Chapter 7

The worm algorithm

This chapter, like chapter 5, presents a Markov chain for MCMC sampling of the model of random spatial permutations (chapter 2) within the framework of chapter 4. The worm algorithm solves the problem of winding-number-parity conservation within the SAR algorithm (section 5.4). However, a stopping time which is poorly bounded in the lattice size prevents the worm algorithm from being our algorithm of choice.

7.1 The idea of the worm algorithm

Worm algorithms have been used heavily in path-integral Monte Carlo (PIMC) simulations: see [BPS06, PST98, GCL97, KPS, NL04, PC87]. The context is that interparticle interactions are modeled using Brownian bridges in the Feynman-Kac approach. A naive, pre-PIMC sampling approach involves generating separate Brownian bridges from point A to point B. The PIMC idea is to generate a single Brownian bridge, then modify a bit at a time using MCMC methods.

For the random-cycle model with true Bose interactions, the Brownian bridges implicit in the V terms of equation (2.1.1) have been integrated out in equation (2.1.3). In fact, this is the key selling point of the random-cycle model in the larger context of the Bose gas (larger meaning beyond the scope of this dissertation). Most of the complexity of PIMC simulations, which are an efficient approach to Brownian bridges, goes away. If one were to adapt a PIMC worm algorithm to the RCM, one would need to spend significant time learning about PIMC. Yet it is likely that most of the complexity will also go away. Instead, it is simpler to ask: If one were to have a worm algorithm for the random-cycle model, what properties would it have? We require the following:

- We have a lattice with a fixed number N of points. There is no desire to work in the grand-canonical ensemble.
- We want the ability to open and close permutation cycles. (An open cycle is a "worm".)
- Given that, tips of open cycles may wander around the 3-torus before closing, permitting arbitrary winding numbers.

Thus, we want to sometimes open a cycle, then modify it with SO-like steps, then close it again. Following PIMC worm algorithms, all our Metropolis steps will involve the worm. This does touch all lattice points: a worm is opened at a site, then modified, then closed. Then, a worm is opened somewhere else, and so on.

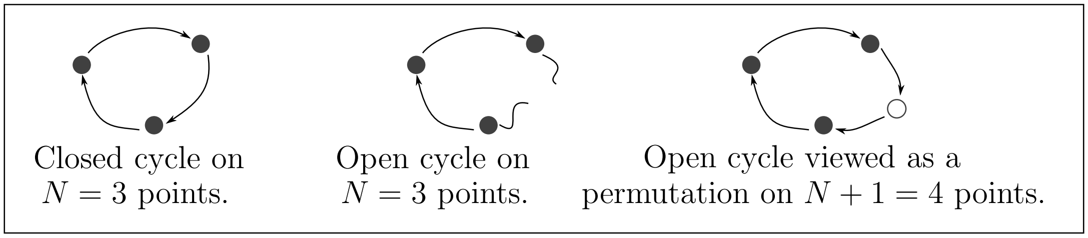

FIGURE 7.1. Open cycles as permutations on $N + 1$ points.

Question: Can we leverage our knowledge of permutations? To see how, consider a closed cycle and an open cycle on $N = 3$ points (see also figure 7.1):

$$\begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{pmatrix}, \qquad \begin{pmatrix} 1 & 2 & 3 & \\ 2 & 3 & & 1 \end{pmatrix}$$

In the open cycle, $1 \mapsto 2$, $2 \mapsto 3$, $3 \mapsto$ nothing, and nothing $\to 1$. Call that nothing *something* — the *wormhole point*. It is an $(N+1)$st point, w:

$$\begin{pmatrix} 1 & 2 & 3 & w \\ 2 & 3 & w & 1 \end{pmatrix}$$

Henceforth, the wormhole point will be written as w or $N + 1$. In diagrams, it will be an open dot while the other N points will be written with filled dots (as in figure 7.1). Now we have permutations on $\mathcal{S}_{N+1}$. Given $\pi \in \mathcal{S}_N$, inject π into $\mathcal{S}_{N+1}$ via $\pi(w) = w$.

Definition 7.1.1. For $\pi \in \mathcal{S}_{N+1}$, we say π is a *closed permutation* if $\pi(w) = w$. We say π is an *open permutation* if $\pi(w) \neq w$. Likewise, a cycle of π is said to be open or closed, respectively, if it does or does not contain w.

Remark. The PIMC jargon is that closed permutations are in the *Z sector* (for partition function), while open permutations are in the *G sector* (for Matsubara Green's function).

The goal is to invent an energy function, Gibbs distribution, and Metropolis algorithm for these extended permutations in $\mathcal{S}_{N+1}$ such that the marginal distribution on $\mathcal{S}_{N+1}$, conditioned on closed permutations, matches the RCM Gibbs distribution (equation (2.1.4)). Then, random variables will be sampled only at closed permutations.

7.2 Extended random-cycle model

Recall that we inject $\pi \in \mathcal{S}_N$ into $\mathcal{S}_{N+1}$ via $\pi(w) = w$. The $(N+1)$st point w is *non-spatial*: it has no distance associated with it.

