Shift in Critical Temperature for Random Spatial Permutations with Cycle Weights

John Kerl

Abstract We examine a phase transition in a model of random spatial permutations which originates in a study of the interacting Bose gas. Permutations are weighted according to point positions; the low-temperature onset of the appearance of arbitrarily long cycles is connected to the phase transition of Bose-Einstein condensates. In our simplified model, point positions are held fixed on the fully occupied cubic lattice and interactions are expressed as Ewens-type weights on cycle lengths of permutations. The critical temperature of the transition to long cycles depends on an interaction-strength parameter α. For weak interactions, the shift in critical temperature is expected to be linear in α with constant of linearity c. Using Markov chain Monte Carlo methods and finite-size scaling, we find $c = 0.618 \pm 0.086$. This finding matches a similar analytical result of Ueltschi and Betz. We also examine the mean longest cycle length as a fraction of the number of sites in long cycles, recovering an earlier result of Shepp and Lloyd for non-spatial permutations.

Keywords Random graph · Phase transition · Branching process

1 Introduction

The model of random spatial permutations arises in the study of the Bose gas. Its history includes Bose-Einstein, Feynman [8], Penrose-Onsager [13], Sütő [18, 19], and Betz-Ueltschi-Gandolfo-Ruiz [3, 4, 9, 20]. Such random permutations arise physically when one symmetrizes the N-boson Hamiltonian with pair interactions, then applies a multi-particle Feynman-Kac formula and a cluster expansion [3, 4]. Specifically, given points $\mathbf{x}_1, \ldots, \mathbf{x}_N$ in the box $[0, L]^3$ and temperature T, permutations π are given probability weights proportional to the Gibbs factor $e^{-H(\pi)}$ where

$$H(\pi) = \frac{T}{4} \sum_{i=1}^{N} \|\mathbf{x}_i - \mathbf{x}_{\pi(i)}\|_\Lambda^2 \qquad (1.1)$$

J. Kerl (✉)
Department of Mathematics, University of Arizona, Tucson, AZ, USA
e-mail: kerl@math.arizona.edu

for the non-interacting case. The notation $\| \cdot \|_\Lambda$ indicates the natural distance on the 3-torus:

$$\|\mathbf{x} - \mathbf{y}\|_\Lambda = \min_{\mathbf{n} \in \mathbb{Z}^3} \{\|\mathbf{x} - \mathbf{y} + L\mathbf{n}\|\} \tag{1.2}$$

(The sum in (1.1) is scaled by temperature rather than reciprocal temperature. This surprising feature is opposite that of many models in statistical mechanics.) These energy terms involve lengths of permutation jumps; additional interaction terms take the form

$$\sum_{i < j} V(\mathbf{x}_i, \mathbf{x}_{\pi(i)}, \mathbf{x}_j, \mathbf{x}_{\pi(j)}), \tag{1.3}$$

i.e. permutation jumps from sites $\mathbf{x}_i$ and $\mathbf{x}_j$ interact pairwise. In the above-cited papers of Betz and Ueltschi, these may be approximated and rearranged such that one obtains interaction terms of the form

$$\sum_{\ell=1}^{N} \alpha_\ell r_\ell(\pi) \tag{1.4}$$

where $r_\ell(\pi)$ counts the number of ℓ-cycles of the permutation π, and the coefficients α_ℓ are *cycle weights*.

When the Bose gas is cooled below a critical temperature T_c, there is a phase transition: a macroscopic fraction of the bosons are found in the ground state of the external potential, and such particles are said to participate in a Bose-Einstein condensate. In the permutation representation, this transition manifests itself as the onset of long permutation cycles. Bose suggested the statistics carrying his name for describing the gas of photons; Einstein developed the notion of what we now call Bose-Einstein condensation, and computed the critical temperature for the non-interacting Bose gas. The critical temperature for liquid helium, where interparticle interactions are strong, is lower than would be expected [1] for non-interacting atoms of the same density. For weakly interacting systems, however, an emerging consensus is that interactions *increase* the critical temperature. See [1, 16] for surveys. Concretely, for interactions parameterized by some α, one defines

$$\Delta T_c(\alpha) = \frac{T_c(\alpha) - T_c(0)}{T_c(0)}.$$

It is well accepted that

$$\lim_{\alpha \to 0} \Delta T_c(\alpha) = c\rho^{1/3}\alpha$$

where ρ is the particle density, i.e. that for small α the shift in critical temperature is linear in α. What is more contentious, as enumerated in the surveys cited above, is the value of the constant c.

The interaction terms (1.3) for the permutation representation of the Bose gas are difficult to compute. Moreover, it is interesting to consider the model of random spatial permutations for its own sake. In [9], a simulational approach is taken for points held fixed on the fully occupied unit lattice in the non-interacting case. In the papers [3, 20], Betz and Ueltschi examine the Bose-gas permutation weights with point positions allowed to vary on the continuum; an exact expression for the critical temperature is stated and proved for a simplified interaction model in which only two-cycles interact. That is, interactions are of the form of (1.4) with $\alpha_2 = \alpha$, where α is related to a hard-core scattering length, and the remaining

cycle weights are zero. In [4], this approach is extended to a model in which all the α_ℓ's may vary, but with the hypothesis that α_ℓ goes to zero faster than $1/\log(\ell)$. In this paper, we take a simulational approach to points on the fully occupied unit lattice, with cycle weights constant in ℓ—removing the decaying-cycle-weight hypothesis. The shift in critical temperature is nonetheless found to match that predicted by Betz and Ueltschi.

An outline of the paper is as follows. Section 2 provides background necessary to understand the results of the paper: the probability model is defined in Sect. 2.1; qualitative and quantitative behavior of long cycles are discussed in Sects. 2.2 and 2.3, respectively. Known results and conjectures are listed in Sect. 2.4. In Sect. 3, the simulational methods are presented. The *swap-only* and *swap-and-reverse* algorithms generate simulational data; these algorithms are proved correct in Sects. 3.1 through 3.3. The finite-size-scaling method, which reduces the raw simulational data, is summarized in Sect. 3.5. Section 4 presents the data and its analysis in full detail: estimation of critical exponents and critical temperature in Sects. 4.1 through 4.3, verification of the finite-size-scaling hypothesis in Sect. 4.4, and final results in Sects. 4.5 through 4.7.

2 The Model of Random Spatial Permutations

Here we review concepts from [3, 4], fixing notation and intuition to be used in the rest of the paper.

2.1 The Probability Model

The state space is $\Omega_{\Lambda,N} = \Lambda^N \times \mathcal{S}_N$, where $\Lambda = [0, L]^3$ with periodic boundary conditions; point positions are $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_N)$ for $\mathbf{x}_1, \ldots, \mathbf{x}_N \in \Lambda$. The Hamiltonian takes one of two forms. In the first, relevant to the Bose gas, we have

$$H(\mathbf{X}, \pi) = \frac{T}{4} \sum_{i=1}^{N} \|\mathbf{x}_i - \mathbf{x}_{\pi(i)}\|_\Lambda^2 + \sum_{i<j} V(\mathbf{x}_i, \mathbf{x}_{\pi(i)}, \mathbf{x}_j, \mathbf{x}_{\pi(j)}) \qquad (2.1)$$

where $T = 1/\beta$ and the V terms are interactions between permutation jumps. (The temperature scale factor $T/4$, not $\beta/4$, is surprising but correct for the Bose-gas derivation of the Hamiltonian.) In the second form of the Hamiltonian, considered in this paper, we use interactions which are dependent solely on cycle length:

$$H(\mathbf{X}, \pi) = \frac{T}{4} \sum_{i=1}^{N} \|\mathbf{x}_i - \mathbf{x}_{\pi(i)}\|_\Lambda^2 + \sum_{\ell=1}^{N} \alpha_\ell r_\ell(\pi), \qquad (2.2)$$

where $r_\ell(\pi)$ is the number of ℓ-cycles in π and the α_ℓ's are free parameters, called *cycle weights*. One ultimately hopes to choose the α_ℓ's appropriately for the Bose gas; even if not, the model is well-defined and of its own interest.

Different choices of α_ℓ result in different models: The *non-interacting model* [9] has $\alpha_\ell \equiv 0$. The *two-cycle model* [3, 20], has $\alpha_2 = \alpha$ and other cycle weights are zero. The *general-cycle model* has no restrictions on α_ℓ. In [4], the *decaying cycle-weight case* of the general-cycle model is considered: the only restriction on α_ℓ is that α_ℓ goes to zero in ℓ faster than $1/\log \ell$. The *Ewens model*, treated in this paper (see also [7]), is another special case of the general-cycle model: it has $\alpha_\ell \equiv \alpha$ constant in ℓ.

