

Shift in critical temperature for random spatial permutations with cycle weights

John Kerl

December 21, 2009

Abstract

We elaborate on a model of random spatial permutations, wherein permutations are weighted according to point positions. This model originates in a study of the interacting Bose gas; the low-temperature-dependent onset of the appearance of arbitrarily long cycles is connected to the phase transition of Bose-Einstein condensates. For our work, we consider a simplified model with point positions held fixed on the cubic lattice, with interactions 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, we find $c = 0.665 \pm 0.067$. 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.

1 Introduction

The random-spatial-permutation model arises in the study of the Bose gas. Its history includes Bose-Einstein, Feynman [Feynman], Penrose-Onsager [PO], Sütő [Sütő1, Sütő2], and Betz-Ueltschi-Gandolfo-Ruiz [GRU, U07, BU07, BU08]. 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 [BU07, BU08]. Specifically, given points $\mathbf{x}_1, \dots, \mathbf{x}_N$ in the box $[0, L]^3$ and temperature T , permutations π are given weights proportional to the Gibbs factor $e^{-H(\pi)}$ where

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

for the non-interacting case. (The sum 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

$$\sum_{i < j} V(\mathbf{x}_i, \mathbf{x}_{\pi(i)}, \mathbf{x}_j, \mathbf{x}_{\pi(j)}), \quad (1.2)$$

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

$$\sum_{\ell=1}^N \alpha_{\ell} r_{\ell}(\pi) \tag{1.3}$$

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 and Einstein 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 [BBHLV] for non-interacting atoms of the same density. For weakly interacting systems, however, an emerging consensus is that interactions *increase* the critical temperature. See [BBHLV, SU09] 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

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

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 (equation 1.2) 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 [GRU], a simulational approach is taken for points held fixed on the cubic unit lattice in the non-interacting case. In the papers [U07, BU07], 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 1.3 with $\alpha_2 = \alpha$, where α is related to a hard-core scattering length, and the remaining cycle weights are zero. In [BU08], 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 lattice, with cycle weights constant in ℓ — removing the small-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 subsection 2.1; qualitative and quantitative behavior of long cycles are discussed in subsections 2.2 and 2.3, respectively. Known results and conjectures are listed in subsection 2.4. In section 3, the simulational methods are presented. The *swap-only* and *swap-and-reverse* algorithms generate simulational data; these algorithms are proved correct in subsections 3.1 through 3.3. The finite-size-scaling method, which reduces the raw simulational data, is summarized in subsection 3.5. Section 4 presents the data and its analysis in full detail: estimation of critical exponents and critical temperature in subsections 4.1 through 4.3, verification of the finite-size-scaling hypothesis in subsection 4.4, and final results in sections 4.5 through 4.7.

2 The model of random spatial permutations

Here we review concepts from [BU07, BU08], fixing notation and intuition to be used in the rest of the paper.




Figure 1: A spatial permutation.

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 $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ for $\mathbf{x}_1, \dots, \mathbf{x}_N \in \Lambda$. The Hamiltonian takes one of two forms. In the first, relevant to the Bose gas, we have

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

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:

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

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* [GRU] has $\alpha_\ell \equiv 0$. The *two-cycle model* [BU07, U07], has $\alpha_2 = \alpha$ and other cycle weights are zero. The *general-cycle model* has no restrictions on α_ℓ . In [BU08], the *small 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 [Ewens]), is another special case of the general-cycle model: it has $\alpha_\ell \equiv \alpha$ constant in ℓ .

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

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

(Alternatively, one may integrate over all positions in Λ , with a resulting Gibbs distribution on \mathcal{S}_N . Here, several analytical results are available [BU07, BU08].) For a random variable $S(\pi)$, we have

$$\mathbb{E}_{\Lambda, \mathbf{X}}[S] = \frac{1}{Y(\Lambda, \mathbf{X})} \sum_{\pi \in \mathcal{S}_N} S(\pi) e^{-H(\mathbf{X}, \pi)}, \quad \text{i.e.} \quad \mathbb{E}_\pi[S(\pi)] = \sum_{\pi \in \mathcal{S}_N} P(\pi) S(\pi). \quad (2.4)$$

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.



Figure 2: Some typical permutations for high T , medium but subcritical T , and low T .

As $T \rightarrow \infty$, 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 $T \rightarrow 0$, 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 $\|\pi(\mathbf{x}) - \mathbf{x}\|$ of each permutation jump remains small, increasing smoothly as T drops.

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 ℓ_{\max} be the length of the longest cycle in π , with $\mathbb{E}[\ell_{\max}]$ its mean over all permutations. We observe that for $T > T_c$, $\mathbb{E}[\ell_{\max}]$ is constant as $N \rightarrow \infty$: cycles remain finite. That is to say, $\mathbb{E}[\ell_{\max}]/N$ goes to zero as $N \rightarrow \infty$. 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 $N \rightarrow \infty$: there are arbitrarily long cycles, or infinite cycles, in the infinite-volume limit. See figure 2 for depictions of typical permutations at high T , subcritical T , and low T ; see figure 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 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*

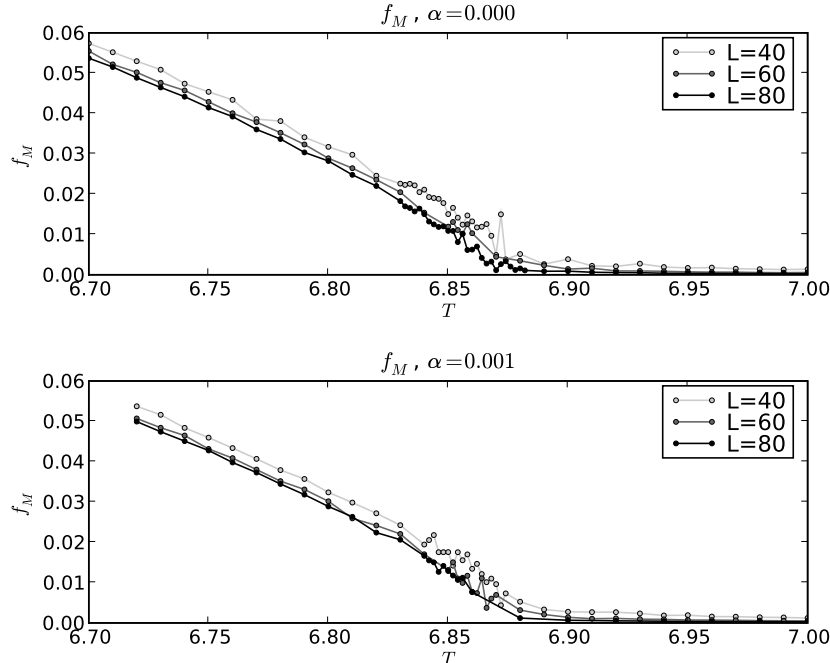


Figure 3: Order parameter $f_M = \mathbb{E}[\ell_{\max}]/N$ for finite systems, with $\alpha = 0.000, 0.001$. Interactions increase the critical temperature.

