density can be computed, and setting it equal to one leads to the condition:

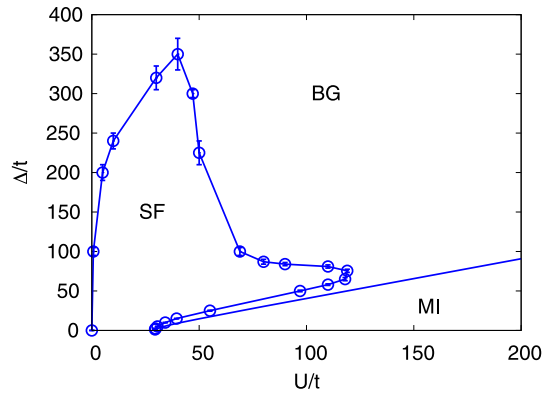$$\mu = -U/2 - \Delta + 2\sqrt{U\Delta} \tag{10}$$

Fig. 4. (Color online.) Phase diagram of the disordered 3d Bose–Hubbard model at unit filling, obtained by a finite-size analysis of winding numbers. In the absence of disorder, the system undergoes a quantum phase transition between SF and MI phases. The presence of disorder allows for a compressible, insulating BG phase, which always intervenes between the MI and SF phases because of the theorem of inclusions [30,31]. The transition between MI and BG is of the Griffiths type, as an exception implied by the theorem of inclusions [31]. At $U/t \to 0$, the SF–BG transition line has an infinite slope [28]. Reprinted figure with permission from Ref. [31]. Copyright (2009) by the American Physical Society.
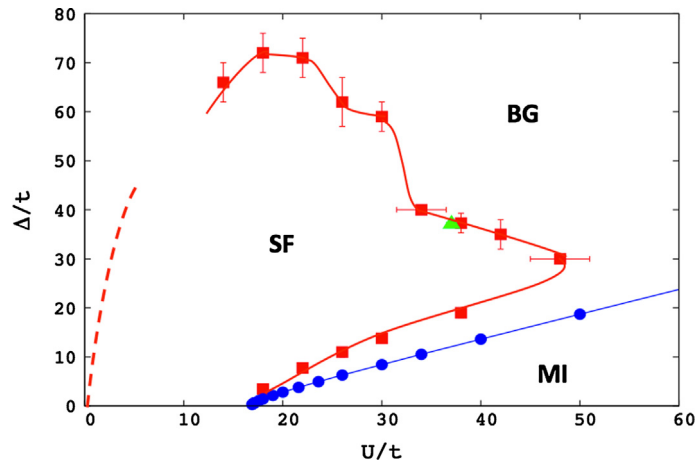

Fig. 5. (Color online.) Zero temperature phase diagram of the disordered 2d Bose–Hubbard model at unit filling. The MI–BG transition at $\Delta = E_{g/2}$ is obtained using the gap data from Ref. [51]. The green triangle is the point on the SF–BG boundary obtained in Ref. [52]. Reprinted figure with permission from Ref. [53]. Copyright (2011) by the American Physical Society.

A site will be occupied if its disorder lies within the $(-\Delta, U/2 + \mu) = (-\Delta, -\Delta + 2\sqrt{U\Delta})$ interval, which occurs with a probability $\sqrt{U/\Delta}$. Superfluidity requires at least that all occupied sites form a percolating cluster, hence we can put the transition line at:

$$\frac{U}{\Delta} \gtrsim \frac{1}{p_c^2} \tag{11}$$

with p_c the percolation threshold, which is $p_c \approx 0.31$ [49] for a simple cubic 3d lattice. This estimate is in good quantitative agreement with the Monte Carlo results shown in Fig. 4 for intermediate interaction strengths $U/t \leqslant U_c/t = 29.34(2)$. We shall see below that similar considerations also hold in lower dimensions. We also note that a quantitatively accurate mean-field study has recently been put forward in Ref. [50] for dimensions $d \geqslant 2$.