Fig. 1 A spatial permutation
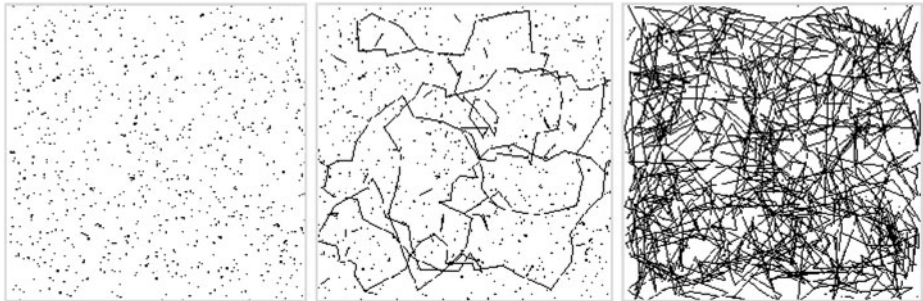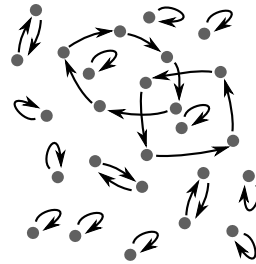
Fig. 2 Some typical permutations for high T, medium but subcritical T, and low T

One may hold point positions fixed, e.g. on the fully occupied unit lattice; this approach has been taken for all simulations done up to the present by the author and by Gandolfo [9], including specifically the work described in this paper. One obtains a Gibbs probability distribution on $\mathcal{S}_N$:

$$Y(\Lambda, \mathbf{X}) = \sum_{\sigma \in \mathcal{S}_N} e^{-H(\mathbf{X}, \sigma)}, \qquad P(\pi) = \frac{e^{-H(\mathbf{X}, \pi)}}{Y(\Lambda, \mathbf{X})}. \tag{2.3}$$

(Alternatively, one may integrate over all positions in Λ, with a resulting Gibbs distribution on $\mathcal{S}_N$. Here, several analytical results are available [3, 4].) For a random variable $S(\pi)$, we have

$$\mathbb{E}[S] = \sum_{\pi \in \mathcal{S}_N} P(\pi) S(\pi). \tag{2.4}$$

2.2 Qualitative Characterization of Long Cycles

One next inquires which permutations are typical in this temperature-dependent probability distribution on $\mathcal{S}_N$. In this section we develop intuition; in the next section, we construct quantitative descriptions of the ideas presented here.

As $T \to \infty$, the probability measure becomes supported only on the identity permutation: the distance-dependent terms are large whenever any jump has non-zero length. For large but finite T, there are tiny islands of 2-cycles, 3-cycles, etc. On the other hand, as $T \to 0$, length-dependent terms go to zero, and the probability measure approaches the uniform distribution on $\mathcal{S}_N$: the distance-dependent terms all go to zero. For intermediate T, one observes that the length $\|\pi(\mathbf{x}) - \mathbf{x}\|_\Lambda$ of each permutation jump remains small, increasing smoothly as T drops.
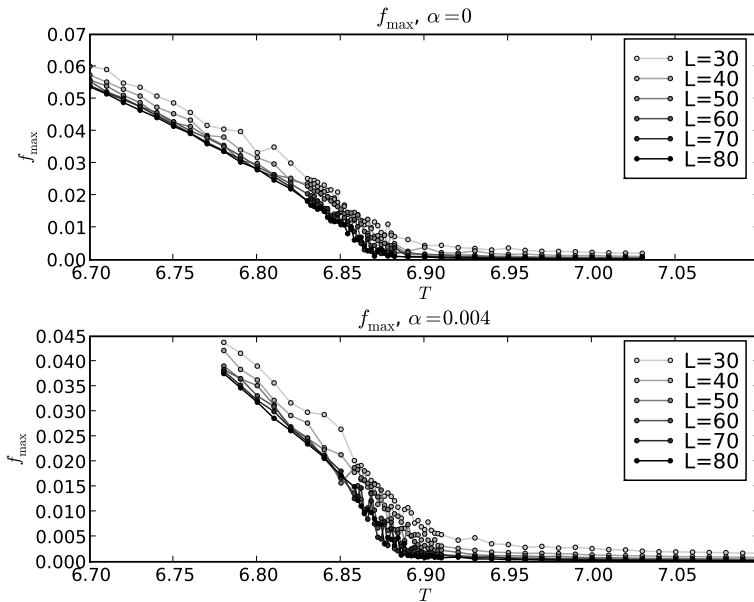
Fig. 3 Order parameter $f_M = \mathbb{E}[\ell_{\max}]/N$ for finite systems, with $\alpha = 0, 0.004$. Interactions increase the critical temperature. The shift is slight, but visible; we work in the regime of small interaction parameters. See Sect. 4 for a quantitative analysis of this shift

For T above a critical temperature T_c, all cycles are short: two-cycles, three-cycles, and so on. We find $T_c \approx 6.86$ at $\alpha = 0$, and positive α terms increase T_c. At T_c, though, there is a phase transition: for $T < T_c$ jump lengths remain short but *long cycles form*. Quantitatively, let $\ell_{\max}$ be the length of the longest cycle in π, with $\mathbb{E}[\ell_{\max}]$ its mean over all permutations. We observe that for $T > T_c$, $\mathbb{E}[\ell_{\max}]$ grows only perhaps as fast as $\log(L)$ as $L \to \infty$. That is to say, $\mathbb{E}[\ell_{\max}]/N$ goes to zero as $N \to \infty$. For $T < T_c$, on the other hand, $\mathbb{E}[\ell_{\max}]$ scales with N, i.e. $\mathbb{E}[\ell_{\max}]/N$ approaches a temperature-dependent constant as $N \to \infty$: there are arbitrarily long cycles, or infinite cycles, in the infinite-volume limit. See Fig. 2 for depictions of typical permutations at high T, subcritical T, and low T; see Fig. 3 for plots of $\mathbb{E}[\ell_{\max}]/N$ as a function of T for various system sizes with $N = L^3$. Note in particular that higher alpha shifts the order-parameter curve to the right, with resulting upward shift in critical temperature T_c.

Feynman's claim for the Bose gas is that Bose-Einstein condensation occurs if and only if there are infinite cycles in the infinite-volume limit. The central point of this approach is that the system energy has been recast in terms of permutations, which are amenable to analysis and simulation. This permits a new perspective on the venerable question: how does the critical temperature of Bose-Einstein condensation depend on inter-particle interaction strength?

Obtaining a full answer to this notoriously difficult question is a long-term project. As an intermediate step, we here consider the Ewens cycle-weight Hamiltonian with point positions on the unit fully occupied unit lattice. Through careful use of MCMC algorithms, statistical analysis, and finite-size scaling, we are able to quantify the dependence of critical temperature on interaction strength.

2.3 Quantitative Characterization of Long Cycles

Various order parameters may be defined; all of them may be used to locate the critical temperature $T_c(\alpha)$. The fraction $\mathbb{E}[\ell_{\max}]/N$ discussed above will, for brevity, be hereafter referred to as f_M. The *fraction of sites in long cycles*, f_I, is described in detail in [9]. The *correlation length* $\xi(T)$ is defined to be the spatial length of the cycle containing a given point $\mathbf{x}$: for $T < T_c$, it blows up as L increases. Namely, we define

$$s_{\mathbf{x}}(\pi) = \|\pi(\mathbf{x}) - \mathbf{x}\|_{\Lambda} \quad \text{and} \quad s(\pi) = \frac{1}{N} \sum_{i=1}^{N} s_{\mathbf{x}_i}(\pi).$$

The expectation over all π of $s_{\mathbf{x}}$ is the same as s, of course; in a Monte Carlo simulation, however, the latter yields a larger sample size and thus a smaller error bar. We use $\xi(T) = \mathbb{E}[s]$.