Definition 7.2.1. The *extended lattice* is

$$\Lambda' := \Lambda \cup \{w\}$$

Definition 7.2.2. For $\pi \in \mathcal{S}_{N+1}$, define

$$H'(\pi) = \frac{T}{4} \sum_{\substack{i=1 \\ \pi(\mathbf{x}_i) \neq w}}^{N} \|\mathbf{x}_i - \mathbf{x}_{\pi(i)}\|_\Lambda^2 + \sum_{\ell=1}^{N} \alpha_\ell r_\ell(\pi) + \gamma \, 1_{\mathcal{S}_{N+1} \setminus \mathcal{S}_N}(\pi). \tag{7.2.3}$$

That is, we add γ to the energy if the permutation is open. Note that this *extended energy* agrees with the RCM energy (equation (2.1.3)) on closed permutations. (The γ term is only one particular choice; one might develop a better choice.) This is used to prove the marginality condition below. The *extended Gibbs distribution* and *extended partition function* are defined in the obvious way, as follows.

Definition 7.2.4. Let

$$P'_{\text{Gibbs}}(\pi) = \frac{e^{-H'(\pi)}}{Z'} \tag{7.2.5}$$

where the partition function is

$$Z' = \sum_{\pi \in \mathcal{S}_{N+1}} e^{-H'(\pi)}. \tag{7.2.6}$$

7.3 Proof of marginality

As long as the energy function for the ERCM and the RCM agree on closed permutations, the following desired marginality condition holds. This means that the interaction in section 2.1 — or any other to-be-invented interaction model — may use the worm algorithm as long as it agrees on closed permutations.

Proposition 7.3.1 (Marginality condition)**.** *Let $\mathcal{S}_N \hookrightarrow \mathcal{S}_{N+1}$ by taking $\pi(w) = w$. Let H, H' be energy functions on $\mathcal{S}_N$ and $\mathcal{S}_{N+1}$, respectively, such that for all $\pi \in \mathcal{S}_N$,*

$$H(\pi) = H'(\pi). \tag{7.3.2}$$

Let $P_{\text{Gibbs}}, P'_{\text{Gibbs}}, Z, Z'$ be as above. Then for $\pi \in \mathcal{S}_N$,

$$P'_{\text{Gibbs}}(\pi \mid \pi \in \mathcal{S}_N) = P_{\text{Gibbs}}(\pi). \tag{7.3.3}$$

Proof. Let $\pi \in \mathcal{S}_N$. The left-hand side of equation (7.3.3) is, by definition of conditional expectation,

$$P'_{\text{Gibbs}}(\pi \mid \pi \in \mathcal{S}_N) = \frac{P'_{\text{Gibbs}}(\pi)\, 1_{\mathcal{S}_N}(\pi)}{P'_{\text{Gibbs}}(\mathcal{S}_N)}.$$

The numerator is the Gibbs probability for closed permutations, or zero for open ones:

$$P'_{\text{Gibbs}}(\pi)\, 1_{\mathcal{S}_N}(\pi) = \frac{1}{Z'} e^{-H'(\pi)}\, 1_{\mathcal{S}_N}(\pi) = \frac{1}{Z'} e^{-H(\pi)}\, 1_{\mathcal{S}_N}(\pi)$$

since H and H' agree on closed permutations. The denominator is the total probability of closed permutations:

$$P'_{\text{Gibbs}}(\mathcal{S}_N) = \frac{1}{Z'} \sum_{\pi \in \mathcal{S}_N} e^{-H'(\pi)} = \frac{1}{Z'} \sum_{\pi \in \mathcal{S}_N} e^{-H(\pi)}.$$

Since $\pi \in \mathcal{S}_N$, the ratio is

$$\frac{\frac{1}{Z'} e^{-H(\pi)}\, 1_{\mathcal{S}_N}(\pi)}{\frac{1}{Z'} \sum_{\pi \in \mathcal{S}_N} e^{-H(\pi)}} = \frac{e^{-H(\pi)}\, 1_{\mathcal{S}_N}(\pi)}{\sum_{\pi \in \mathcal{S}_N} e^{-H(\pi)}} = \frac{e^{-H(\pi)}\, 1_{\mathcal{S}_N}(\pi)}{Z} = P_{\text{Gibbs}}(\pi).$$

□

7.4 The worm algorithm

Now that we have the correct Gibbs distribution for the ERCM, the next step is to devise a Metropolis algorithm to sample from it. Below, we will prove correctness. The *worm algorithm*, within the context of the recipe in section 4.5, is as follows:

- A sweep begins with a closed permutation π.
- The permutation is now closed, so $\pi(w) = w$. Select a lattice site $\mathbf{x}$ at uniform random. With probability proportional to $1 \wedge e^{-\Delta H}$, open the permutation by swapping the arrows of $\mathbf{x}$ and w. This is called an *open* move. (See figure 7.2.)
- Now that the permutation is open, do a *head swap*, *tail swap*, or *close*.
- Head swap: Pick a lattice site $\mathbf{x}$ nearest-neighbor to the lattice site $\pi^{-1}(w)$. With probability proportional to $1 \wedge e^{-\Delta H}$, swap arrows as in figure 7.2. The head swap is trivial if $\mathbf{x} = \pi^{-1}(w)$, which happens only if the head swap is rejected. The head swap would be a close if $\mathbf{x} = w$, but we choose $\mathbf{x}$ to be a lattice site. Thus, the permutation remains open on a head swap.