4.3. Disordered Bose–Hubbard model in 2d

In two dimensions, the topology of the phase diagram is the same as in 3d (see Fig. 5) [53]. Reentrant behavior is seen, there is also a superfluid "finger", and for intermediate interactions the superfluid survives up to disorder strengths roughly 10 times the bandwidth, $\Delta/t = 72$. The finite temperature phase transition between the superfluid and the normal phase is of Kosterlitz–Thouless type, with a very strong system-size dependence. The superfluid is hence very fragile against temperature fluctuations, but robust against strong interactions and disorder. Interestingly, the authors of Ref. [53] report that they were unable to make an unambiguous case for the form of the transition line $U_c(\Delta)$ in the vicinity of the tip of the Mott lobe (see Fig. 6). For very small Δ, the first critical points could not be distinguished from the critical value
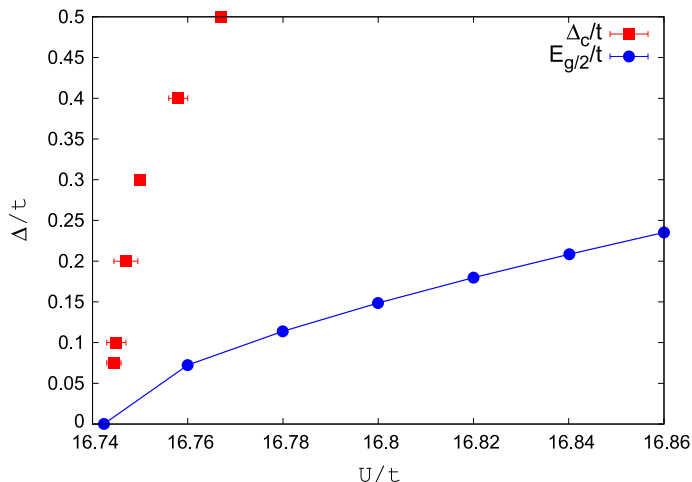
Fig. 6. (Color online.) Zero-temperature phase diagram of the disordered 2d Bose–Hubbard model at unit filling in the vicinity of the tip of the Mott lobe. The MI–BG transition at $\Delta = E_{g/2}$ is obtained using the gap data from Ref. [51]. Note the extremely weak dependence of the SF–BG critical line on disorder for small Δ. Reprinted figure with permission from Ref. [53]. Copyright (2011) by the American Physical Society.
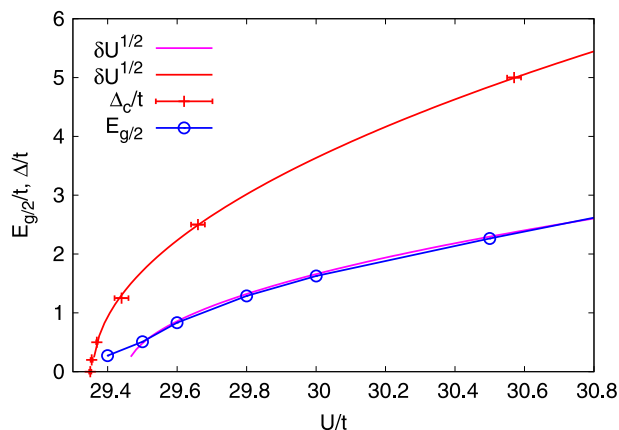

Fig. 7. (Color online.) Critical disorder bound Δ_c/t near the tip to the Mott insulator lobe for the 3d case. Also shown are the half gaps $E_{g/2}$ for the disorder-free system taken from Ref. [51]. The fits through the data are square roots. For Δ_c, the fit extends at least till $\Delta = 10$. For the disorder-free system, the mean-field exponent $\nu = 1/2$ fits through the data only for relatively small gaps, $E_{g/2} \leqslant 2t$. Close to the tip of the lobe, the gaps deviate from the fit, possibly due to the logarithmic corrections to scaling (the quantum critical point is exactly at the upper critical dimension). The fits for both quantities extrapolated to $\Delta = E_{g/2} = 0$ yield a small difference in the location of the critical point.

in the clean system at $U_c(\Delta = 0) = 16.7424(5)$ despite the very low error bars. The authors argue that disorder becomes relevant on length scales $L > (\kappa \Delta)^{-2/d}$, but this implies an exponentially large system size in 2d, $L > \xi > \exp(U/\Delta)$ when the phase of a vortex loop using Popov's hydrodynamic action is estimated. The SF–BG transition line is hence determined by the non-universal microscopic physics. Only in 1d does disorder determine the shape of the transition line [13].