   Winding numbers count the integer number of x, y, z wraps around the 3-torus (Λ with periodic boundary conditions). Specifically, the winding number of a permutation π is the triple

$$\mathbf{W} = (W_x, W_y, W_z) = \frac{1}{L} \sum_{i=1}^{N} \mathbf{d}_{\Lambda}(\mathbf{x}_{\pi(i)}, \mathbf{x}_i), \tag{2.5}$$

where $\mathbf{d}_{\Lambda}$ is the difference vector defined as follows. For $\mathbf{z} \in \Lambda$, we define a zero-centered modulus vector $\mathbf{m}_L(\mathbf{z})$. For $\mathbf{x}, \mathbf{y} \in \Lambda$, this gives rise to a difference vector $\mathbf{d}_{\Lambda}(\mathbf{x}, \mathbf{y})$:

$$\mathbf{m}_L(\mathbf{z}) = \begin{pmatrix} m_L(z_1) \\ m_L(z_2) \\ m_L(z_3) \end{pmatrix} \tag{2.6}$$

$$n_L(z) = n \in \mathbb{Z} \quad \text{which minimizes } |z + nL| \tag{2.7}$$

$$m_L(z) = z + n_L(z)L \tag{2.8}$$

$$\mathbf{d}_{\Lambda}(\mathbf{x}, \mathbf{y}) = \mathbf{m}_{\Lambda}(\mathbf{x} - \mathbf{y}). \tag{2.9}$$

We also write

$$\mathbf{W}^2 = \mathbf{W} \cdot \mathbf{W} = W_x^2 + W_y^2 + W_z^2.$$

   The *scaled winding number* [14] is

$$f_S = \frac{\langle \mathbf{W}^2 \rangle L^2}{3\beta N}.$$

Lastly, the order parameter f_W is the fraction of sites which participate in winding cycles.
   The order parameters $f_I(T)$, $f_S(T)$, and $f_W(T)$ show behavior similar to $f_M :=$ $\mathbb{E}[\ell_{\max}]/N$ (Fig. 3): asymptotically as $N \to \infty$, they are zero for $T \geq T_c$ and non-zero for $T < T_c$. For finite N, the curves remain analytic: finite-size effects persist. The inverse correlation length $1/\xi(T)$, on the other hand, is zero for $T \leq T_c$ and non-zero for $T > T_c$.
   Our goal is to quantify the dependence of T_c on α, where

$$\Delta T_c(\alpha) = \frac{T_c(\alpha) - T_c(0)}{T_c(0)}. \tag{2.10}$$

Known results and conjectures are formulated quantitatively in terms of $\lim_{\alpha \to 0} \Delta T_c(\alpha)$.

2.4 Known Results and Conjectures

Known results for point locations averaged over the continuum are obtained largely using Fourier methods [4], which are unavailable for point positions held fixed on the lattice. Betz and Ueltschi have determined $\Delta T_c(\alpha)$, to first order in α, for two-cycle interactions [3] and decaying cycle weights [4]. (This taps into a long and controversial history in the physics literature: see [1] or [16] for surveys.) The critical (ρ, T, α) manifold relates ρ_c to T_c. Specifically,

$$\rho_c(\alpha_1, \alpha_2, \ldots) = \sum_{\ell \geq 1} e^{-\alpha_\ell} \int_{\mathbb{R}^3} e^{-\ell 4\pi^2 \beta \|\mathbf{k}\|^2} \, d\mathbf{k} = \frac{1}{(4\pi\beta)^{3/2}} \sum_{\ell \geq 1} e^{-\alpha_\ell} \ell^{-3/2} \qquad (2.11)$$

$$\Delta T_c(\alpha) = c\rho^{1/3}\alpha, \quad \text{for } \alpha \approx 0. \qquad (2.12)$$

Using this formula for constant cycle weights $\alpha_\ell \equiv \alpha$ and for lattice density $\rho = 1$, we have

$$\rho_c = \frac{\zeta(3/2)e^{-\alpha}}{(4\pi\beta)^{3/2}}, \qquad T_c = \frac{4\pi e^{2\alpha/3}}{\zeta(3/2)^{2/3}} \approx 6.626 e^{2\alpha/3},$$
$$\Delta T_c(\alpha) = \frac{T_c(\alpha) - T_c(0)}{T_c(0)} = e^{2\alpha/3} - 1 \approx \frac{2\alpha}{3}, \quad c \approx 0.667. \qquad (2.13)$$

We inquire whether this result, obtained for decaying cycle weights with point positions varying on the continuum, holds for Ewens weights with point positions held fixed on the lattice.

For $\alpha_\ell \equiv 0$ (the non-interacting model), $\mathbb{E}[\ell_{\max}]/Nf_I$ is constant for T below but near T_c. (That is, the two order parameters f_I and $\mathbb{E}[\ell_{\max}]/N$ have the same critical exponent.) For uniform-random permutations (Shepp and Lloyd 1966 [17] solved Golomb's 1964 question [10]), $\mathbb{E}[\ell_{\max}]/N \approx 0.6243$; unpublished work of Betz and Ueltschi has found $\mathbb{E}[\ell_{\max}]/Nf_I$ is that same number for the non-interacting case $\alpha_\ell \equiv 0$. The intuition is that long cycles are uniformly distributed within the zero Fourier mode. (This was proved in Sect. 5 of [18]. Other results on the distribution of the length and number of cycles for probabilities depending only on the conjugacy class can be found in Sects. 2 and 5 of [18], and in [19].) We conjecture that $\mathbb{E}[\ell_{\max}]/Nf_I$ is α-dependent but constant in T (for T below but near T_c) for all interaction models.

We suspect that the fine details of point positions are unimportant for the shift in critical temperature. Thus, $\Delta T_c(\alpha)$ on the lattice should be similar to that on the continuum, if decaying cycle weights are used. For Ewens interactions, though, $\Delta T_c(\alpha)$ is theoretically unknown for Ewens interactions with points either on the continuum or on the lattice. The simulational treatment in this paper is the only known attack on this question.

3 Simulational Methods

We run Markov chain Monte Carlo experiments for various values of L, T, and interaction strength α. For each parameter combination, we generate M typical permutations $\pi_1, \ldots, \pi_M$ from the stationary distribution, using MCMC algorithms described below, and we compute random variables $X_i = X(\pi_i)$. (The values of M used are 10^5 away from T_c, and 10^6 near T_c where sample variance is higher.) We find the sample mean and estimate the variance of the sample mean. The correlation of the X_i's complicates the latter. Finite-size scaling compensates for finite-size effects: mathematically, we are interested in estimating infinite-volume quantities based on finite-volume numerical experiments.

Fig. 4 Metropolis moves for the swap-only algorithm

$$\text{swap w.p.} \sim 1 \wedge e^{-\Delta H}$$
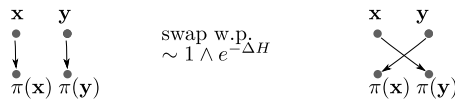
3.1 The Swap-Only Algorithm

Recall from Sect. 2.1 that the expectation of a random variable S (such as ξ, f_M, f_W, f_I, f_S) is

$$\mathbb{E}[S] = \sum_{\pi \in \mathcal{S}_N} P(\pi) S(\pi).$$

The number of permutations, $N!$, grows intractably in N. As is typical in Markov chain Monte Carlo methods [2, 12], one contents oneself with a smaller number of samples: the expectation is instead estimated by summing over some number M (10^5 or 10^6) of typical permutations.

The *swap-only algorithm* for sampling from the Gibbs distribution (2.3) is as follows:

- Start with the identity or uniform-random permutation.
- Sweep through sites $\mathbf{x}$ of the lattice in either lexical or uniform-random order.
- For each site $\mathbf{x}$, do a *Metropolis step*:
  - Choose a site $\pi(\mathbf{y})$ from among the six nearest neighbors of $\pi(\mathbf{x})$.
  - Propose to change π to the permutation π' which has $\pi'(\mathbf{z}) = \pi(\mathbf{z})$ for all $\mathbf{z} \neq \mathbf{x}, \mathbf{y}$ but $\pi'(\mathbf{x}) = \pi(\mathbf{y})$ and $\pi'(\mathbf{y}) = \pi(\mathbf{x})$. (See Fig. 4.)
  - With probability proportional to $\min\{1, e^{-\Delta H}\}$ where $\Delta H = H(\pi') - H(\pi)$, accept the change. (If the change is rejected, $\pi' = \pi$.)
- After each sweep, obtain a value of each random variable for inclusion in computation of its sample mean.

One starts accumulating data only after a suitable number of *thermalization sweeps*. The idea is that the initial, identity permutation is not typical, nor are the first few afterward. The integrated autocorrelation time [2] of system energy H gives an idea of how many Metropolis sweeps should be discarded before the permutations become typical. Also, one may examine H to ensure that it has reached its long-term average value. This is explained in detail in [11]. We next prove correctness of this algorithm.

3.2 Explicit Construction of the Markov Matrix

For Sect. 3.3 we will need an explicit construction of the Markov matrix corresponding to the swap-only algorithm as described in Sect. 3.1. The Markov perspective on the algorithm is that the distribution $P^{(0)}(\pi)$ of the first permutation is either supported solely on the identity, or uniform on all $N!$ permutations. The distribution for subsequent permutations is

$$P^{(k+1)}(\pi') = \sum_{\pi \in \mathcal{S}_N} P^{(k)}(\pi) M(\pi, \pi')$$

or, in matrix/vector notation,

$$\mathbf{P}^{(k+1)} = \mathbf{P}^{(k)} \mathbf{M}.$$

In this section we precisely describe the matrix $\mathbf{M}$; in Sect. 3.3 we show that $\mathbf{P}^{(k)}$ approaches the Gibbs distribution (2.3).