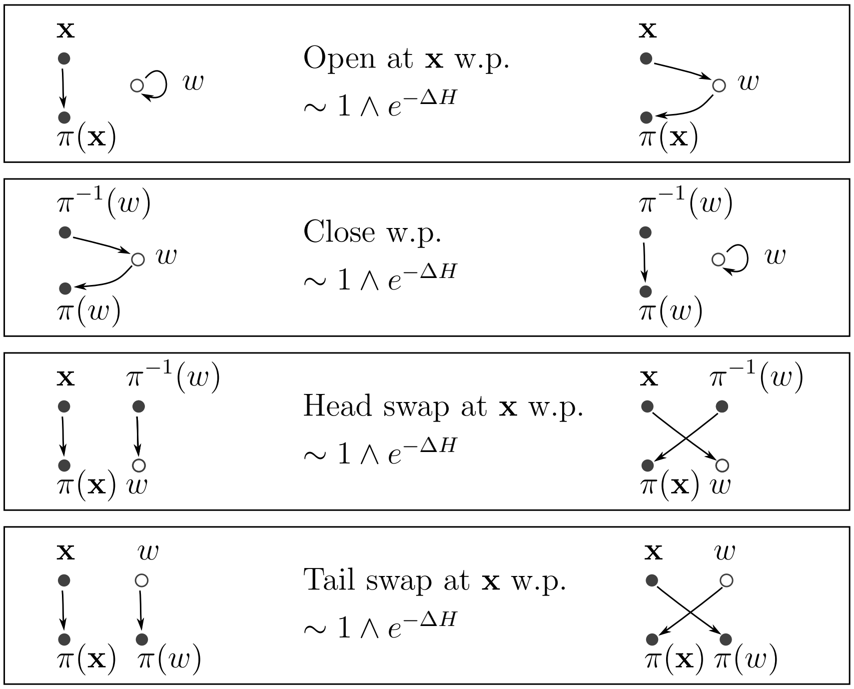

FIGURE 7.2. Metropolis moves for the worm algorithm.

- Tail swap: Pick a lattice site $\pi(\mathbf{x})$ nearest-neighbor to the lattice site $\pi(w)$. With probability proportional to $1 \wedge e^{-\Delta H}$, swap arrows as in figure 7.2. The tail swap is trivial if $\pi(\mathbf{x}) = \pi(w)$, which happens only if the tail swap is rejected. The tail swap would be a close if $\pi(\mathbf{x}) = w$, but we choose $\pi(\mathbf{x})$ to be a lattice site. Thus, the permutation remains open on a tail swap.

- Close: with probability proportional to $1 \wedge e^{-\Delta H}$, swap arrows as in figure 7.2. The permutation is now closed.

- Once the permutation is closed — after an open, some number of head/tail swaps, and a close, or after a rejected open — a *worm sweep* has been completed. At every sweep, one may obtain a value of any desired random variables for inclusion in computation of their sample means.

Definition 7.4.1. A head swap at $\mathbf{x}$ is *trivial* if $\mathbf{x} = \pi^{-1}(w)$; a tail swap at $\mathbf{x}$ is trivial if $\pi(\mathbf{x}) = \pi(w)$.

7.5 Fibration of $\mathcal{S}_{N+1}$ over $\mathcal{S}_N$

The definitions and lemmas in this section facilitate explicit construction of the Markov matrix, and are necessary for proving correctness of the worm algorithm.

As suggested by figure 7.3, we may separate all of $\mathcal{S}_{N+1}$ into the closed permutations $\mathcal{S}_N$ and the open permutations $\mathcal{S}_{N+1} \setminus \mathcal{S}_N$. Furthermore, for each of the $N!$ closed permutations π, we may open π at any of the N sites $\mathbf{x}_1, \ldots, \mathbf{x}_N$. Collecting each of the N open permutations obtained from each closed permutation creates a *fibration* of $\mathcal{S}_{N+1}$. The key points about the structure of this fibration, formalized by the lemmas below, are as follows.

- Each open permutation is one opener move away from a base closed permutation. The N open permutations above a base closed permutation π are the fiber over π.
- This induces a disjoint partition of the open permutations $\mathcal{S}_{N+1} \setminus \mathcal{S}_N$.
- Opens and closes, as defined in section 7.4, stay within fibers; non-trivial head swaps and tail swaps cross fibers.
- For each open permutation, the six non-trivial head swaps and six tail swaps result in twelve distinct permutations.
- Head swaps and tail swaps are transitive on fibers.

We first define maps corresponding to worm Metropolis moves.

Definition 7.5.1. The four worm Metropolis moves of figure 7.2 may be viewed in terms of maps. Throughout, $\mathbf{z} \in \Lambda \cup \{w\}$.