in long cycles, f_I , is described in detail in [GRU]. 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

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

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 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

$$\mathbf{W} = (W_x, W_y, W_z) = \sum_{i=1}^N \tilde{d}(\mathbf{x}_{\pi(i)} - \mathbf{x}_i). \quad (2.5)$$

We also write

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

Equation 2.5 is more subtle than it first appears. We work in the regime where single permutation jumps are on the order of 1, 2, 3, or so — specifically, jumps are of length less than $L/2$. The distances \tilde{d} in equation 2.5 are to be interpreted in terms of lifts to a separate Euclidean chart for each site \mathbf{x}_i . The scaled winding number [PC87] 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 $f_M := \mathbb{E}[\ell_{\max}]/N$ (figure 3): asymptotically as $N \rightarrow \infty$, they are zero to the left of T_c and non-zero to the right. For finite N , the curves remain analytic: finite-size effects persist. The inverse correlation length $1/\xi(T)$, on the other hand, is zero to the left of T_c and non-zero to the right.

Our goal is to quantify the dependence of T_c on α , where

$$\Delta T_c(\alpha) = \frac{T_c(\alpha) - T_c(0)}{T_c(0)}. \quad (2.6)$$

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

2.4 Known results and conjectures

Known results for point locations averaged over the continuum are obtained largely using Fourier methods [BU08], 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 [BU07] and small cycle weights [BU08]. (This taps into a long and controversial history in the physics literature: see [BBHLV] or [SU09] for surveys.) The critical (ρ, T, α) manifold relates ρ_c to T_c . Specifically,

$$\rho_c(\alpha) \approx \sum_{\ell \geq 1} e^{-\alpha \ell} \int_{\mathbb{R}^3} e^{-\ell 4\pi^2 \beta \|\mathbf{k}\|^2} d\mathbf{k} = \frac{1}{(4\pi\beta)^{3/2}} \sum_{\ell \geq 1} e^{-\alpha \ell} \ell^{-3/2} \quad (2.7)$$

$$\Delta T_c(\alpha) \approx c\rho^{1/3}\alpha, \quad \text{for } \alpha \approx 0. \quad (2.8)$$

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

$$\rho_c \approx \frac{\zeta(3/2)e^{-\alpha}}{(4\pi\beta)^{3/2}}, \quad T_c \approx \frac{4\pi e^{2\alpha/3}}{\zeta(3/2)^{2/3}} \approx 6.626e^{2\alpha/3}, \quad c \approx 0.667. \quad (2.9)$$

We inquire whether this result, obtained for small 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), $\mathbb{E}[\ell_{\max}]/N f_I$ 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 [SL] solved Golomb’s 1964 question [Golomb]), $\mathbb{E}[\ell_{\max}]/N \approx 0.6243$; unpublished work of Betz and Ueltschi has found $\mathbb{E}[\ell_{\max}]/N f_I$ 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. We conjecture that $\mathbb{E}[\ell_{\max}]/N f_I$ 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 small 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 π_1, \dots, π_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.

3.1 The swap-only algorithm

Recall from equation 2.4 that the expectation of a random variable S (such as ξ , f_M , f_W , f_I , f_S) is

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

The number of permutations, $N!$, grows intractably in N . As is typical in Markov chain Monte Carlo methods [Berg, LB], 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.




Figure 4: Metropolis moves for the swap-only algorithm.

The *swap-only algorithm* for sampling from the Gibbs distribution (equation 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 figure 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 [Berg] 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 [Kerl]. We next prove correctness of this algorithm.

3.2 Explicit construction of the Markov matrix

For section 3.3 we will need an explicit construction of the Markov matrix corresponding to the swap-only algorithm as described in section 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 section 3.3 we show that $\mathbf{P}^{(k)}$ approaches the Gibbs distribution (equation 2.3).

The matrix \mathbf{M} is $N! \times N!$: rows are indexed by $\pi_1, \dots, \pi_{N!}$ and columns are indexed by $\pi'_1, \dots, \pi'_{N!}$. 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

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

Remark. Note that $d(\pi, \pi') \neq 1$ since if two permutations agree on $N - 1$ sites, they must agree on the remaining site.

Lemma 3.2. *The function $d(\pi, \pi')$ is a metric on \mathcal{S}_N .*

Proof. Symmetry is obvious, as is non-negativity. For positive definiteness, note that $d(\pi, \pi') = 0$ iff $\pi = \pi'$. For the triangle inequality, let $\pi, \pi', \pi'' \in \mathcal{S}_N$. Partition the set $\{1, 2, \dots, N\}$ into the four disjoint sets

$$\begin{aligned} A &= \{i = 1, 2, \dots, N : \pi(i) = \pi'(i), \pi'(i) = \pi''(i)\}, \\ B &= \{i = 1, 2, \dots, N : \pi(i) = \pi'(i), \pi'(i) \neq \pi''(i)\}, \\ C &= \{i = 1, 2, \dots, N : \pi(i) \neq \pi'(i), \pi'(i) = \pi''(i)\}, \\ D &= \{i = 1, 2, \dots, N : \pi(i) \neq \pi'(i), \pi'(i) \neq \pi''(i)\}. \end{aligned}$$

Then $\pi = \pi''$ on all of A ; $\pi \neq \pi''$ on all of B and C ; and π, π'' may or may not agree on various elements of D :

A	B	C	D
$\pi = \pi'$	$\pi = \pi'$	$\pi \neq \pi'$	$\pi \neq \pi'$
$\pi' = \pi''$	$\pi' \neq \pi''$	$\pi' = \pi''$	$\pi' \neq \pi''$
$\pi = \pi''$	$\pi \neq \pi''$	$\pi' \neq \pi''$	Varies

That is,

$$\begin{aligned} d(\pi, \pi') &= \#C + \#D, \\ d(\pi', \pi'') &= \#B + \#D, \\ \#B + \#C &\leq d(\pi, \pi'') \leq \#B + \#C + \#D. \end{aligned}$$

Then

$$d(\pi, \pi'') \leq \#B + \#C + \#D \leq \#B + \#C + 2\#D = d(\pi, \pi') + d(\pi', \pi'').$$

□

Definition 3.3. A swap on sites \mathbf{x} and \mathbf{y} is *trivial* if $\mathbf{x} = \mathbf{y}$.