The matrix $\mathbf{M}$ is $N! \times N!$: rows are indexed by $\pi_1, \ldots, \pi_{N!}$ and columns are indexed by $\pi'_1, \ldots, \pi'_{N!}$. Most of the entries of $\mathbf{M}$ are zero: Metropolis steps change only two permutation sites whereas most π, π' differ at more than two sites.

Definition 3.1 For $\pi, \pi' \in \mathcal{S}_N$, define

$$d(\pi, \pi') = \#\{i = 1, 2, \ldots, N : \pi(i) \neq \pi'(i)\}.$$

Remark Note that $d(\pi, \pi') \neq 1$ since if two permutations agree on $N - 1$ sites, they must agree on the remaining site. It is easily shown that the function $d(\pi, \pi')$ is a metric on $\mathcal{S}_N$.

Definition 3.2 Lattice sites $\mathbf{x}, \mathbf{y}$ are *nearest-neighbor* if $\|\mathbf{x} - \mathbf{y}\|_\Lambda = 1$.

Definition 3.3 For $\pi \in \mathcal{S}_N$, define

$$R(\pi) = \{\pi' \in \mathcal{S}_N : d(\pi, \pi') = 2 \text{ and } \|\pi(\mathbf{x}) - \pi(\mathbf{y})\|_\Lambda = 1\}$$

where the $\mathbf{x}$ and $\mathbf{y}$ are taken to be the two points at which π, π' differ. Then $R(\pi)$ is the set of permutations π' reachable from π on a swap.

We construct the Markov matrix for use when sites $\mathbf{x}$ are selected at uniform random. (The matrices for use when $\mathbf{x}$ is selected sequentially are similar.) For each $\pi \in \mathcal{S}_N$,

$$M(\pi, \pi') = \begin{cases} \frac{1}{3N}\left(1 \wedge e^{-H(\pi') + H(\pi)}\right), & \pi' \in R(\pi), \\ 1 - \sum_{\pi'' \in R(\pi)} \frac{1}{3N}\left(1 \wedge e^{-H(\pi'') + H(\pi)}\right), & \pi = \pi'; \\ 0, & \text{otherwise.} \end{cases} \tag{3.4}$$
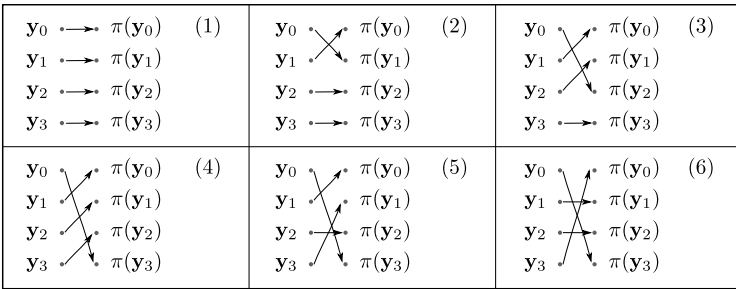
To justify the choice of prefactor $1/3N$, note that there are N choices of lattice points $\mathbf{x}$. For each $\mathbf{x}$, there are 6 choices of $\pi(\mathbf{y})$ which are nearest neighbors to $\pi(\mathbf{x})$. This double-counts the $3N$ distinct choices of π' reachable from π in a single Metropolis step, since choosing $\mathbf{x}$ and then $\mathbf{y}$ results in the same Metropolis step as choosing $\mathbf{y}$ and then $\mathbf{x}$.

3.3 Correctness of the Swap-Only Algorithm

It is clear that the swap-only algorithm produces a sequence of permutations, but with what distribution? From Markov-chain theory, we know the following: If the chain is irreducible, aperiodic, and satisfies detailed balance, then the chain has the Gibbs distribution (2.3) as its unique invariant distribution.

We note the following terminology: *detailed balance* is the same as *reversibility*. Also, an irreducible, aperiodic chain on a finite state space is called *ergodic*. Also note from Markov-chain theory that all states in a recurrence class have the same period. Thus, if we can show that the chain is irreducible (i.e. the entire state space is a single recurrence class), then for aperiodicity of the chain it suffices to show that a single state (e.g. the identity permutation) has period 1.

Proposition 3.5 (Irreducibility) *For all π, π', there is an n such that $M^n(\pi, \pi') > 0$. That is, any permutation is reachable from any other.*

$\mathbf{y}_0 \longrightarrow \pi(\mathbf{y}_0)$ (1)	$\mathbf{y}_0 \pi(\mathbf{y}_0)$ (2)	$\mathbf{y}_0 \pi(\mathbf{y}_0)$ (3)
$\mathbf{y}_1 \longrightarrow \pi(\mathbf{y}_1)$	$\mathbf{y}_1 \pi(\mathbf{y}_1)$	$\mathbf{y}_1 \pi(\mathbf{y}_1)$
$\mathbf{y}_2 \longrightarrow \pi(\mathbf{y}_2)$	$\mathbf{y}_2 \longrightarrow \pi(\mathbf{y}_2)$	$\mathbf{y}_2 \pi(\mathbf{y}_2)$
$\mathbf{y}_3 \longrightarrow \pi(\mathbf{y}_3)$	$\mathbf{y}_3 \longrightarrow \pi(\mathbf{y}_3)$	$\mathbf{y}_3 \longrightarrow \pi(\mathbf{y}_3)$
$\mathbf{y}_0 \pi(\mathbf{y}_0)$ (4)	$\mathbf{y}_0 \pi(\mathbf{y}_0)$ (5)	$\mathbf{y}_0 \pi(\mathbf{y}_0)$ (6)
$\mathbf{y}_1 \pi(\mathbf{y}_1)$	$\mathbf{y}_1 \pi(\mathbf{y}_1)$	$\mathbf{y}_1 \pi(\mathbf{y}_1)$
$\mathbf{y}_2 \pi(\mathbf{y}_2)$	$\mathbf{y}_2 \pi(\mathbf{y}_2)$	$\mathbf{y}_2 \pi(\mathbf{y}_2)$
$\mathbf{y}_3 \pi(\mathbf{y}_3)$	$\mathbf{y}_3 \pi(\mathbf{y}_3)$	$\mathbf{y}_3 \pi(\mathbf{y}_3)$

Fig. 5 A sequence of (nearest-neighbor) swaps which results in a non-nearest-neighbor swap

Proof Transpositions generate $\mathcal{S}_N$: for all $\pi \in \mathcal{S}_N$, there exist transpositions $\sigma_1, \ldots, \sigma_m$ such that $\pi = \prod_{j=1}^m \sigma_j$. Thus, it suffices to show that given any permutation π and any two points $\mathbf{x}$ and $\mathbf{z}$, so $\pi : \mathbf{x} \mapsto \pi(\mathbf{x})$ and $\pi : \mathbf{z} \mapsto \pi(\mathbf{z})$, we can construct a sequence of swaps sending π to π' so that $\pi' : \mathbf{x} \mapsto \pi(\mathbf{z})$, $\pi' : \mathbf{z} \mapsto \pi(\mathbf{x})$, and $\pi'(\mathbf{y}) = \pi(\mathbf{y})$ for all $\mathbf{y} \neq \mathbf{x}, \mathbf{z}$. (If $\pi(\mathbf{x})$ and $\pi(\mathbf{z})$ are nearest-neighbor lattice sites, of course, then a single swap does the job.)

Define $G_{\mathbf{a},\mathbf{b}} : \mathcal{S}_N \to \mathcal{S}_N$ to be the swap operator for nearest-neighbor lattice sites $\pi(\mathbf{a})$ and $\pi(\mathbf{b})$. Write $\pi' = G_{\mathbf{a},\mathbf{b}}\pi$. Given $\mathbf{x}$ and $\mathbf{z}$, there is a (non-unique) sequence of lattice sites $\mathbf{y}_0, \mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n$ such that $\mathbf{y}_0 = \mathbf{x}$, $\mathbf{y}_n = \mathbf{z}$, and $\|\pi(\mathbf{y}_{i+1}) - \pi(\mathbf{y}_i)\|_\Lambda = 1$ for $i = 0, 1, \ldots, n-1$. (See Fig. 5.) We will construct a sequence of swaps along this nearest-neighbor path whose end result is to swap the permutation arrows starting at $\mathbf{x}$ and $\mathbf{z}$, leaving all other arrows unchanged. We first need a lemma about compositions of swaps.