Let $O: \ \mathcal{S}_N \times \Lambda \to \mathcal{S}_{N+1} \setminus \mathcal{S}_N$ send $O(\pi, \mathbf{x}) = \pi'$ such that

$$\begin{aligned} \pi'(\mathbf{x}) &= w, \\ \pi'(w) &= \pi(\mathbf{x}), \\ \pi'(\mathbf{z}) &= \pi(\mathbf{z}), \quad \mathbf{z} \neq \mathbf{x}, w. \end{aligned}$$

Let $C: \ \mathcal{S}_{N+1} \setminus \mathcal{S}_N \to \mathcal{S}_N$ send $C(\pi) = \pi'$ such that

$$\begin{aligned} \pi'(\pi^{-1}(w)) &= \pi(w), \\ \pi'(w) &= w, \\ \pi'(\mathbf{z}) &= \pi(\mathbf{z}), \quad \mathbf{z} \neq \pi^{-1}(w), w. \end{aligned}$$

Let $S: \ \mathcal{S}_{N+1} \setminus \mathcal{S}_N \times \Lambda \to \mathcal{S}_{N+1} \setminus \mathcal{S}_N$ send $S(\pi, \mathbf{x}) = \pi'$ such that

$$\begin{aligned} \pi'(\mathbf{x}) &= w, \\ \pi'(\pi^{-1}(w)) &= \pi(\mathbf{x}), \\ \pi'(\mathbf{z}) &= \pi(\mathbf{z}), \quad \mathbf{z} \neq \mathbf{x}, \pi^{-1}(w). \end{aligned}$$

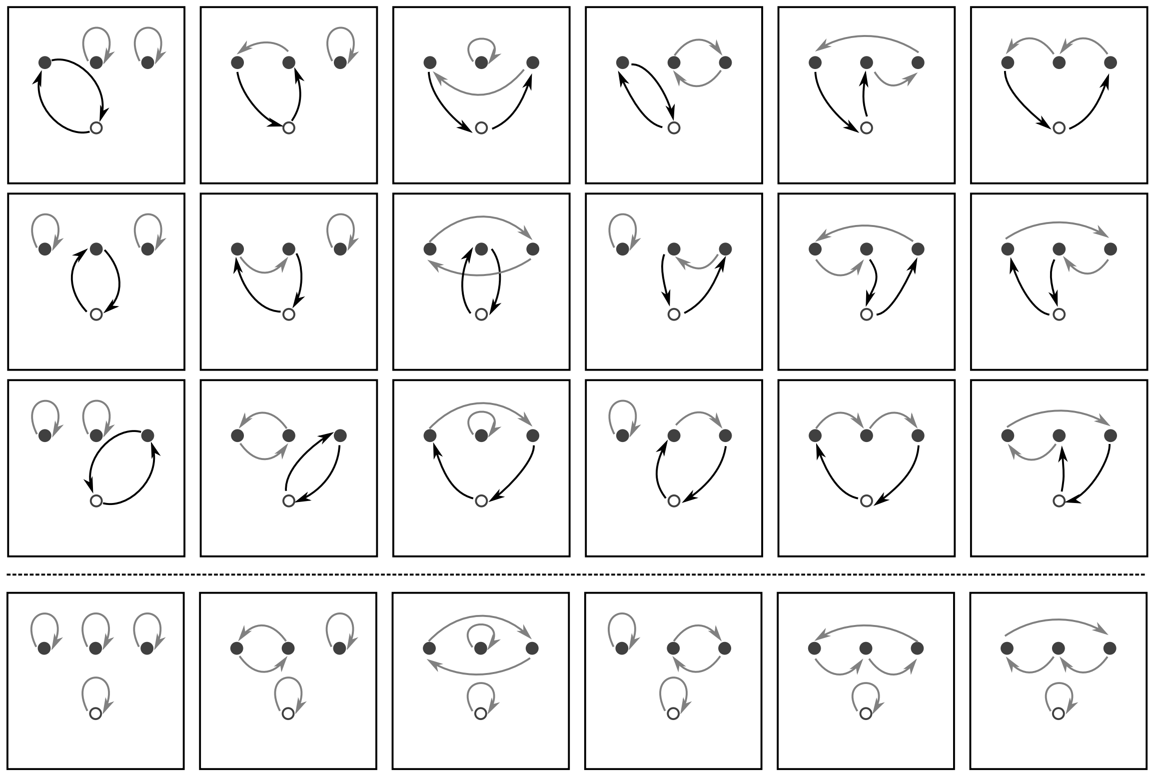

FIGURE 7.3. Fibration of $\mathcal{S}_4$ over $\mathcal{S}_3$. Closed permutations (i.e. $\mathcal{S}_3$) are along the bottom row; open permutations (i.e. $\mathcal{S}_4 \setminus \mathcal{S}_3$) are above the bottom row. The column, or fiber, above each closed permutation π contains the open permutations obtained from π by an opener move. Arrows modified by opener moves are shown in black.

Let $T : \ \mathcal{S}_{N+1} \setminus \mathcal{S}_N \times \Lambda \rightarrow \mathcal{S}_{N+1} \setminus \mathcal{S}_N$ send $T(\pi, \mathbf{x}) = \pi'$ such that

$$\begin{aligned} \pi'(\mathbf{x}) &= \pi(w), \\ \pi'(w) &= \pi(\mathbf{x}), \\ \pi'(\mathbf{z}) &= \pi(\mathbf{z}), \quad \mathbf{z} \neq \mathbf{x}, w. \end{aligned}$$

Throughout the proofs of the fibration-structure lemmas, we will use the following fact.

Lemma 7.5.2. *If* $\mathbf{x} \neq \mathbf{y}$*, then* $\pi(\mathbf{x}) \neq \pi(\mathbf{y})$ *and* $\pi^{-1}(\mathbf{x}) \neq \pi^{-1}(\mathbf{y})$*.*

Proof. If $\mathbf{x} \neq \mathbf{y}$ and $\pi(\mathbf{x}) = \pi(\mathbf{y})$, then π is not 1-1 which is a contradiction since π is a permutation. This applies to π^{-1} as well, since π^{-1} is also a permutation. □

Now we may prove the fibration-structure lemmas.