Definition 3.4. Lattice sites \mathbf{x}, \mathbf{y} are *nearest-neighbor* if $\|\mathbf{x} - \mathbf{y}\| = 1$.

Definition 3.5. For $\pi \in \mathcal{S}_N$, define

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

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 also write

$$\pi' \circ\text{-}\circ \pi$$

if $\pi' \in R(\pi)$ or $\pi' = \pi$.

The Metropolis steps are then described as follows. For each $\pi \in \mathcal{S}_N$,

$$M(\pi, \pi') = \begin{cases} C \left(\mathbf{1} \wedge e^{-H(\pi') + H(\pi)} \right), & \pi' \in R(\pi), \\ 1 - \sum_{\pi' \in R(\pi)} C \left(\mathbf{1} \wedge e^{-H(\pi') + H(\pi)} \right), & \pi = \pi'; \\ 0, & \text{otherwise.} \end{cases}$$

To find out what the normalization factor C is, we observe the following:

- 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} .

We require rows sums to be 1 for stochastic matrices so we take

$$C = \frac{1}{3N}.$$

This prefactor of $1/3N$ is the same for all non-zero entries of \mathbf{M} for Metropolis. Then

$$M(\pi, \pi') = \begin{cases} \frac{1}{3N} \left(\mathbf{1} \wedge e^{-H(\pi') + H(\pi)} \right), & \pi' \in R(\pi), \\ 1 - \sum_{\pi' \in R(\pi)} \frac{1}{3N} \left(\mathbf{1} \wedge e^{-H(\pi') + H(\pi)} \right), & \pi = \pi'; \\ 0, & \text{otherwise.} \end{cases} \quad (3.6)$$

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 (equation 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 3.7 (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.*

Proof. Transpositions generate \mathcal{S}_N : for all $\pi \in \mathcal{S}_N$, there exist transpositions $\sigma_1, \dots, \sigma_m$ 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 $G_{\mathbf{a}, \mathbf{b}} : \mathcal{S}_N \rightarrow \mathcal{S}_N$ 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 $\mathbf{y}_0, \mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n$ such that $\mathbf{y}_0 = \mathbf{x}$, $\mathbf{y}_n = \mathbf{z}$, and $\|\pi(\mathbf{y}_{i+1}) - \pi(\mathbf{y}_i)\| = 1$ for $i = 0, 1, \dots, n-1$. (See figure 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.




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

Notation 3.8. Given $\mathbf{x}_1, \dots, \mathbf{x}_N$ 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:

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

Lemma 3.9. *The composition $G_{\mathbf{a}, \mathbf{c}} \circ G_{\mathbf{a}, \mathbf{b}}$ behaves as follows:*

$$\pi = \begin{pmatrix} \dots & \mathbf{a} & \mathbf{b} & \mathbf{c} & \dots \\ \dots & \pi(\mathbf{a}) & \pi(\mathbf{b}) & \pi(\mathbf{c}) & \dots \end{pmatrix} \mapsto \pi' = \begin{pmatrix} \dots & \mathbf{a} & \mathbf{b} & \mathbf{c} & \dots \\ \dots & \pi(\mathbf{c}) & \pi(\mathbf{a}) & \pi(\mathbf{b}) & \dots \end{pmatrix}$$

Proof. The first map, $G_{\mathbf{a}, \mathbf{b}}$, does

$$\begin{pmatrix} \dots & \mathbf{a} & \mathbf{b} & \mathbf{c} & \dots \\ \dots & \pi(\mathbf{a}) & \pi(\mathbf{b}) & \pi(\mathbf{c}) & \dots \end{pmatrix} \mapsto \begin{pmatrix} \dots & \mathbf{a} & \mathbf{b} & \mathbf{c} & \dots \\ \dots & \pi(\mathbf{b}) & \pi(\mathbf{a}) & \pi(\mathbf{c}) & \dots \end{pmatrix};$$

$G_{\mathbf{a}, \mathbf{c}}$ sends this to

$$\begin{pmatrix} \dots & \mathbf{a} & \mathbf{b} & \mathbf{c} & \dots \\ \dots & \pi(\mathbf{c}) & \pi(\mathbf{a}) & \pi(\mathbf{b}) & \dots \end{pmatrix}.$$

□

Corollary 3.10. *The composition $G_{\mathbf{y}_0, \mathbf{y}_n} \circ G_{\mathbf{y}_0, \mathbf{y}_{n-1}} \circ \dots \circ G_{\mathbf{y}_0, \mathbf{y}_2} \circ G_{\mathbf{y}_0, \mathbf{y}_1}$, sending $\pi \mapsto \pi'$, performs the right cyclic shift on images of $\mathbf{y}_0, \dots, \mathbf{y}_n$ given by*

$$\begin{pmatrix} \mathbf{y}_0 & \mathbf{y}_1 & \dots & \mathbf{y}_{n-1} & \mathbf{y}_n \\ \pi(\mathbf{y}_0) & \pi(\mathbf{y}_1) & \dots & \pi(\mathbf{y}_{n-1}) & \pi(\mathbf{y}_n) \end{pmatrix} \mapsto \begin{pmatrix} \mathbf{y}_0 & \mathbf{y}_1 & \dots & \mathbf{y}_{n-1} & \mathbf{y}_n \\ \pi(\mathbf{y}_n) & \pi(\mathbf{y}_0) & \dots & \pi(\mathbf{y}_{n-2}) & \pi(\mathbf{y}_{n-1}) \end{pmatrix}.$$