These arguments imply that also in 3d we expect a vertical slope at U_c. It is seen in Fig. 7 that Δ_c is indeed insensitive to small changes in $U - U_c(\Delta = 0) = \delta U$, and the fit suggests $\Delta_c \propto (\delta U)^{1/2}$ (which is the most natural situation) and hence a vertical slope at U_c. The critical line $\Delta_c(U)$ has been determined on the basis of winding number fluctuations. To this end, we needed a more accurate determination of the tip of the lobe in the disorder-free case, namely $U_c = 29.350(5)$ (determined on the basis of winding number fluctuations for system sizes up to $L = 48$), in agreement with the previously published result $U_c = 29.34(2)$ [51] (determined on the basis of the gaps in the MI for system sizes up to $L = 20$). It is interesting to note in Fig. 7 that the extrapolation of the gaps in the pure MI and the extrapolation of the SF–BG transition points in case of disorder lead to a small difference in the location of the critical point. This may be related to logarithmic corrections to scaling in the disorder-free case, where we are exactly at the upper critical dimension.

In Ref. [52] the dynamic conductivity as a function of Matsubara frequency for a specific point on the SF–BG transition line in 2d was computed (see the point on the line $\Delta = U$ indicated by a triangle in Fig. 5). The universal conductivity in the high-frequency limit [54] was estimated to be around $0.17\sigma_Q$ with $\sigma_Q = e^{*2}/\hbar^2$ the quantum conductivity unit expressed in terms of the effective charge e^* of the bosons. This value was in agreement with previous calculations on quantum rotor
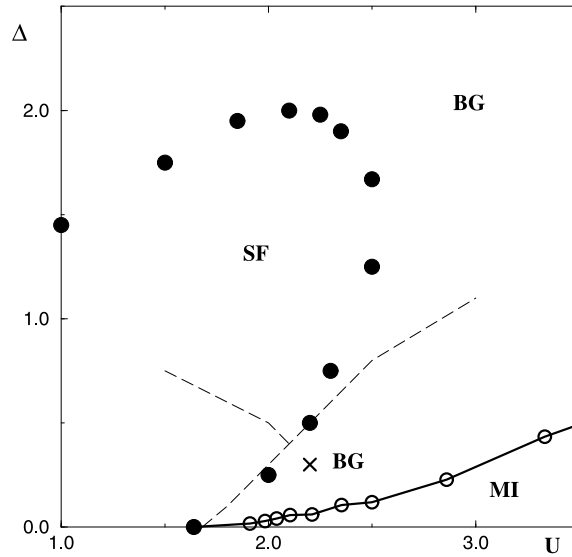
Fig. 8. Phase diagram of the disordered 1d Bose–Hubbard model at unit filling. For comparison with the units used in this paper, the values of U and Δ need to be multiplied by 2 in the figure. Error bars are smaller than the pointsize. The BG phase was verified for a system of size $L = 1000$ at the point marked by a cross. The dashed line indicates results obtained in Ref. [12] in violation of Section 3.3. Reprinted figure with permission from Ref. [56]. Copyright (1998) by the American Physical Society.

models [55]. The dynamic conductivity shows however clear deviations from scaling with ω, consistent instead with the expected scaling with ω/T.

4.4. Phase diagram of the disordered Bose–Hubbard model in 1d

In one dimension, Svistunov could show 10 years before the theorem of inclusions was known that the BG phase must intervene between the MI and SF [13]. The phase diagram was computed by Monte Carlo simulations [56] (shown in Fig. 8) and by the density matrix renormalization group [57]. (The DMRG data in Fig. 8 are not the ones of Ref. [57], whose data are not in violation of Section 3.3 and are closer to the Monte Carlo data, but the agreement is rather qualitative.) Interestingly, controversy remains about the nature of the SF–BG transition in one dimension. Giamarchi and Schulz performed a lowest-order RG scaling and found that the transition is of Kosterlitz–Thouless type with the Luttinger parameter taking an universal value $K = 3/2$ (instead of $K = 2$ for the SF–MI transition in the clean system) [2]. In Refs. [58,13], this was shown to hold for any finite disorder by invoking a picture of the unbinding of instanton–anti-instanton pairs. A recent calculation to two loops in the RG showed that the Luttinger parameter remains unchanged at the transition and also correlation functions keep universal exponents as long as bosonization applies [63]. However, speculation remained about a possible strong-disorder scenario that would be disorder driven instead of interaction driven. In the scenarios put forward in Refs. [59–61], the transition for $\Delta/U \gg 1$ is believed to be still of Kosterlitz–Thouless type, but with strong power-law distributions of the Luttinger parameter, whose average is then non-universal on the transition line. This RG procedure is also explained in the paper by Refael and Altman published in the same dossier of this journal [62].