Lemma 7.5.3. *Each open permutation π is one opener move away from a base closed permutation π'. That is, for all $\pi \in \mathcal{S}_{N+1} \setminus \mathcal{S}_N$, there exists $\pi' \in \mathcal{S}_N$ such that $C(\pi) = \pi'$.*

Proof. Let $\pi \in \mathcal{S}_{N+1}$. Since π is open, $\pi(w) \neq w$ and $\pi^{-1}(w) \neq w$. Let $\mathbf{a} = \pi^{-1}(w)$ and $\mathbf{b} = \pi(w)$. Both are lattice points. Applying C, we have $C(\pi) = \pi'$ where $\pi'(\mathbf{a}) = \mathbf{b}$, $\pi'(w) = w$, and $\pi'(\mathbf{z}) = \pi(\mathbf{z})$ for all remaining lattice points $\mathbf{z} \neq \mathbf{a}, \mathbf{b}$. Since $\pi'(w) = w$, π' is closed. □

Definition 7.5.4. For $\pi \in \mathcal{S}_N$, $C^{-1}(\pi) \subset \mathcal{S}_{N+1} \setminus \mathcal{S}_N$ is the **fiber** of open permutations over π.

Lemma 7.5.5. *Opens and closes stay within fibers, and each fiber has N elements.*

Proof. Closes stay within fibers by definition of fiber. Next, fix $\pi \in \mathcal{S}_N$ and let $\mathbf{x}_1, \mathbf{x}_2 \in \Lambda$. (These are two different ways to open the same closed permutation.) Let

$$\pi_1' = O(\pi, \mathbf{x}_1), \quad \pi_2' = O(\pi, \mathbf{x}_2).$$

Then π_1' and π_2' have

$$\mathbf{x}_1 \mapsto w \mapsto \pi(\mathbf{x}_1), \qquad \mathbf{x}_2 \mapsto w \mapsto \pi(\mathbf{x}_2),$$

respectively, agreeing with π at all other lattice points $\mathbf{z}$. Now, $C(\pi_1')$ and $C(\pi_2')$ have

$$\mathbf{x}_1 \mapsto \pi(\mathbf{x}_1), w \mapsto w, \qquad \mathbf{x}_2 \mapsto \pi(\mathbf{x}_2), w \mapsto w$$

respectively, agreeing with π at all other lattice points $\mathbf{z}$. But this means $C(\pi_1')$ agrees with $C(\pi_2')$ agree at all points of Λ', so $C(\pi_1') = C(\pi_2')$. Thus, π_1' and π_2' are in the same fiber.

For the last claim, fix $\pi \in \mathcal{S}_N$ and enumerate the N lattice points of Λ as $\mathbf{x}_1, \ldots, \mathbf{x}_N$. We claim that the N permutations

$$\pi_1' = O(\pi, \mathbf{x}_1), \ldots, \pi_N' = O(\pi, \mathbf{x}_N),$$

which are all now known to be in the same fiber, are all distinct. To see this, fix $i \neq j$ from out of $\{1, 2, \ldots, N\}$. Then π_i' and π_j' have

$$\mathbf{x}_i \mapsto w \mapsto \pi(\mathbf{x}_i), \qquad \mathbf{x}_j \mapsto w \mapsto \pi(\mathbf{x}_j).$$

Since $\mathbf{x}_i \neq \mathbf{x}_j$, by lemma 7.5.2 $\pi(\mathbf{x}_i) \neq \pi(\mathbf{x}_j)$. Since

$$\pi_i'(w) = \pi(\mathbf{x}_i) \neq \pi(\mathbf{x}_j) = \pi_j'(w),$$

π_i' and π_j' send w to different points. Therefore, the permutations π_i' and π_j' are distinct. □

Lemma 7.5.6. *The above fibration induces a disjoint partition of the open permutations $\mathcal{S}_{N+1} \setminus \mathcal{S}_N$. That is, for $\pi'_1, \pi'_2 \in \mathcal{S}_N$,*

$$\pi'_1 \neq \pi'_2 \implies C^{-1}(\pi'_1) \cap C^{-1}(\pi'_2) = \emptyset \qquad \text{and} \qquad \bigcup_{\pi \in \mathcal{S}_N} C^{-1}(\pi) = \mathcal{S}_{N+1} \setminus \mathcal{S}_N.$$

Proof. For the first claim, suppose the intersection is non-empty. Let $\pi \in \mathcal{S}_{N+1} \setminus \mathcal{S}_N$ be such that $\pi \in C^{-1}(\pi'_1)$ and $\pi \in C^{-1}(\pi'_2)$. This means $C(\pi) = \pi'_1$ and $C(\pi) = \pi'_2$ with $\pi'_1 \neq \pi'_2$, which is a contradiction since the map C is uniquely defined for all $\pi \in \mathcal{S}_{N+1} \setminus \mathcal{S}_N$.