Likewise, $G_{\mathbf{y}_n, \mathbf{y}_1} \circ G_{\mathbf{y}_n, \mathbf{y}_2} \circ \dots \circ G_{\mathbf{y}_n, \mathbf{y}_{n-2}} \circ G_{\mathbf{y}_n, \mathbf{y}_{n-1}}$ leaves the image of \mathbf{y}_0 unchanged and performs the left cyclic shift on images of $\mathbf{y}_1, \dots, \mathbf{y}_n$ given by

$$\begin{pmatrix} \mathbf{y}_0 & \mathbf{y}_1 & \dots & \mathbf{y}_{n-1} & \mathbf{y}_n \\ \pi(\mathbf{y}_0) & \pi(\mathbf{y}_1) & \dots & \pi(\mathbf{y}_{n-1}) & \pi(\mathbf{y}_n) \end{pmatrix} \mapsto \begin{pmatrix} \mathbf{y}_0 & \mathbf{y}_1 & \dots & \mathbf{y}_{n-1} & \mathbf{y}_n \\ \pi(\mathbf{y}_0) & \pi(\mathbf{y}_2) & \dots & \pi(\mathbf{y}_n) & \pi(\mathbf{y}_1) \end{pmatrix}.$$

Proof. These follow from the lemma by induction on n . □

Composing these two maps, we find that

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

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

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

□

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 3.11. A permutation π has period p if any return to π must occur at multiples of p steps. Precisely, let Π_n be the random variable which is the n th value of the Markov chain. Then

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

Furthermore, the chain is said to be aperiodic if all states are aperiodic.

Proposition 3.12 (Aperiodicity). *The swap-only algorithm's Markov chain is aperiodic.*

Proof. From Markov-chain theory, we know that all states in a recurrence class have the same period. Thus, it suffices to show that one state, e.g. the identity, is aperiodic. If there is non-zero probability of transitioning from the identity to itself, the identity is an aperiodic state. But this is in fact the case: The identity has zero energy (equations 2.2). Any state proposed as a transition from the identity is necessarily a two-cycle, for which the distance terms in equation 2.2 are positive. The cycle-weight terms or interaction terms are non-negative. Thus $\Delta H > 0$, and the Metropolis transition probability is strictly less than 1. Therefore there is a non-zero probability of staying at the identity permutation. □

Proposition 3.13 (Detailed balance). *For all $\pi, \pi' \in \mathcal{S}_N$,*

$$P(\pi)M(\pi, \pi') = P(\pi')M(\pi', \pi). \tag{3.14}$$

Proof. The detailed-balance statement in terms of the Gibbs distribution (equation 2.3) and the Metropolis transition matrix (equation 3.6) is

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

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

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

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

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

In all cases, detailed balance holds. □

Lemma 3.15. For all $\pi, \pi' \in \mathcal{S}_N$,

$$M(\pi, \pi') \neq 0 \iff M(\pi', \pi) \neq 0.$$

Proof. If $d(\pi, \pi') > 2$, then $M(\pi, \pi') = 0$ and $M(\pi', \pi) = 0$. If $d(\pi, \pi') = 2$, then π and π' differ at two sites \mathbf{x} and \mathbf{y} . Put

$$\begin{aligned} \pi(\mathbf{x}) &= \mathbf{u}, & \pi'(\mathbf{x}) &= \mathbf{v}, \\ \pi(\mathbf{y}) &= \mathbf{v}, & \pi'(\mathbf{y}) &= \mathbf{u}. \end{aligned}$$

If $\|\mathbf{u} - \mathbf{v}\| = 1$ then $M(\pi, \pi') \neq 0$ and $M(\pi', \pi) \neq 0$ since \mathbf{u} and \mathbf{v} are nearest-neighbor lattice sites. If $\|\mathbf{u} - \mathbf{v}\| > 1$ then $M(\pi, \pi') = 0$ and $M(\pi', \pi) = 0$ since \mathbf{u} and \mathbf{v} are not nearest-neighbor lattice sites. Last, if $\pi = \pi'$, then $M(\pi, \pi') = M(\pi', \pi)$. □

This lemma completes the proof that the swap-only algorithm satisfies detailed balance and thus has the Gibbs distribution as its invariant distribution. The following proposition is not a correctness result, but rather a sanity check. It shows that cycles may grow or shrink upon swap-only moves.



Figure 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.

Definition 3.16. If \mathbf{x} and \mathbf{y} are in a common cycle of π , we say that \mathbf{x} is **connected to** \mathbf{y} ; otherwise we say that \mathbf{x} is not connected to \mathbf{y} . These are written

$$\mathbf{x} \circ\text{-}\circ\mathbf{y} \quad \text{and} \quad \mathbf{x} \circ\not\text{-}\circ\mathbf{y}.$$

In the former case, we write

$$\ell_{\mathbf{x}, \mathbf{y}}(\pi)$$

for the smallest positive integer a such that $\pi^a(\mathbf{x}) = \mathbf{y}$. This is the number of permutation jumps from \mathbf{x} to \mathbf{y} .

Proposition 3.17. If \mathbf{x} and \mathbf{y} are in disjoint cycles before a non-trivial swap then they are in the same cycle afterward and vice versa (see figure 6).

Proof. First suppose that \mathbf{x} and \mathbf{y} are in disjoint cycles. Let the respective cycle lengths be $\ell(\mathbf{x}) = a$ and $\ell(\mathbf{y}) = b$. Recall that $1 \leq a, b \leq N$. Those cycles are

$$\mathbf{x} \mapsto \pi(\mathbf{x}) \mapsto \pi^2(\mathbf{x}) \mapsto \dots \mapsto \pi^{a-1}(\mathbf{x}) \mapsto \mathbf{x} \quad \text{and} \quad \mathbf{y} \mapsto \pi(\mathbf{y}) \mapsto \pi^2(\mathbf{y}) \mapsto \dots \mapsto \pi^{b-1}(\mathbf{y}) \mapsto \mathbf{y}.$$

Since these are disjoint cycles, all elements listed are distinct lattice sites. After the swap, we have

$$\mathbf{y} \mapsto \pi(\mathbf{x}) \mapsto \pi^2(\mathbf{x}) \mapsto \dots \mapsto \pi^{a-1}(\mathbf{x}) \mapsto \mathbf{x} \quad \text{and} \quad \mathbf{x} \mapsto \pi(\mathbf{y}) \mapsto \pi^2(\mathbf{y}) \mapsto \dots \mapsto \pi^{b-1}(\mathbf{y}) \mapsto \mathbf{y}.$$

This is a single cycle of length $a + b$, starting with \mathbf{y} , including \mathbf{x} , and returning to \mathbf{y} .