The phase diagram shown in Fig. 8 was determined by locating where the Luttinger parameter takes a value $3/2$ on a system size of $L = 100$. Close to the MI, this is certainly fine and a full Kosterlitz–Thouless scaling analysis will not shift the data points outside the error bars. One notices that the transition line has not been determined for $U/t \leqslant 2$ (in the units used in this paper, not in the figure). Indeed, simulations show that the distribution of the Luttinger parameter becomes very wide, making simulations very cumbersome. Very recently, new insights revealed that a strong disorder fixed point leads to a prolonged flow of a renormalized classical field where the fugacity of the instanton–anti-instanton pairs is so low that they are ineffective on mesoscopic length scales [64]. It was also shown that self-averaging of the inverse stiffness, formulated in terms of a relative characteristic width determined by the percentiles of the distribution, always takes place in the superfluid phase and on the critical line (even if the variance would diverge), implying that the quantum transition in the thermodynamic limit always occurs for $K = 3/2$, but it may be masked by a logarithmically slow classical flow on all numerically accessible system sizes. This theory was illustrated for a J-current model [10] in $(1 + 1)$d in the vicinity of a first-order transition point in the disorder-free case (which speeds up the calculations substantially compared to the quantum models), but all the arguments equally apply to the 1d Bose–Hubbard model. Such a calculation has not been completed yet, but an analysis along the lines of Ref. [64] can now be done at moderate computational cost. The implications are that $K = 3/2$ holds everywhere on the SF–BG line in the thermodynamic limit (thus contrasting with Ref. [62]), but that the mesoscopic flow may be of greater physical relevance, and that the law in Eq. (4) also applies for weak U.

5. Experimental systems

Bosons with disorder can be realized in a number of different ways. They include Cooper pairs in thin superconducting films [65], Josephson junction arrays [66], and ^{4}He in porous media [67] and on substrates [68]. The detection of the Bose glass phase in these systems has remained ambiguous however. The experiments on ^{4}He in porous media, for example, were more convincingly explained by a model that has a constant density of states for low and for high energies, with a gap in between [67]. Recently, also cold atom systems came at our disposal, with their unique properties of tunability and control. In the experiments of Refs. [69,70] optical speckles were used to generate the disorder. It was not possible to distinguish between a Mott insulator and a Bose glass phase; only insulating phases could be distinguished from superfluid phases on the basis of time-of-flight interference images and transport measurements. They found insulating phases for disorder strengths several hundred times the tunneling amplitude, in agreement with the quantum Monte Carlo simulations for $U < U_c$ ($\Delta = 0$) in Fig. 4. However, they saw no signs of disorder-induced superfluidity (missing the 'finger' in Fig. 4). Although the disorder distributions are different in experiment and in simulations (see Ref. [71] for more realistic system parameters), the topology of the phase diagram should be the same. The discrepancy is attributed to the low transition temperature (or equivalently, the low superfluid density at zero temperature in the finger [31]), while the temperature in the experiment is estimated to be well above it. Finite system sizes (they are currently not larger than what can be done numerically) also play a role, certainly when only a single-disorder realization is used (as is currently the case). However, localization due to disorder but in the absence of interactions was successful [72], as well as in the context of the Aubry–André model [73].

Finally, we wish to mention that there have been a number of recent experiments on disordered quantum antiferromagnets [74–76], supported by simulations in Ref. [77]. These systems are, however, not without criticism either [78,79] and we refer to the paper by Zheludev and Roscilde that is published in the same dossier of this journal for a detailed discussion [80].

6. Significant others

There have been many interesting studies on disordered bosonic systems that are omitted here for reasons of space. We mention just a few examples.

Disordered bosonic systems in continuous space were studied in Refs. [81,82], where localization effects due to interaction (MI) are absent. The authors looked at the suppression of T_c of the SF caused by disorder, modeled by an isotropic 3d speckle potential. Unlike in the disorder-free system, T_c changed considerably between $na^3 = 10^{-4}$ and $na^3 = 10^{-6}$ (an older study found no substantial drop in T_c [83]). Agreement with a perturbative approach for δ-correlated disorder could not unambiguously be established, since the precision in the weak disorder limit was not high enough. For strong interactions and low-enough temperature, a normal phase with energy scaling as $\sim T^2$ was found, compatible with a BG at $T = 0$.