Notation 3.6 *Given* $\mathbf{x}_1, \ldots, \mathbf{x}_N$ *and a permutation* π, *we may write* π *as an* image map *with the* $\mathbf{x}_i$*'s along the top row and their images along the bottom row*:

$$\begin{pmatrix} \mathbf{x}_1 & \ldots & \mathbf{x}_N \\ \pi(\mathbf{x}_1) & \ldots & \pi(\mathbf{x}_N) \end{pmatrix}.$$

We find that the composition of maps

$$(G_{\mathbf{y}_n,\mathbf{y}_1} \circ G_{\mathbf{y}_n,\mathbf{y}_2} \circ \ldots \circ G_{\mathbf{y}_n,\mathbf{y}_{n-2}} \circ G_{\mathbf{y}_n,\mathbf{y}_{n-1}}) \circ (G_{\mathbf{y}_0,\mathbf{y}_n} \circ G_{\mathbf{y}_0,\mathbf{y}_{n-1}} \circ \ldots \circ G_{\mathbf{y}_0,\mathbf{y}_2} \circ G_{\mathbf{y}_0,\mathbf{y}_1}) \tag{3.7}$$

swaps the images of $\mathbf{x} = \mathbf{y}_0$ and $\mathbf{z} = \mathbf{y}_n$ while leaving all other images unchanged, that is,

$$\begin{pmatrix} \mathbf{y}_0 & \mathbf{y}_1 & \ldots & \mathbf{y}_{n-1} & \mathbf{y}_n \\ \pi(\mathbf{y}_0) & \pi(\mathbf{y}_1) & \ldots & \pi(\mathbf{y}_{n-1}) & \pi(\mathbf{y}_n) \end{pmatrix}$$

$$\mapsto \begin{pmatrix} \mathbf{y}_0 & \mathbf{y}_1 & \ldots & \mathbf{y}_{n-1} & \mathbf{y}_n \\ \pi(\mathbf{y}_n) & \pi(\mathbf{y}_1) & \ldots & \pi(\mathbf{y}_{n-1}) & \pi(\mathbf{y}_0) \end{pmatrix}. \qquad \square$$

Remark Below we will discuss winding numbers, and the empirical fact that the swap-only algorithm changes them only rarely. The chain is irreducible but various non-zero transition probabilities can still be very small.

Definition 3.8 The *period* of π is

$$p(\pi) = \gcd\{n : P(\Pi_n = \pi \mid \Pi_0 = \omega) > 0\}$$

where Π_k is the random variable which is the permutation appearing at the kth step of the Markov chain. We say that π has period p if it reappears with probability 1 after every p steps. A permutation π is *aperiodic* if $p(\pi) = 1$. The chain is *aperiodic* if $p(\pi) = 1$ for every π.

Proposition 3.9 (Aperiodicity) *The swap-only algorithm's Markov chain is aperiodic.*

Proof This follows from irreducibility, which says in particular that for every π, there is an integer m such that $M^m(\pi, \pi) > 0$. Then $M^n(\pi, \pi) > 0$ for all $n > m$, implying $p(\pi) = 1$. □

Proposition 3.10 (Detailed balance) *For all $\pi, \pi' \in \mathcal{S}_N$,*

$$P(\pi)M(\pi, \pi') = P(\pi')M(\pi', \pi). \tag{3.8}$$

Proof The detailed-balance statement in terms of the Gibbs distribution (2.3) and the Metropolis transition matrix (3.4) is

$$\frac{e^{-H(\pi)}}{Z}\left(1 \wedge e^{-H(\pi')}e^{H(\pi)}\right) \stackrel{?}{=} \frac{e^{-H(\pi')}}{Z}\left(1 \wedge e^{-H(\pi)}e^{H(\pi')}\right).$$

The Z's cancel. The lemma below shows that $M(\pi, \pi') \neq 0$ iff $M(\pi', \pi) \neq 0$. If $M(\pi, \pi') = 0$, then detailed balance holds. If $M(\pi, \pi') \neq 0$, then there are two cases. If $H(\pi') \leq H(\pi)$, then

$$e^{-H(\pi)}(1) = e^{-H(\pi')}\left(e^{-H(\pi)}e^{H(\pi')}\right).$$

If $H(\pi') > H(\pi)$,

$$e^{-H(\pi)}\left(e^{-H(\pi')}e^{H(\pi)}\right) = e^{-H(\pi')}(1).$$

In all cases, detailed balance holds. □

Lemma 3.12 *For all $\pi, \pi' \in \mathcal{S}_N$,*

$$M(\pi, \pi') \neq 0 \quad \Longleftrightarrow \quad M(\pi', \pi) \neq 0.$$

Proof This is true since $\pi' \in R(\pi)$ if and only if $\pi \in R(\pi')$, which is a direct consequence of Definition 3.3 of $R(\pi)$. □

This lemma completes the proof that the swap-only algorithm satisfies detailed balance and thus has the Gibbs distribution as its invariant distribution. It is not hard to show that if swaps sites $\mathbf{x} \neq \mathbf{y}$ are in the same cycle before a swap, they are in different cycles after the swap, and vice versa. This is not a correctness result, but rather a sanity check: it shows that cycles may grow or shrink upon swap-only moves.
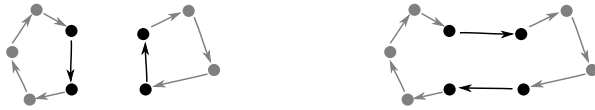
Fig. 6 Swaps merge disjoint cycles and split single cycles. The left-hand permutation can be reached from the right-hand permutation via a swap, and vice versa
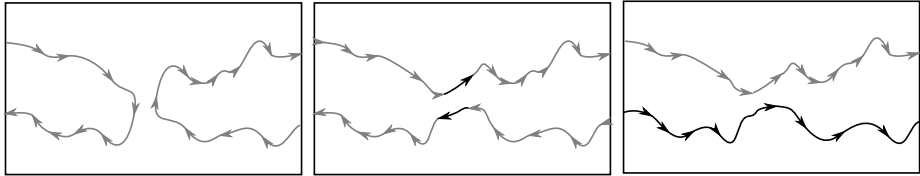
Fig. 7 Conservation of winding number in the swap-only algorithm

3.4 Winding Cycles and the Swap-and-Reverse Algorithm

The propositions of Sect. 3.3 showed that the swap-only algorithm is correct—in particular, any permutation is reachable from any other with non-zero probability. However, in practice some of these non-zero transition probabilities can be quite small. In particular, we observe that the swap-only algorithm almost always generates permutations with zero winding number.

This problem, and a partial solution, is explained intuitively by Fig. 7 and rigorously in [11]. Part 1 of the figure shows a permutation π with a long cycle on the torus which almost meets itself in the x direction. In part 2, after a Metropolis step sending π to π', one cycle winds by $+1$ and the other by -1. Metropolis steps create winding cycles only in opposite-direction pairs; the total $W_x(\pi)$ is still zero. Part 3 of the figure shows that if we reverse one cycle (which is a zero-energy move), $W_x(\pi)$ is now 2. In general (with full details in [11]), winding numbers of even parity can be generated.

Our current best algorithm (swap-and-reverse) has two types of sweeps: (1) For each lattice site, do a Metropolis step as above. (2) For each cycle in the permutation, reverse the direction of the cycle with probability 1/2. This permits winding numbers of even parity in each of the three axes.

We have experimented with various methods to obtain winding numbers of all parities. The creation or destruction of a winding cycle is a non-local update; one is reminded of the Swendsen-Wang algorithm for the Ising model. However, our attempt at non-local updates has an unreasonably low acceptance rate, namely, on the order of e^{-L} where L is the box length.

We have also created a *worm algorithm*, inspired by approaches to this same winding-number problem in path-integral Monte Carlo methods [5, 15]. That is, a permutation loop is selected at random and then cut open at a randomly selected point. The resulting worm is allowed to move around Λ via Metropolis moves; eventually, it closes again. This worm algorithm has an elegant theory and correctness proof [11]; yet, it has an unacceptably long stopping time for loop closure, and none of our attempts to remedy the stopping-time problem have satisfied detailed balance.

At present, we content ourselves with the swap-and-reverse algorithm; it is used to generate all the results discussed in Sect. 4. The order parameters f_S and f_W depend on winding phenomena, but the other three, $1/\xi$, f_I, and $f_{\max}$, do not; furthermore, results obtained in

Sect. 4 using each of the five order parameters are, for the most part, compatible. Yet, as we will see, f_S and f_W do not permit successful finite-size scaling.

3.5 Finite-Size Scaling

Finite-size scaling takes the form of a hypothesis, or rather a set of hypotheses, which is tested against the data. See also [6] for a nice survey.