Second, suppose that \mathbf{x} and \mathbf{y} are in the same cycle. Let $a = \ell_{\mathbf{x},\mathbf{y}}(\pi)$ and $b = \ell_{\mathbf{y},\mathbf{x}}(\pi)$ (see definition 3.16). (These numbers are both positive since the swap is non-trivial, i.e. $\mathbf{x} \neq \mathbf{y}$.) Then we have

$$\mathbf{x} \mapsto \pi(\mathbf{x}) \mapsto \pi^2(\mathbf{x}) \mapsto \dots \mapsto \pi^{a-1}(\mathbf{x}) \mapsto \mathbf{y} \mapsto \pi(\mathbf{y}) \mapsto \pi^2(\mathbf{y}) \mapsto \dots \mapsto \pi^{b-1}(\mathbf{y}) \mapsto \mathbf{x}.$$

This is a single cycle of length $a + b$; all lattice sites listed are distinct. After the swap, we have

$$\mathbf{y} \mapsto \pi(\mathbf{x}) \mapsto \pi^2(\mathbf{x}) \mapsto \dots \mapsto \pi^{a-1}(\mathbf{x}) \mapsto \mathbf{y} \quad \text{and} \quad \mathbf{x} \mapsto \pi(\mathbf{y}) \mapsto \pi^2(\mathbf{y}) \mapsto \dots \mapsto \pi^{b-1}(\mathbf{y}) \mapsto \mathbf{x}.$$

These are disjoint cycles of length a and b , respectively; the first contains \mathbf{x} and the second contains \mathbf{y} . \square

Remark 3.18. If $\mathbf{x} \circ \circ \mathbf{y}$ (i.e. the swap splits their common cycle): the old cycle lengths $\ell_{\mathbf{x}}(\pi)$ and $\ell_{\mathbf{y}}(\pi)$ are equal, and the new cycle lengths after the swap are

$$\ell_{\mathbf{x}}(\pi') = \ell_{\mathbf{y},\mathbf{x}}(\pi) \quad \text{and} \quad \ell_{\mathbf{y}}(\pi') = \ell_{\mathbf{x},\mathbf{y}}(\pi).$$

Otherwise (i.e. the swap merges the disjoint cycles):

$$\ell_{\mathbf{x}}(\pi') = \ell_{\mathbf{y}}(\pi) = \ell_{\mathbf{x}}(\pi) + \ell_{\mathbf{y}}(\pi).$$

3.4 Winding cycles and the swap-and-reverse algorithm

The propositions of section 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 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 figure 7 and rigorously in [Kerl]. 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 [Kerl]), 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 [BPS06, PST98]. That is, a permutation loop is selected at random and

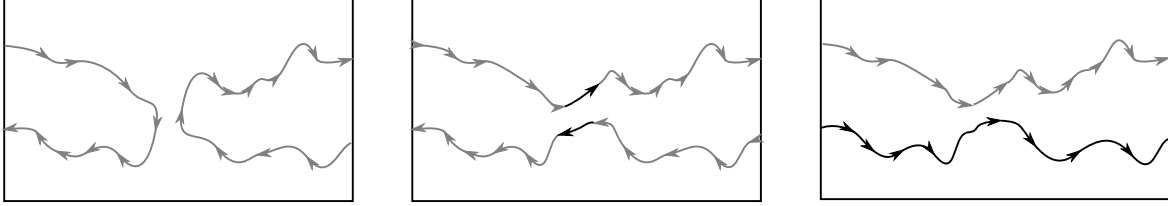


Figure 7: Conservation of winding number in the swap-only algorithm.

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 [Kerl]; 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. We believe that sampling permutations with even winding number does not harm our final result. The order parameters f_S and f_W depend on winding phenomena, but the other three, $1/\xi$, f_I , and f_M , do not; results obtained in section 4 using each of the five order parameters are compatible.

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 [CGGP] for a nice survey.

We have an infinite-volume random variable $S(T)$, e.g. any of the order parameters defined in section 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

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

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 :

$$S_L(T) = L^{-\rho/\nu} Q_S(L^{1/\nu} t) \sim L^{-\rho/\nu} Q_S((L/\xi)^{1/\nu}). \quad (3.19)$$

4 Results

Here we complete the steps sketched in section 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 [Berg], 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 .

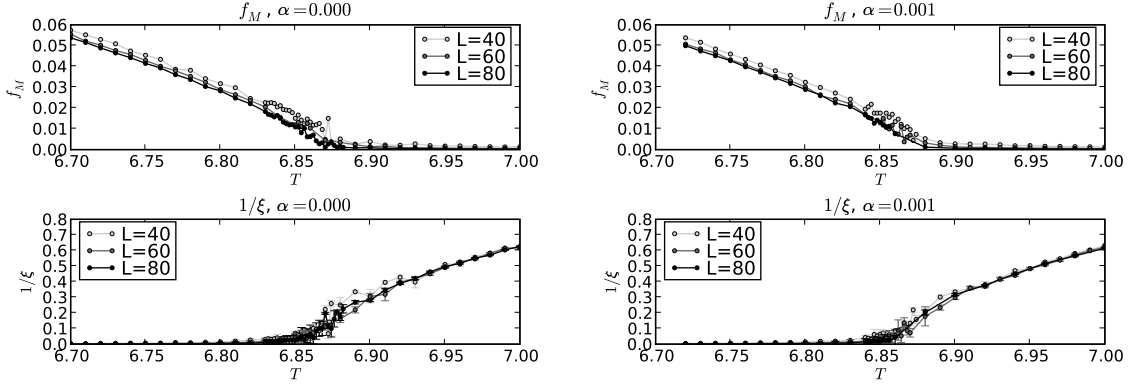


Figure 8: Order parameters f_M and $1/\xi$ for $L = 40, 60, 80$ and $\alpha = 0.000$ and 0.001 . The remaining order parameters f_S , f_W , and f_I behave similarly to f_M but with not all with the same critical exponents.