Instead of a microscopic quantum system, there have been simulations on classical J-current systems [10] in a higher dimension that have the same universal physics near the transitions, but different microscopics unrelated to the Bose–Hubbard model. An older study [84] in 1d found a Kosterlitz–Thouless transition compatible with a universal Luttinger value (cf. [1,2]), but the disorder was too weak to see the physics mentioned in Ref. [64]. In 2d, the $z = d = 2$ universality class on the SF–BG transition line could be established for diagonal disorder [19] with a finite and non-singular compressibility, but the universal behavior sets in only at very large space–time distances. The authors looked also at off-diagonal disorder, where the symmetry is different, resulting in the vanishing of the compressibility of the SF–BG transition line (the BG phase remains gapless) and a dynamical exponent numerically determined as $z = 1.5(2)$.

7. Conclusion

Following up on studies of helium-4 in porous media in the 1980s, Giamarchi and Schulz introduced the Bose–Glass phase in a disordered Bose–Hubbard Hamiltonian in one dimension [1,2]. Fisher et al. argued the existence of this phase in any dimension and provided an in-depth analysis [3]. The model has remained subject of debate ever since, even at the qualitative level. Especially the relevance of weak disorder at the SF–MI transition line remained controversial. Over the last five years, however, a number of analytical arguments and numerical simulations have been presented, which result in a fairly complete phase diagrams of the disordered Bose–Hubbard model in three [31] and two dimensions [53], with a final verdict on the issue whether a direct transition between the SF and the MI is possible in the presence of disorder (it is not). An analytical understanding of the shape of the superfluid 'finger' is still lacking, and also the shape of the SF–BG transition line near the tip of the Mott lobe requires further understanding. Even in one dimension, the completion of the phase diagram along the lines of Ref. [64] now seems possible. With the help of Monte Carlo simulations, we have thus arrived at a comprehensive understanding of the thermodynamic properties of the disordered Bose–Hubbard system, as well on mesoscopic scales as in the thermodynamic limit.