We have an infinite-volume random variable $S(T)$, e.g. any of the order parameters defined in Sect. 2.3. The finite-volume quantity is $S_L(T)$. Define $t = (T - T_c)/T_c$. Examine, say, $0.99 < t < 1.01$. The first hypothesis is that the correlation length $\xi(T)$ follows a power law

$$\xi(T) \sim |t|^{-\nu}, \quad T \to T_c.$$

For the infinite-volume quantity, we expect a power-law behavior

$$S(T) \sim t^\rho, (-t)^\rho, \quad \text{or} \quad |t|^\rho.$$

(The domain of validity is $t < 0$ or $t > 0$ depending on whether the order parameter S is left-sided or right-sided, respectively.) One moreover hypothesizes that for T near T_c, $S_L(T)$ and $S(T)$ are related by a *universal function* Q_S which depends on T only through the ratio L/ξ:

$$S_L(T) = L^{-\rho/\nu} Q_S(L^{1/\nu} t) \sim L^{-\rho/\nu} Q_S((L/\xi)^{1/\nu}). \tag{3.13}$$

4 Results

Here we complete the steps sketched in Sect. 3.5. The flow of data and uncertainties are as follows:

- Markov chain Monte Carlo simulations, with error bars determined using the method of integrated autocorrelation time [2], yield $S_L(T, \alpha)$ data points. There are five order parameters S, six values of L (30, 40, 50, 60, 70, 80), nine values of α, and a few dozen values of T for each α.
- CPU time per L, T, α experiment, with 10^5 Metropolis sweeps, is approximately 1.3 hours for $L = 40$ and 20 hours for $L = 80$. For the work described in this paper and in [11], a total of 5.4 CPU years was used.
- For each S, L, and α, we use $S_L(T, \alpha)$ values for all available values of T and α to estimate $\hat{\rho}_S(L)$. (Critical exponents are assumed to be independent of α for small α, or with weak enough dependence on α that that dependence is lost in the noise.) Error bars may be propagated from the MCMC simulations, or computed from regression uncertainties.
- Extrapolating $\hat{\rho}_S(L)$ in $L \to \infty$ results in the five estimated critical exponents $\hat{\rho}_S$. Uncertainties are computed from the regression analysis.
- Once the critical exponents are estimated, we obtain $\hat{T}_{c,S}(\alpha)$ for each of the five order parameters S and for each α. Uncertainties are computed by visual inspection of the crossing plots discussed in Sect. 4.3.
- Once the critical exponents and T_c are known, one should be able to obtain plots of the universal function Q_S which is, up to sampling variability, independent of L, T, and α. This verifies the finite-size-scaling hypothesis.
- The shift in reduced critical temperature is as in (2.10). Error bars are computed from regression uncertainties.
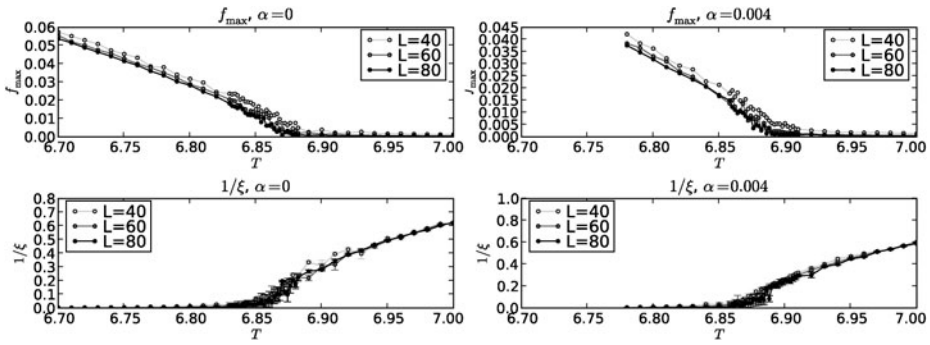
Fig. 8 Order parameters f_M and $1/\xi$ for $L = 40, 60, 80$ and $\alpha = 0$ and 0.004. The remaining order parameters f_S, f_W, and f_I behave similarly to f_M but with not all with the same critical exponents
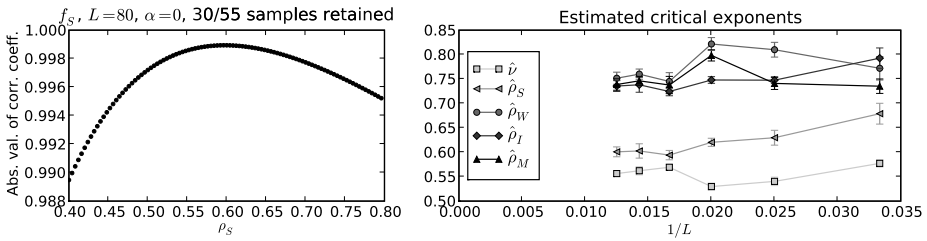
Fig. 9 On the *left*: determination of critical exponent $\hat{\rho}_S(L, \alpha)$ for order parameter f_S, as the value which minimizes linear-regression error for $S_L(T, \alpha)^{1/\rho}$. Visually, one sees $\hat{\rho}_S(L = 80, \alpha = 0.0) \approx 0.59$. On the *right*: estimated critical exponents for $L = 30, 40, 50, 60, 70, 80$

4.1 Determination of L-Dependent Critical Exponents

For each of order parameter S, interaction parameter α, and box length L, we examine all $S(L, T, \alpha)$ data for which $S > \varepsilon$, with ε taken from the plots to ensure that we examine the portions of the curves corresponding to non-zero order parameter in the infinite limit (see Fig. 8). For $1/\xi$, this means $T > T_c$; for the other four order parameters, this means $T < T_c$. From plots such as those in Fig. 8, we choose ε to be 0.1 for $1/\xi$, 0.01 for f_M, 0.01 for f_I, 0.05 for f_S, and 0.01 for f_W. For each S, α, and L, we then apply linear regression to $S(L, T)^{1/\rho_S}$ for varying ρ_S. We find $\hat{\rho}_S(L)$ which optimizes the correlation coefficient [21] of the linear regression. Results are shown in Fig. 9. Given $\hat{\rho}_S(L)$ along with its corresponding linear-regression parameters m and b, we may plot a power-law fit to the simulational data. One such comparison plot is shown in Fig. 10.

4.2 Extrapolation of Critical Exponents for the Infinite-Volume Limit

Next, for each S, given estimates $\hat{\rho}_S(L)$ for increasing values of L, we plot $\hat{\rho}_S(L)$ versus $1/L$. The vertical intercept of this plot estimates the infinite-volume exponent $\hat{\rho}_S(\alpha)$. (See Fig. 9.) Results are shown in Table 2.

4.3 Determination of Critical Temperature

Given the above estimators of the critical exponents, the *crossing method* [6] estimates $T_c(\alpha)$. Namely, we plot $L^{\hat{\rho}/\hat{\nu}} S_L(T)$ as a function of T. At $T = T_c$ we have $t = 0$ and
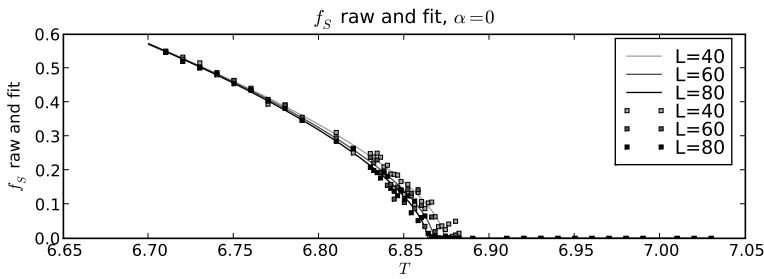
Fig. 10 Power-law fit vs. simulational data for order parameter f_S, $\alpha = 0$

**Table 1** f_M/f_I as a function of α. An upward trend is visible, but it is not pronounced

α	Mean	Std.err.	Count
0.000	0.6242981	0.0000897	78
0.0001	0.6243312	0.0001079	78
0.0002	0.6245691	0.0000921	72
0.0005	0.6245402	0.0001062	66
0.0008	0.6244347	0.0000856	72
0.001	0.6244779	0.0001020	60
0.002	0.6246345	0.0001154	42
0.003	0.6245906	0.0001559	48
0.004	0.6245966	0.0001964	42

Table 2 Extrapolated estimates of the infinite-volume critical exponents, found from the vertical intercept of Fig. 9

$\hat{\nu}$	0.5559	± 0.0037
$\hat{\rho}_S$	0.6201	± 0.0065
$\hat{\rho}_W$	0.7750	± 0.0073
$\hat{\rho}_I$	0.7451	± 0.0052
$\hat{\rho}_M$	0.7486	± 0.0059

$L^{\rho/\nu} S_L(T) = Q_S(0)$, regardless of L (3.13). Thus, these curves will cross (approximately, due to sampling variability) at $T = T_c$. If they do not, the finite-size-scaling hypothesis is not verified. (Note in particular that for order parameter $1/\xi$ whose critical exponent is ν, we apply the crossing method to $LS_L(T)$ as a function of T: thus, the $T_c(\alpha)$ estimate using $1/\xi$ is independent of $\hat{\nu}$.) See for example Fig. 11. (We acknowledge that larger values of L, would improve the visual effect. Results presented here are those obtained within the timeframe of the author's doctoral dissertation work.) Results are in Fig. 13.