- 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 $L \rightarrow \infty$ 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 section 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 equation 2.6. Error bars are computed from regression uncertainties.

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 figure 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 figure 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 [Young] of the linear regression. Results are shown in figure 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 figure 10.

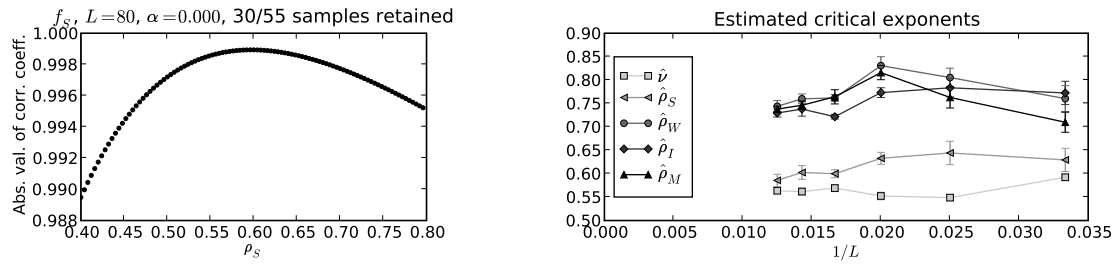


Figure 9: On the left: determination of critical exponent $\hat{\rho}_S(L, \alpha)$ for order parameter f_S , as the value which minimizes linear-regression error for $S_L(T, \alpha)^{1/\rho}$. Visually, one sees $\hat{\rho}_S(L = 80, \alpha = 0.0) \approx 0.59$. On the right: estimated critical exponents for $L = 30, 40, 50, 60, 70, 80$.

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	0.6246290	0.0001281	36
0.003	0.6247003	0.0001728	40
0.004	0.6247119	0.0002234	35

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

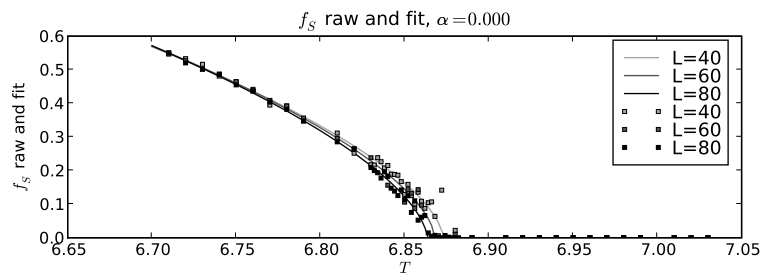


Figure 10: Power-law fit vs. simulational data for order parameter f_S , $\alpha = 0.000$

$\hat{\nu}$	0.5647	± 0.0035
$\hat{\rho}_S$	0.6168	± 0.0074
$\hat{\rho}_W$	0.7770	± 0.0083
$\hat{\rho}_I$	0.7534	± 0.0068
$\hat{\rho}_M$	0.7546	± 0.0079

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

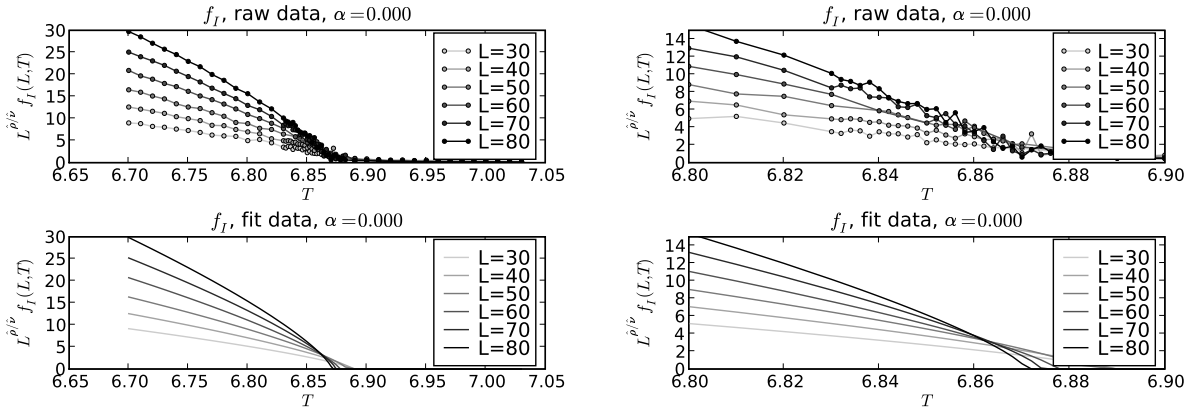


Figure 11: The crossing method to estimate $T_c(\alpha)$ for order parameter f_I , with $\hat{\rho}$ and $\hat{\nu}$ as above: $T_c(\alpha)$ corresponds to the horizontal coordinate of the intersection point of the plots. The upper-right-hand plot is a close-up of the upper-left-hand plot. The lower pair have raw MCMC data replaced with power-law-fit data.

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 figure 9.) Results are shown in table 2.

4.3 Determination of critical temperature

Given the above estimators of the critical exponents, the *crossing method* [CGGP] estimates $T_c(\alpha)$. Namely, we plot $L^{\hat{\rho}/\nu} S_L(T)$ as a function of T . At $T = T_c$ we have $t = 0$ and $L^{\hat{\rho}/\nu} S_L(T) = Q_S(0)$, regardless of L (equation 3.19). 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 figure 11. Results are in figure 13.