For the second claim: there are $N!$ closed permutations. We know from the first claim that the $N!$ fibers, one above each closed permutation, are all disjoint. From lemma 7.5.5, we know that each fiber has N elements. We have accounted for all $N \cdot N! = (N+1)! - N!$ open permutations, so we must have all of $\mathcal{S}_{N+1} \setminus \mathcal{S}_N$. □

Lemma 7.5.7. *Non-trivial head swaps and tail swaps (definition 7.4.1) cross fibers.*

Proof. First consider head swaps. Let $\pi, \pi' \in \mathcal{S}_{N+1} \setminus \mathcal{S}_N$ differ by a non-trivial head swap, namely, there is $\mathbf{x} \neq \pi^{-1}(w)$ such that $\pi' = S(\pi, \mathbf{x})$. Then π and π' have

$$\begin{array}{llllllllll} \pi: & \mathbf{x} & \mapsto & \pi(\mathbf{x}) & \mapsto & \pi^2(\mathbf{x}), & \pi^{-1}(w) & \mapsto & w & \mapsto & \pi(w), \\ \pi': & \pi^{-1}(w) & \mapsto & \pi(\mathbf{x}) & \mapsto & \pi^2(\mathbf{x}), & \mathbf{x} & \mapsto & w & \mapsto & \pi(w), \end{array}$$

respectively. Now apply C to each: $C(\pi)$ and $C(\pi')$ have

$$\begin{array}{llllllllll} C(\pi): & \mathbf{x} & \mapsto & \pi(\mathbf{x}) & \mapsto & \pi^2(\mathbf{x}), & \pi^{-1}(w) & \mapsto & \pi(w), & w & \mapsto & w, \\ C(\pi'): & \pi^{-1}(w) & \mapsto & \pi(\mathbf{x}) & \mapsto & \pi^2(\mathbf{x}), & \mathbf{x} & \mapsto & \pi(w), & w & \mapsto & w, \end{array}$$

respectively. Since $\mathbf{x} \neq \pi^{-1}(w)$, $C(\pi) \neq C(\pi')$.

Next, consider tail swaps. Let $\pi, \pi' \in \mathcal{S}_{N+1} \setminus \mathcal{S}_N$ differ by a non-trivial tail swap, namely, there is $\pi(\mathbf{x}) \neq \pi(w)$ such that $\pi' = T(\pi, \mathbf{x})$. Then π and π' have

$$\begin{array}{lllllllllll} \pi: & \pi^{-1}(\mathbf{x}) & \mapsto & \mathbf{x} & \mapsto & \pi(\mathbf{x}), & \pi^{-1}(w) & \mapsto & w & \mapsto & \pi(w), \\ \pi': & \pi^{-1}(\mathbf{x}) & \mapsto & \mathbf{x} & \mapsto & \pi(w), & \pi^{-1}(w) & \mapsto & w & \mapsto & \pi(\mathbf{x}), \end{array}$$

respectively. Now apply C to each: $C(\pi)$ and $C(\pi')$ have

$$\begin{array}{llllllllll} C(\pi): & \pi^{-1}(\mathbf{x}) & \mapsto & \mathbf{x} & \mapsto & \pi(\mathbf{x}), & \pi^{-1}(w) & \mapsto & \pi(w), & w & \mapsto & w, \\ C(\pi'): & \pi^{-1}(w) & \mapsto & \mathbf{x} & \mapsto & \pi(w), & \pi^{-1}(w) & \mapsto & \pi(\mathbf{x}), & w & \mapsto & w, \end{array}$$

respectively. Since $\pi(\mathbf{x}) \neq \pi(w)$, $C(\pi) \neq C(\pi')$. □

Lemma 7.5.8. *For each open permutation, the six non-trivial head swaps and six non-trivial tail swaps result in twelve distinct permutations.*

Proof. Fix $\pi \in \mathcal{S}_{N+1} \setminus \mathcal{S}_N$. Let $\mathbf{x}_1, \ldots, \mathbf{x}_6$ be the six nearest-neighbor lattice sites to the lattice site $\pi^{-1}(w)$; let $\mathbf{y}_1, \ldots, \mathbf{y}_6$ be the six lattice sites such that $\pi(\mathbf{y}_1), \ldots, \pi(\mathbf{y}_6)$ are nearest-neighbor lattices site to the lattice site $\pi(w)$. (See figure 7.2.)

First, we show that the six permutations $S(\pi, \mathbf{x}_1), \ldots, S(\pi, \mathbf{x}_6)$ are distinct. Let $i \neq j$ for $i, j = 1, \ldots, 6$; let $\pi_i = S(\pi, \mathbf{x}_i)$ and $\pi_j = S(\pi, \mathbf{x}_j)$. Then π has

$$\pi: \quad \mathbf{x}_i \mapsto \pi(\mathbf{x}_i), \qquad \mathbf{x}_j \mapsto \pi(\mathbf{x}_j), \qquad \pi^{-1}(w) \mapsto w;$$

π_i, and π_j have

$$\begin{array}{lllllll} \pi_i: & \mathbf{x}_i & \mapsto & w, & \pi^{-1}(w) & \mapsto & \pi(\mathbf{x}_i), \\ \pi_j: & \mathbf{x}_j & \mapsto & w, & \pi^{-1}(w) & \mapsto & \pi(\mathbf{x}_j), \end{array}$$

respectively. Since $\mathbf{x}_i \neq \mathbf{x}_j$, $\pi_i \neq \pi_j$.