Acknowledgements

References

[1] T. Giamarchi, H.J. Schulz, Europhys. Lett. 3 (1987) 1287.
[2] T. Giamarchi, H.J. Schulz, Phys. Rev. B 37 (1998) 325.
[3] M.P.A. Fisher, P.B. Weichman, G. Grinstein, D.S. Fisher, Phys. Rev. B 40 (1989) 546.
[4] J.K. Freericks, H. Monien, Phys. Rev. B 53 (1996) 2691.
[5] R.T. Scalettar, G.G. Batrouni, G.T. Zimanyi, Phys. Rev. Lett. 66 (1991) 3144.
[6] W. Krauth, N. Trivedi, D. Ceperley, Phys. Rev. Lett. 67 (1991) 2307.
[7] L. Zhang, M. Ma, Phys. Rev. B 45 (1992) 4855.
[8] K.G. Singh, D.S. Rokhsar, Phys. Rev. B 46 (1992) 3002.
[9] M. Makivić, N. Trivedi, S. Ullah, Phys. Rev. Lett. 71 (1993) 2307.
[10] M. Wallin, E.S. Sørensen, S.M. Girvin, A.P. Young, Phys. Rev. B 49 (1994) 12115.
[11] F. Pazmandi, G. Zimanyi, R. Scalettar, Phys. Rev. Lett. 75 (1995) 1356.
[12] R.V. Pai, R. Pandit, H.R. Krishnamurthy, S. Ramasesha, Phys. Rev. Lett. 76 (1996) 2937.
[13] B.V. Svistunov, Phys. Rev. B 54 (1996) 16131.
[14] F. Pazmandi, G.T. Zimanyi, Phys. Rev. B 57 (1998) 5044.
[15] J. Kisker, H. Rieger, Phys. Rev. B 55 (R11) (1997) 981.
[16] I.F. Herbut, Phys. Rev. Lett. 79 (1997) 3502.
[17] I.F. Herbut, Phys. Rev. B 57 (1998) 13729.
[18] P. Sen, N. Trivedi, D.M. Ceperley, Phys. Rev. Lett. 86 (2001) 4092.
[19] N. Prokof'ev, B. Svistunov, Phys. Rev. Lett. 92 (2004) 015703.
[20] J. Wu, P. Phillips, Phys. Rev. B 78 (2008) 014515.
[21] U. Bissbort, W. Hofstetter, EPL 86 (2009) 50007.
[22] P.B. Weichman, R. Mukhopadhyay, Phys. Rev. B 77 (2008) 214516.
[23] P.B. Weichman, Mod. Phys. Lett. B 22 (2008) 2623.
[24] J. Carrasquilla, F. Becca, A. Trambettoni, M. Fabrizio, Phys. Rev. B 81 (2010) 195129.
[25] J. Carrasquilla, F. Becca, M. Fabrizio, Phys. Rev. B 83 (2011) 245101.
[26] S. Sachdev, Quantum Phase Transitions, Cambridge University Press, New York, 1999.
[27] T. Giamarchi, P. Le Doussal, E. Orignac, Phys. Rev. B 64 (2001) 245119.
[28] G.M. Falco, T. Nattermann, V.L. Pokrovsky, Phys. Rev. B 80 (2009) 104515.
[29] P. Lugan, D. Clément, P. Bouyer, A. Aspect, M. Lewenstein, L. Sanchez-Palencia, Phys. Rev. Lett. 98 (2007) 170403.
[30] L. Pollet, N.V. Prokof'ev, B.V. Svistunov, M. Troyer, Phys. Rev. Lett. 103 (2009) 140402.
[31] V. Gurarie, L. Pollet, N.V. Prokof'ev, B.V. Svistunov, M. Troyer, Phys. Rev. B 80 (2009) 214519.
[32] L. Pollet, Rep. Prog. Phys. 75 (2012) 094501.
[33] N.V. Prokof'ev, B.V. Svistunov, I.S. Tupitsyn, JETP 87 (1998) 310.
[34] G.D. Mahan, Many Particle Physics, 3rd ed., Springer Verlag, 2000.
[35] J.W. Negele, H. Orland, Quantum Many-Particle Systems, Perseus Books, 1988.
[36] A.L. Fetter, J.D. Walecka, Quantum Theory of Many-Particle Systems, McGraw-Hill, San Francisco, 1971.
[37] E.L. Pollock, D.M. Ceperley, Phys. Rev. B 36 (1987) 8343.
[38] A.W. Sandvik, Phys. Rev. B 59 (1999) 14157R.
[39] O.F. Syljuåsen, A.W. Sandvik, Phys. Rev. E 66 (2002) 046701.
[40] L. Pollet, K. Van Houcke, S. Rombouts, J. Comput. Phys. 225 (2007) 2249.
[41] V.G. Rousseau, Phys. Rev. E 77 (2007) 056705.
[42] V.G. Rousseau, Phys. Rev. E 78 (2008) 056707.
[43] S.M.A. Rombouts, K. Van Houcke, L. Pollet, Phys. Rev. Lett. 96 (2006) 180603.
[44] K. Van Houcke, S.M.A. Rombouts, L. Pollet, Phys. Rev. E 73 (2006) 056703.
[45] A. Priyadarshee, S. Chandrasekharan, J.-W. Lee, H.U. Baranger, Phys. Rev. Lett. 97 (2006) 115703.
[46] P.B. Weichman, R. Mukhopadhyay, Phys. Rev. Lett. 98 (2007) 245701.
[47] W. Krauth, N. Trivedi, Europhys. Lett. 14 (1991) 627.
[48] B. Bulka, M. Schreiber, B. Kramer, Z. Phys. B, Condens. Matter 66 (1987) 21.
[49] M.B. Isichenko, Rev. Mod. Phys. 64 (1992) 961.
[50] A. Niederle, H. Rieger, New J. Phys. 15 (2013) 075029.
[51] B. Capogrosso-Sansone, N.V. Prokof'ev, B.V. Svistunov, Phys. Rev. B 75 (2007) 134302.
[52] F. Lin, E.S. Sørensen, D.M. Ceperley, Phys. Rev. B 84 (2011) 094507.
[53] S.G. Söyler, M. Kiselev, N.V. Prokof'ev, B.V. Svistunov, Phys. Rev. Lett. 107 (2011) 185301.
[54] K. Damle, S. Sachdev, Phys. Rev. B 56 (1997) 8714.
[55] E.S. Sørensen, M. Wallin, S.M. Girvin, A.P. Young, Phys. Rev. B 46 (1992) 3002.
[56] N.V. Prokofev, B.V. Svistunov, Phys. Rev. Lett. 80 (1998) 4355.
[57] S. Rapsch, U. Schollwöck, W. Zwerger, Europhys. Lett. 46 (1999) 559.
[58] V.A. Kashurnikov, A.I. Podlivaev, N.V. Prokof'ev, B.V. Svistunov, Phys. Rev. B 53 (1996) 13091.
[59] E. Altman, Y. Kafri, A. Polkovnikov, G. Refael, Phys. Rev. Lett. 93 (2004) 150402.
[60] E. Altman, Y. Kafri, A. Polkovnikov, G. Refael, Phys. Rev. Lett. 100 (2008) 170402.
[61] E. Altman, Y. Kafri, A. Polkovnikov, G. Refael, Phys. Rev. B 81 (2010) 174528.
[62] G. Refael, A. Altman, to be published in CRS.
[63] Z. Ristivojevic, A. Petkovic, P. Le Doussal, T. Giamarchi, Phys. Rev. Lett. 109 (2012) 026402.
[64] L. Pollet, N.V. Prokof'ev, B.V. Svistunov, Phys. Rev. B 87 (2013) 144203.
[65] A.M. Goldman, Y. Liu, Physica D 83 (1995) 163.