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,

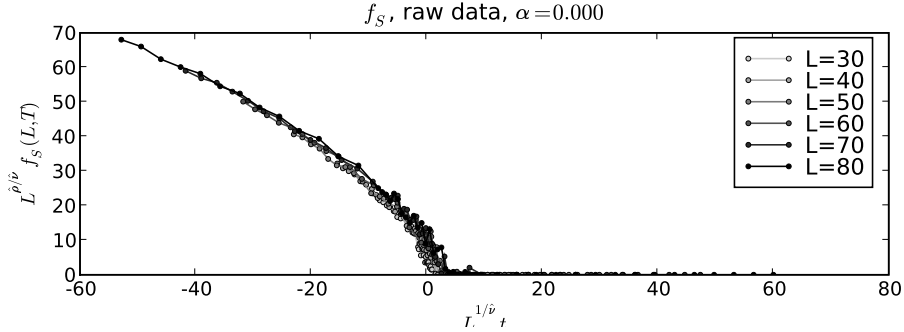


Figure 12: Collapse plot for order parameter f_S .

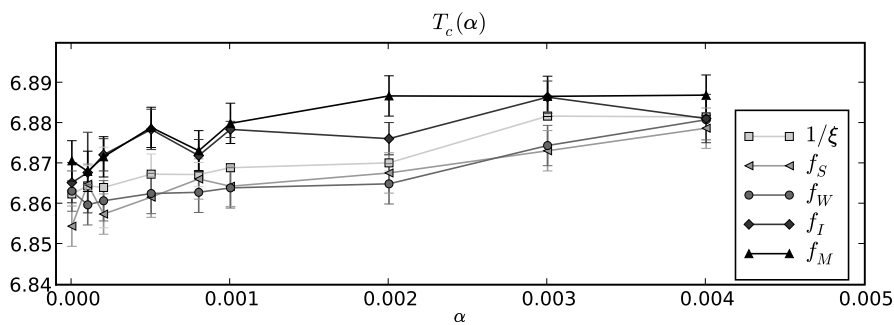


Figure 13: Critical temperature as function of α .

the curves for all L should coincide, or collapse, to within sampling error — which they do (e.g. figure 12).

4.5 Determination of the shift in critical temperature

As discussed in section 2.4, we are seeking a linear relationship between $\Delta T_c(\alpha)$ and α , with constant c . This can be visualized in figure 14, which is obtained from the $T_{c,S}(\alpha)$ data of figure 13 using equation 2.6. Performing a linear regression with error bars [Young] on all the $(\alpha, \Delta T_c(\alpha))$ data points, we find

$$c = 0.665 \pm 0.067.$$

This is the slope of the linear fit; the vertical intercept is 0.00026 ± 0.00012 which is consistent with equation 2.6. Within experimental uncertainty, this result, for points on the lattice with Ewens cycle-weights, matches the c value of equation 2.9 for point positions varying on the continuum with small-cycle-weight interactions.

4.6 Constancy of the macroscopic-cycle quotient

As discussed in section 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 $T < T_c$ 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 (figure 15). In the infinite-volume limit, f_I is




Figure 14: Shift in critical temperature as function of α . Recall from equation 2.6 that $\Delta T_c(\alpha) = \frac{T_c(\alpha) - T_c(0)}{T_c(0)}$.

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 section 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 $\hat{\rho}_{M,I}$:

$$\begin{aligned}
 \hat{\rho}_M &= 0.7560 & \hat{\rho}_I &= 0.7522 \\
 s_M &= 0.0564 & s_I &= 0.0481 \\
 n_M &= 48 & n_I &= 48 \\
 s_M/\sqrt{n_M} &= 0.008136 & s_I/\sqrt{n_I} &= 0.006938.
 \end{aligned}$$

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 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

$$\text{Cov}(\hat{\rho}_M, \hat{\rho}_I) = 0.0014 \quad s_d/\sqrt{n} = 0.0075.$$

Normalizing, we find

$$\hat{\rho}_M - \hat{\rho}_I = 0.0038 \quad \frac{\hat{\rho}_M - \hat{\rho}_I}{s_d/\sqrt{n}} = \frac{0.0038}{0.0075} = 0.5095.$$

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 figure 16). This merits further investigation.

4.7 Conclusions

(1) For annealed point positions, equation 2.9 gives $T_c(0) \approx 6.625$. Our result $T_c(0) = 6.863 \pm 0.003$ unambiguously shows that the lattice structure modifies the critical temperature, even in the non-interacting ($\alpha = 0$) case.

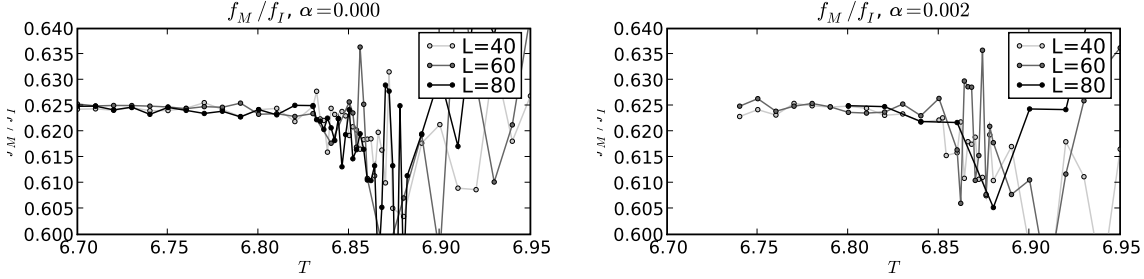


Figure 15: Macroscopic-cycle quotient f_M/f_I for $\alpha = 0.000, 0.002$.

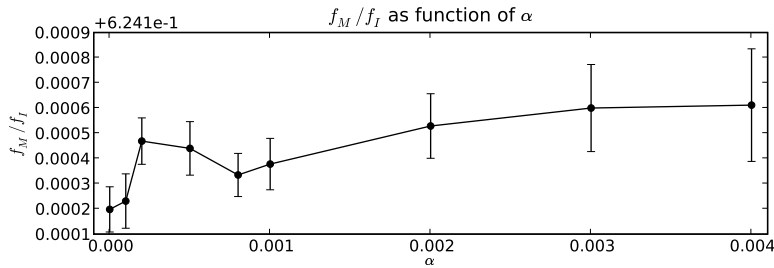


Figure 16: f_M/f_I as a function of α .

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

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

with

$$c = 0.665 \pm 0.067.$$

This is compatible (section 2.4) with the related result of [BU08]. Even though the lattice structure changes the critical temperature (conclusion 1), the *shift* in critical temperature is unaffected.

(3) As described in section 2.4, Shepp and Lloyd [SL] 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 .