Second, we show that the six permutations $T(\pi, \mathbf{y}_1), \ldots, T(\pi, \mathbf{y}_6)$ are distinct. Let $i \neq j$ for $i, j = 1, \ldots, 6$; let $\pi_i = T(\pi, \mathbf{y}_i)$ and $\pi_j = T(\pi, \mathbf{y}_j)$. Then π has

$$\pi: \quad \mathbf{y}_i \mapsto \pi(\mathbf{y}_i), \qquad \mathbf{y}_j \mapsto \pi(\mathbf{y}_j), \qquad w \mapsto \pi(w);$$

π_i, and π_j have

$$\begin{array}{lllllll} \pi_i: & \mathbf{y}_i & \mapsto & \pi(w), & w & \mapsto & \pi(\mathbf{y}_i), \\ \pi_j: & \mathbf{y}_j & \mapsto & \pi(w), & w & \mapsto & \pi(\mathbf{y}_j), \end{array}$$

respectively. Since $\mathbf{y}_i \neq \mathbf{y}_j$, by lemma 7.5.2 $\pi(\mathbf{y}_i) \neq \pi(\mathbf{y}_j)$. Since π_i, π_j send w to to different sites, $\pi_i \neq \pi_j$.

Third, we show that the head-swaps of π are distinct from the tail-swaps of π. Fix $\pi \in \mathcal{S}_{N+1} \setminus \mathcal{S}_N$ and let $i, j \in \{1, \ldots, 6\}$. Then π has

$$\pi: \quad \mathbf{x}_i \mapsto \pi(\mathbf{x}_i), \qquad \mathbf{y}_j \mapsto \pi(\mathbf{y}_j), \qquad \pi^{-1}(w) \mapsto w \mapsto \pi(w);$$

$S(\pi, \mathbf{x}_i)$ and $T(\pi, \mathbf{y}_j)$ have

$$\begin{array}{lllllllll} S(\pi, \mathbf{x}_i): & \pi^{-1}(w) & \mapsto & \pi(\mathbf{x}_i), & \mathbf{x}_i & \mapsto & w & \mapsto & \pi(w); \\ T(\pi, \mathbf{y}_j): & \mathbf{y}_j & \mapsto & \pi(w), & \pi^{-1}(w) & \mapsto & w & \mapsto & \mathbf{y}_j; \end{array}$$

respectively. Under these two permutations, w has images $\pi(w)$ and $\mathbf{y}_j$, respectively, and preimages $\mathbf{x}_i$ and $\pi^{-1}(w)$. By definition 7.4.1, the non-trivial head swap $S(\pi, \mathbf{x}_i)$ has $\mathbf{x}_i \neq \pi^{-1}(w)$ and the non-trivial tail swap $T(\pi, \mathbf{y}_j)$ has $\pi(w) \neq \mathbf{y}_j$. Thus, $S(\pi, \mathbf{x}_i)$ and $T(\pi, \mathbf{y}_j)$ are distinct permutations. □

7.6 Explicit construction of the Markov matrix

Transition probabilities were described in section 7.4 as being proportional to $1 \wedge e^{-\Delta H}$. We put the constants of proportionality to be the following:

- a for head swaps and tail swaps;

- b for closer moves;
- c for opener moves.

For SO/SAR, we chose the normalizing factor easily. Here, with a more complicated algorithm, we will choose the normalizing factors to satisfy detailed balance. In particular, in this section we will obtain $c = b = 1/N$ and $a = (1-b)/12$.

The Markov matrix at each Metropolis step is now $(N+1)! \times (N+1)!$:

- A closed permutation transitions only to itself, or to any of the N open permutations in the fiber above it. Thus, there are $N+1$ non-zero entries in π's row of A'.
- An open permutation transitions to any of the 12 open permutations available by head-swapping or tail-swapping, or to itself, or to the closed permutation at the base of its fiber. Thus, there are 14 non-zero entries in π's row of A'.

Definition 7.6.1. For open π, let

$$\{\mathbf{x}_1, \dots, \mathbf{x}_6\} = \{\mathbf{x} \in \Lambda : \|\mathbf{x}, \pi^{-1}(w)\|_\Lambda = 1\}$$

and

$$\{\mathbf{y}_1, \dots, \mathbf{y}_6\} = \{\mathbf{y} \in \Lambda : \|\pi(\mathbf{y}), \pi(w)\|_\Lambda = 1\}.$$

Then define

$$\begin{aligned} R_S(\pi) &= \{S(\pi, \mathbf{x}_1), \dots, S(\pi, \mathbf{x}_6)\}, \\ R_T(\pi) &= \{T(\pi, \mathbf{y}_1), \dots, T(\pi, \mathbf{y}_6)\}. \end{aligned}$$

These are the twelve open permutations reachable from π via head swaps and tail swaps, respectively (lemma 7.5.8). For closed π, define

$$R_O(\pi) = \{O(\pi, \mathbf{x}_1), \dots, O(\pi, \mathbf{x}_N)\}.$$

These are the N open permutations reachable from π via opener moves.

With these definitions, the entries of the transition matrix are as follows. In analogy with H, P_{Gibbs}, Z, etc. for the random-cycle model and H', P'_{Gibbs}, Z', etc. for the extended random-cycle model, we call this worm-algorithm transition matrix A' to distinguish it from the matrices A and $A_{\mathbf{x}}$ (equations (5.2.5) and (5.2.6)) for the swap-only algorithm.