　　Using order parameters f_S and f_W, which depend on winding phenomena, one does not see clear crossing behavior. We suggest that either this is related to the even-winding-number issue discussed in Sect. 3.4, or f_S and f_W are not good order parameters for this model. We suspect the former; in every manner except this crossing issue, f_S and f_W behave as expected. (In the absence of clear crossing behavior for f_S and f_W, for the sake of discussion we nonetheless provide best visual estimates for $\hat{T}_c(\alpha)$ for f_S and f_W. These will not be used for further analysis toward our final result.)
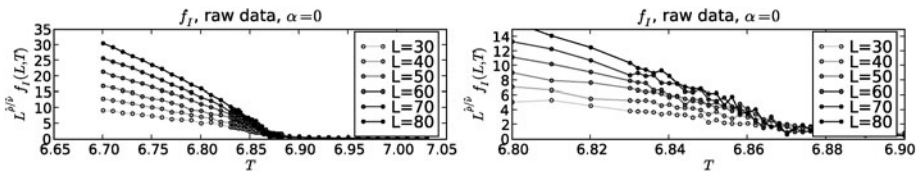
Fig. 11 The crossing method to estimate $T_c(\alpha)$ for order parameter f_I, with $\hat{\rho}$ and $\hat{\nu}$ as above: $T_c(\alpha)$ corresponds to the horizontal coordinate of the intersection point of the plots. The *upper-right-hand* plot is a close-up of the *upper-left-hand* plot. Order parameters f_S and f_W, which depend on winding phenomena, do not exhibit clear crossing behavior
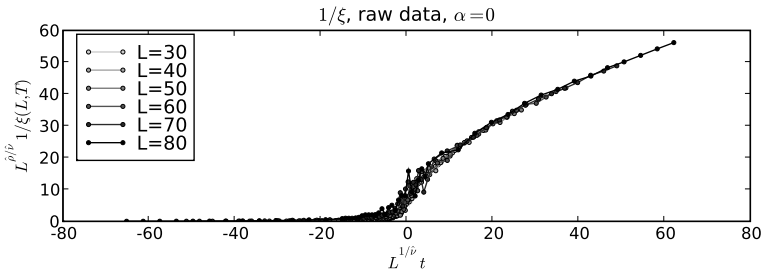
Fig. 12 Collapse plot for order parameter $1/\xi$

4.4 Verification of Finite-Size-Scaling Hypothesis

Now that we have estimated ρ_S, ν, and $T_c(\alpha)$ for each of the five order parameters S, we may plot $L^{\rho_S/\nu} S_L(T, \alpha)$ as a function of $L^{1/\nu} t$. This is a plot of the scaling function Q_S. If the hypothesis is correct, the curves for all L should coincide, or collapse, to within sampling error—which they do (e.g. Fig. 12).

4.5 Determination of the Shift in Critical Temperature

As discussed in Sect. 2.4, we are seeking a linear relationship between $\Delta T_c(\alpha)$ and α, with constant c. This can be visualized in Fig. 14, which is obtained from the $T_{c,S}(\alpha)$ data of Fig. 13 using (2.10). We start with all the $(\alpha, \Delta T_c(\alpha))$ data points from Sect. 4.3. We omit values obtained using f_S and f_W, due to the aforementioned lack of crossing behavior. We also omit values obtained using $\alpha = 0.004$, since the critical-temperature plots of Fig. 13 suggests that this starts to exceed the domain of linear approximation. We perform a linear regression with error bars [21] on the $(\alpha, \Delta T_c(\alpha))$ data points. We use a slope-only fit, rather than a slope-intercept fit, since $\Delta T_c(\alpha)$ has zero intercept by its definition. We find

$$c = 0.618 \pm 0.086 \quad (2\sigma \text{ error bar}).$$

Within experimental uncertainty, this result, for points on the lattice with Ewens cycle-weights, matches the c value of (2.11) for point positions varying on the continuum with decaying-cycle-weight interactions.
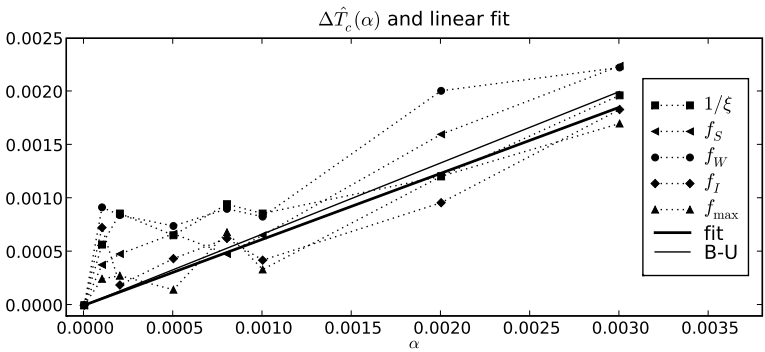
4.6 Constancy of the Macroscopic-Cycle Quotient

As discussed in Sect. 2.4, we hypothesize that the macroscopic-cycle quotient f_M/f_I in the infinite-volume limit is dependent on α but is constant in T where it is defined, i.e. for
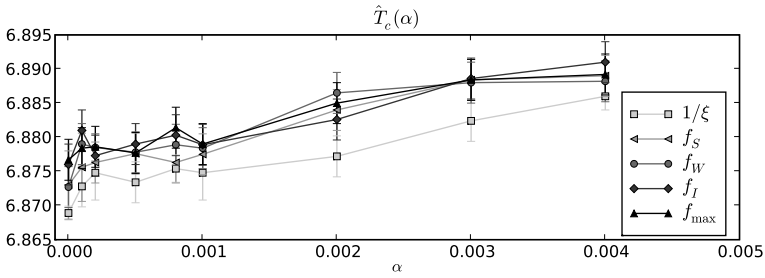
**Fig. 13** Critical temperature as function of α

Fig. 14 Shift in critical temperature, and linear fit, as function of α. Recall from (2.10) that $\Delta T_c(\alpha) = \frac{T_c(\alpha) - T_c(0)}{T_c(0)}$. Order parameters f_S and f_W were omitted from the fit, due to lack of crossing behavior; $\alpha = 0.004$ was omitted due to onset of curvature of $T_c(\alpha)$. The *heavy solid line* shows a linear fit with empirically determined constant of proportionality; the *lighter solid line* is the comparison value of Betz and Ueltschi (slope 2/3) for decaying cycle weights and continuum point positions

$T < T_c$ since $f_I = 0$ for $T > T_c$. This may be visualized by comparing figures such as 3: one sees that f_M and f_I appear to have the same critical exponent. Alternatively, one may plot the ratio f_M/f_I (Fig. 15). In the infinite-volume limit, f_I is zero for $T > T_c$ and so we are interested only in the values of the quotient for $T < T_c$. In that region, the quotient does indeed appear to be constant in T.

We test this constancy hypothesis as follows. The respective critical exponents are ρ_M and ρ_I. The estimators are $\hat{\rho}_M$ and $\hat{\rho}_I$, computed by averaging over several different values of L and α as described in Sect. 4.2. Treating these estimators as normally distributed (as justified by the raw data), we obtain the standard deviations of the $\hat{\rho}_{M,I}(L, \alpha)$ samples, along with the standard deviations of the means $\hat{\rho}_{M,I}$:

$$\hat{\rho}_M = 0.7482 \qquad \hat{\rho}_I = 0.7445$$

$$s_M = 0.0428 \qquad s_I = 0.0374$$

$$n_M = 50 \qquad n_I = 50$$

$$s_M/\sqrt{n_M} = 0.006059 \qquad s_I/\sqrt{n_I} = 0.005295.$$

The difference $\hat{\rho}_M - \hat{\rho}_I$ is also normally distributed about the true mean $\rho_M - \rho_I$, but $\hat{\rho}_M$ and $\hat{\rho}_I$ are not independent since they are sample means of random variables computed
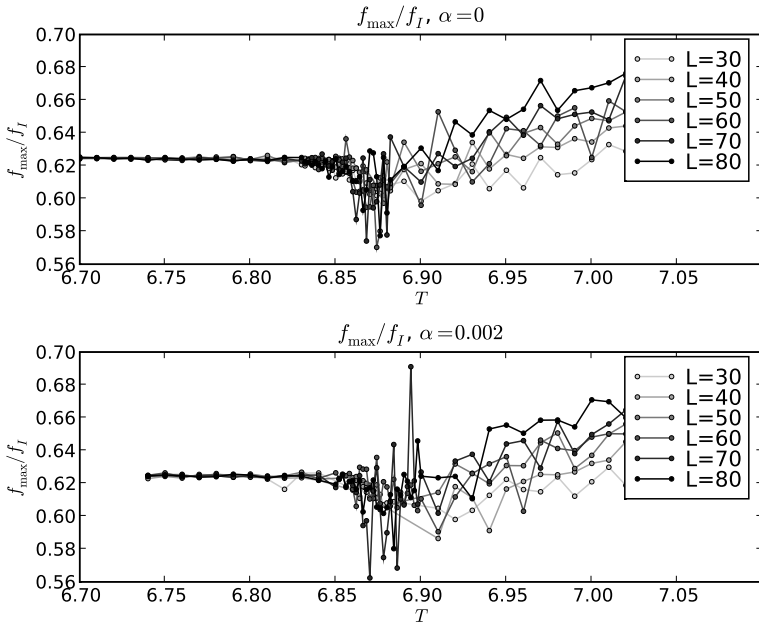
Fig. 15 Macroscopic-cycle quotient f_M/f_I for $\alpha = 0, 0.002$

from the same Markov chain Monte Carlo sequence of permutations. Thus

$$\text{Var}(\hat{\rho}_M - \hat{\rho}_I) = \text{Var}(\hat{\rho}_M) + \text{Var}(\hat{\rho}_I) - 2\text{Cov}(\hat{\rho}_M, \hat{\rho}_I).$$