5 Future work

Macroscopic-cycle quotient: 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.

Non-asymptotic algorithm correctness: The detailed-balance correctness proof of the swap-and-reverse algorithm shows that there is a non-zero transition probability between all pairs of permutations. However, those non-zero transition probabilities can be quite small. As the number of Metropolis steps goes to infinity, asymptotically all permutations can be reached; more interesting is the question of which permutations are actually reachable in reasonable simulation time. One answers this question empirically simply by running simulations; perhaps this suffices. For the system discussed in this paper, as well as for other systems studied using MCMC methods, it would be useful to have a non-asymptotic correctness theory.

Winding numbers of all parities: Ideally, one would have an algorithm to permit odd winding numbers, as discussed in section 3.4.

Bose-gas Hamiltonian: 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 (equation 2.1) rather than the cycle-weight interaction (equation 2.2). The interaction term V is a CPU-intensive Brownian-bridge computation [BU07]; 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 section 3.1. Namely, one picks a point and moves it to a new position nearby, using the detailed-balance condition to choose the acceptance probability [Kerl].

Third, software efficiency requires a hierarchical partitioning of Λ . The Metropolis step of section 3.1 relies on picking $\pi(\mathbf{y})$ near to $\pi(\mathbf{x})$. For points on the lattice, this is easy: each site has six nearest neighbors. For freely placed points, though, one must remember which sites are close to which. The most naive implementation involves computing the distances between all $N(N-1)/2$ pairs of points; the $O(N^2)$ computation time is overwhelming. Instead, the lattice may be partitioned into subcubes. Distances need to be computed only between each given point \mathbf{x} and those in \mathbf{x} 's subcube and the 26 nearest-neighbor subcubes.

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. This is a mathematical effort.

6 Acknowledgements

The author's doctoral dissertation work was done with co-advisors Daniel Ueltschi and Tom Kennedy. Daniel Gandolfo contributed several helpful discussions. The author was supported for three semesters by the National Science Foundation, through NSF grant DMS-0601075 as well as the University of Arizona Department of Mathematics VIGRE grant.

References

- [BBHLV] Baym, G., Blaizot, J.-P., Holzmann, M., Laloë, F., and Vautherin, D. *Bose-Einstein transition in a dilute interacting gas*. [arXiv:cond-mat/0107129v2](https://arxiv.org/abs/cond-mat/0107129v2). Eur. Phys. J. B 24, 107-124 (2001).
- [Berg] Berg, B. *Markov Chain Monte Carlo Simulations and Their Statistical Analysis*. World Scientific Publishing (2004).

- [BPS06] Boninsegni, M., Prokof'ev, N.V., and Svistunov, B.V. *Worm algorithm and diagrammatic Monte Carlo: A new approach to continuous-space path integral Monte Carlo simulations*. Physical Review E **74**, 036701 (2006).
- [BU07] Betz, V. and Ueltschi, D. *Spatial random permutations and infinite cycles*. arXiv:0711.1188. Commun. Math. Phys. 285, 469-501 (2009).
- [BU08] Betz, V. and Ueltschi, D. *Spatial random permutations with small cycle weights*. arXiv:0812.0569v1. Probabl. Th. Rel. Fields (2010).
- [CGGP] Caracciolo, S., Gambassi, A., Gubinelli, M., and Pelissetto, A. *Finite-Size Scaling in the Driven Lattice Gas*. arXiv:cond-mat/0312175. Journal of Statistical Physics, vol. 115, Nos. 1/2, April 2004.
- [Ewens] Ewens, W.J. *The sampling theory of selectively neutral alleles*. Theor. Popul. Biol. 3, 87-112 (1972).
- [Feynman] Feynman, R.P. *Atomic Theory of the λ Transition in Helium*. The Physical Review, vol. 91, no. 6 (1953).
- [Golomb] Golomb, S.W. *Random permutations*. Bull. Amer. Math. Soc. 70 (1964), 747.
- [GRU] Gandolfo, D., Ruiz, J., and Ueltschi, D. *On a model of random cycles*. arXiv:cond-mat/0703315. Statist. Phys. 129, 663-676 (2007).
- [Kerl] Kerl, J. *Critical behavior for the model of random spatial permutations*. Doctoral dissertation, University of Arizona, May 2010 (in progress).
- [LB] Landau, D.P. and Binder, K. *A Guide to Monte Carlo Simulations in Statistical Physics* (2nd ed.). Cambridge University Press (2005).
- [NR] Press, W. et al. *Numerical Recipes* (2nd ed.). Cambridge University Press (1992).
- [PC87] Pollock, E.L. and Ceperley, D.M. *Path-integral computation of superfluid densities*. Physical Review B, vol. 36, no. 16 (1987).
- [PO] Penrose, O. and Onsager, L. *Bose-Einstein Condensation and Liquid Helium*. The Physical Review, vol. 104, no. 3 (1956).
- [PST98] Prokof'ev, N.V., Svistunov, B.V., and Tupitsyn, I.S. *Exact, complete, and universal continuous-time worldline Monte Carlo approach to the statistics of discrete quantum systems*. Journal of Experimental and Theoretical Physics, vol. 87, no. 2 (1998).
- [SL] Shepp, L.A. and Lloyd, S.P. *Ordered Cycle Length in a Random Permutation*. Trans. Amer. Math. Soc. 121, (1966), 340-357.
- [Sütő1] Sütő, A. *Percolation transition in the Bose gas*. J. Phys. A: Math. Gen. **26** (1993) 4689-4710.
- [Sütő2] Sütő, A. *Percolation transition in the Bose gas II*. J. Phys. A: Math. Gen. **35** (2002) 6995-7002.
- [SU09] Seiringer, R. and Ueltschi, D. *Rigorous upper bound on the critical temperature of dilute Bose gases*. arXiv.org:0904.0050. Phys. Rev. B 80, 014502 (2009).
- [U07] Ueltschi, D. *The model of interacting spatial permutations and its relation to the Bose gas*. arXiv:0712.2443v3. Mathematical Results in Quantum Mechanics, pp. 225-272, World Scientific (2008).
- [Young] Young, H.D. *Statistical Treatment of Experimental Data*. McGraw-Hill (1962).