[66] H.S.J. van der Zant, et al., Phys. Rev. B 54 (1996) 10081.
[67] P.A. Crowell, F.W. Van Keuls, J.D. Reppy, Phys. Rev. B 55 (1997) 12620.
[68] G.A. Csáthy, J.D. Reppy, M.W.H. Chan, Phys. Rev. Lett. 91 (2003) 235301.
[69] M. White, M. Pasienski, C. McKay, S.Q. Zhou, D.M. Ceperley, B. DeMarco, Phys. Rev. Lett. 102 (2009) 055301.
[70] M. Pasienski, D. McKay, M. White, B. DeMarco, Nat. Phys. 6 (2010) 677.
[71] S.Q. Zhou, D.M. Ceperley, Phys. Rev. A 81 (2010) 013402.
[72] J. Billy, V. Josse, Z. Zuo, A. Bernard, B. Hambrecht, P. Lugan, D. Clément, L. Sanchez-Palencia, P. Bouyer, A. Aspect, Nature 453 (2008) 891.
[73] G. Roati, C. D'Errico, L. Fallani, M. Fattori, C. Fort, M. Zaccanti, G. Modugno, M. Modugno, M. Inguscio, Nature 453 (2008) 895.
[74] F. Yamada, H. Tanaka, T. Ono, T. Nojiri, Phys. Rev. B 83 (2011) 020409.
[75] T. Hong, A. Zheludev, H. Manaka, L.P. Regnault, Phys. Rev. B 81 (2010) 060410.
[76] R. Yu, L. Yin, N.S. Sullivan, J.S. Xia, C. Huan, A. Paduan-Filho, N.F. Oliveira Jr., S. Haas, A. Steppke, C.F. Miclea, F. Weickert, R. Movshovich, E.D. Mun, V.S. Zapf, T. Roscilde, Nature 489 (2012) 379.
[77] T. Roscilde, S. Haas, Phys. Rev. Lett. 99 (2007) 047205.
[78] A. Zheludev, D. Hüvonen, Phys. Rev. B 83 (2011) 216401.
[79] E. Wulf, S. Mühlbauer, T. Yankova, A. Zheludev, Phys. Rev. B 84 (2011) 174414.
[80] A. Zheludev, T. Roscilde, arXiv:1305.1194 (to be published in CRS).
[81] S. Pilati, S. Giorgini, N.V. Prokof'ev, Phys. Rev. Lett. 102 (2009) 150402.
[82] S. Pilati, S. Giorgini, M. Modugno, N.V. Prokof'ev, New J. Phys. 12 (2010) 073003.
[83] M.C. Gordillo, D.M. Ceperley, Phys. Rev. Lett. 85 (2000) 4735.
[84] K.G. Balabanyan, N.V. Prokof'ev, B.V. Svistunov, Phys. Rev. Lett. 95 (2005) 055701.