Computing the sample covariance of the $\hat{\rho}_M(L, \alpha)$ and $\hat{\rho}_I(L, \alpha)$ data series, we obtain the covariance and resulting standard error s_d of the difference

$$\text{Cov}(\hat{\rho}_M, \hat{\rho}_I) = 0.0004 \qquad s_d/\sqrt{n} = 0.0070.$$

Normalizing, we find

$$\hat{\rho}_M - \hat{\rho}_I = 0.0037 \qquad \frac{\hat{\rho}_M - \hat{\rho}_I}{s_d/\sqrt{n}} = \frac{0.0037}{0.0070} = 0.5293.$$

We hypothesize $\rho_M - \rho_I = 0$; the estimated value $\hat{\rho}_M - \hat{\rho}_I$ lies comfortably within a standard deviation of this. We note, moreover, that the value of f_M/f_I, while constant in T, trends upward with α (see Table 1 and Fig. 16). This merits further investigation.

4.7 Conclusions

(1) For annealed point positions, (2.13) gives $T_c(0) \approx 6.625$. Our result $T_c(0) = 6.873 \pm 0.006$ (2σ error bar) unambiguously shows that the lattice structure modifies the critical temperature, even in the non-interacting ($\alpha = 0$) case.

(2) As detailed in Sect. 4.5, we find that the reduced shift in critical temperature as a function of interaction parameter α is

$$\Delta T_c(\alpha) \approx \frac{T_c(\alpha) - T_c(0)}{T_c(0)} = c\alpha$$
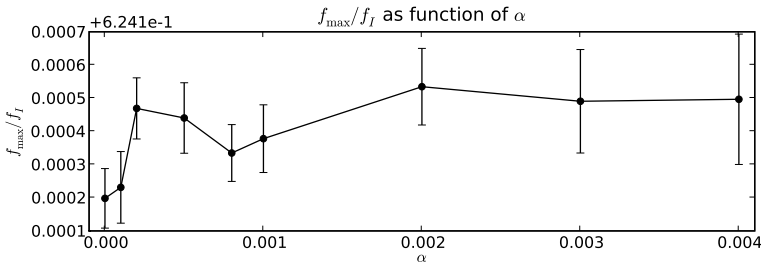
Fig. 16 f_M/f_I as a function of α

with

$$c = 0.618 \pm 0.086 \quad (2\sigma \text{ error bar}).$$

This is compatible (Sect. 2.4) with the related result of [4]. Even though the lattice structure changes the critical temperature (conclusion 1), the *shift* in critical temperature is unaffected.

(3) As described in Sect. 2.4, Shepp and Lloyd [17] find that $\mathbb{E}[\ell_{\max}]/N \approx 0.6243$ for uniform-random (non-spatial) permutations. For spatial permutations, we define a macroscopic-cycle quotient $\mathbb{E}[\ell_{\max}]/Nf_I$ which is the ratio of mean maximum cycle length as a fraction of the number of sites in long cycles. Our result (Table 1) is compatible with that of Shepp and Lloyd for the non-interacting case, with an increase which appears to be linear as a function of interaction parameter α. Our result is also compatible with [9], which (among other conclusions) recovered the Shepp and Lloyd result for the $\alpha = 0$ case.

5 Future Work

Now that the α-dependence of the macroscopic-cycle quotient's constant upon α has been found empirically, one would next like to explain that dependence analytically.

Ideally, one would have an algorithm to permit odd winding numbers, as discussed in Sect. 3.4.

Sampling from the true Bose-gas distribution using the random-cycle model requires three changes. First, one needs to conduct simulations using the Bose-gas interaction (2.1) rather than the cycle-weight interaction (2.2). The interaction term V is a CPU-intensive Brownian-bridge computation [3]; unpublished work of Ueltschi and Betz shows that it may be approximated in the weak-interaction case by a simpler Riemann integral. Precomputed tables and interpolation may make use of this integral feasible. Second, point positions must be allowed to vary on the continuum. This entails a second type of Metropolis step, in addition to that shown in Sect. 3.1. Third, since points are no longer held fixed on the lattice, it is no longer trivial to find nearest neighbors. Software efficiency requires a hierarchical partitioning of Λ. The second and third points simply require a software effort. Implementing them will be worthwhile only if the interaction terms can be simplified to the point that they are computationally feasible, which is a mathematical effort.

References

1. Baym, G., Blaizot, J.-P., Holzmann, M., Laloë, F., Vautherin, D.: Bose-Einstein transition in a dilute interacting gas. Eur. Phys. J. B **24**, 107–124 (2001). arXiv:cond-mat/0107129v2
2. Berg, B.: Markov Chain Monte Carlo Simulations and Their Statistical Analysis. World Scientific, Singapore (2004)
3. Betz, V., Ueltschi, D.: Spatial random permutations and infinite cycles. Commun. Math. Phys. **285**, 469–501 (2009). arXiv:0711.1188
4. Betz, V., Ueltschi, D.: Spatial random permutations with small cycle weights. Probab. Theory Relat. Fields (2010). arXiv:0812.0569v1
5. Boninsegni, M., Prokof'ev, N.V., Svistunov, B.V.: Worm algorithm and diagrammatic Monte Carlo: A new approach to continuous-space path integral Monte Carlo simulations. Phys. Rev. E **74**, 036701 (2006)
6. Caracciolo, S., Gambassi, A., Gubinelli, M., Pelisetto, A.: Finite-size scaling in the driven lattice gas. J. Stat. Phys. **115**(1/2) (2004). arXiv:cond-mat/0312175
7. Ewens, W.J.: The sampling theory of selectively neutral alleles. Theor. Popul. Biol. **3**, 87–112 (1972)
8. Feynman, R.P.: Atomic theory of the λ transition in helium. Phys. Rev. **91**(6) (1953)
9. Gandolfo, D., Ruiz, J., Ueltschi, D.: On a model of random cycles. Stat. Phys. **129**, 663–676 (2007). arXiv:cond-mat/0703315
10. Golomb, S.W.: Random permutations. Bull. Am. Math. Soc. **70**, 747 (1964)
11. Kerl, J.: Critical behavior for the model of random spatial permutations. Doctoral dissertation, University of Arizona (2010)
12. Landau, D.P., Binder, K.: A Guide to Monte Carlo Simulations in Statistical Physics, 2nd edn. Cambridge University Press, Cambridge (2005)
13. Penrose, O., Onsager, L.: Bose-Einstein condensation and liquid helium. Phys. Rev. **104**(3) (1956)
14. Pollock, E.L., Ceperley, D.M.: Path-integral computation of superfluid densities. Phys. Rev. B **36**(16) (1987)
15. Prokof'ev, N.V., Svistunov, B.V., Tupitsyn, I.S.: Exact, complete, and universal continuous-time world-line Monte Carlo approach to the statistics of discrete quantum systems. J. Exp. Theor. Phys. **87**(2) (1998)
16. Seiringer, R., Ueltschi, D.: Rigorous upper bound on the critical temperature of dilute Bose gases. Phys. Rev. B **80**, 014502 (2009). arXiv:0904.0050
17. Shepp, L.A., Lloyd, S.P.: Ordered cycle length in a random permutation. Trans. Am. Math. Soc. **121**, 340–357 (1966)
18. Sütő, A.: Percolation transition in the Bose gas. J. Phys. A, Math. Gen. **26**, 4689–4710 (1993)
19. Sütő, A.: Percolation transition in the Bose gas II. J. Phys. A, Math. Gen. **35**, 6995–7002 (2002)
20. Ueltschi, D.: The model of interacting spatial permutations and its relation to the Bose gas. In: Mathematical Results in Quantum Mechanics, pp. 225–272. World Scientific, Singapore (2008). arXiv:0712.2443v3
21. Young, H.D.: Statistical Treatment of Experimental Data. McGraw-Hill, New York (1962)