If π is closed:

$$A'(\pi, \pi') = \begin{cases} c\left(1 \wedge e^{-H(\pi')+H(\pi)}\right), & \pi' \in R_O(\pi); \\ 1 - \displaystyle\sum_{\pi' \in R_O(\pi)} c\left(1 \wedge e^{-H(\pi')+H(\pi)}\right), & \pi' = \pi; \\ 0, & \text{otherwise.} \end{cases}$$

If π is open:

$$A'(\pi,\pi') = \begin{cases} a\left(1 \wedge e^{-H(\pi')+H(\pi)}\right), & \pi' \in R_S(\pi); \\ a\left(1 \wedge e^{-H(\pi')+H(\pi)}\right), & \pi' \in R_T(\pi); \\ b\left(1 \wedge e^{-H(\pi')+H(\pi)}\right), & \pi' = C(\pi); \\ 1 - t(\pi), & \pi' = \pi; \\ 0, & \text{otherwise} \end{cases}$$

where

$$t(\pi) = \left(\sum_{\pi'' \in R_S(\pi) \cup R_T(\pi)} a\left(1 \wedge e^{-H(\pi'')+H(\pi)}\right) \right) - b\left(1 \wedge e^{-H(C(\pi))+H(\pi)}\right)$$

For row normalization for closed π, note that $c(1 \wedge e^{-\Delta H})$ is between 0 and c so $\sum$ rest is between 0 and cN. Take

$$c = 1/N. \tag{7.6.2}$$

Row normalization for open π then gives

$$12a + b \leq 1. \tag{7.6.3}$$

In practice, we set $12a + b = 1$. That is, we do the following on open permutations: with probability $1/N$, propose a close; else, propose a head or tail swap with equal probability $\frac{1}{2}(1 - \frac{1}{N})$.

The Markov chain for worm Metropolis steps is homogeneous: we use the same transition matrix A' at each step. The correctness proofs of the following section will then imply (by the machinery of chapter 4) that we sample from the extended Gibbs distribution for $\mathcal{S}_{N+1}$. Then by proposition 7.3.1 we will sample from the Gibbs distribution for $\mathcal{S}_N$ whenever the permutation closes.

7.7 Correctness

As discussed in sections 4.5 and 5.3, we need to prove irreducibility, aperiodicity, and detailed balance for the worm Markov chain.

Proposition 7.7.1 (Irreducibility)**.** *The worm algorithm's Markov chain is irreducible.*

Proof. This follows immediately from propositions 5.3.1 and 7.7.2: namely, the worm's chain is irreducibility if the SO's chain is, and moreover the SO's chain is irreducible. □

Proposition 7.7.2. *The worm algorithm's Markov chain is irreducible if the SO algorithm's Markov chain is irreducible.*

Proof. The key point is that the composition of an open, head swap, and close are precisely an SO swap. Let $\mathbf{x}$ and $\mathbf{y}$ be lattice points such that $\pi(\mathbf{x})$ and $\pi(\mathbf{y})$ are nearest neighbors. Starting with π, then applying an open at $\mathbf{x}$, a head swap at $\mathbf{y}$, and a close, we have

$$\begin{array}{llllllllll}
\pi: & \mathbf{x} & \mapsto & \pi(\mathbf{x}), & \mathbf{y} & \mapsto & \pi(\mathbf{y}), & w & \mapsto & w; \\
\pi' = O(\pi,\mathbf{x}): & \mathbf{x} & \mapsto & w, & \mathbf{y} & \mapsto & \pi(\mathbf{y}), & w & \mapsto & \pi(\mathbf{x}); \\
\pi'' = S(\pi',\mathbf{y}): & \mathbf{x} & \mapsto & \pi(\mathbf{y}), & \mathbf{y} & \mapsto & w, & w & \mapsto & \pi(\mathbf{x}); \\
\pi''' = C(\pi''): & \mathbf{x} & \mapsto & \pi(\mathbf{y}), & \mathbf{y} & \mapsto & \pi(\mathbf{x}), & w & \mapsto & w.
\end{array}$$

This shows that, if the SO algorithm is irreducible on $\mathcal{S}_N$, the worm algorithm is irreducible on $\mathcal{S}_N$. But then the worm algorithm is also irreducible on $\mathcal{S}_{N+1}$: fix an initial and final permutation; close the initial permutation, if it is open, to obtain a closed permutation; use the preceding argument to reach the closed permutation which lies under the fiber of the desired final open permutation; do an open move (see lemma 7.5.5) if the final permutation is open. □

Remark 7.7.3. The worm algorithm has an additional degree of freedom. If $\mathbf{x}$ and $\mathbf{y}$ are nearest-neighbor lattice sites, then the composition of an open at $\mathbf{x}$, a tail swap at $\mathbf{y}$, and a close results in a similar swap of the jump targets of $\mathbf{x}$ and $\mathbf{y}$:

$$\begin{array}{llllllllll}
\pi: & \mathbf{x} & \mapsto & \pi(\mathbf{x}), & \mathbf{y} & \mapsto & \pi(\mathbf{y}), & w & \mapsto & w; \\
\pi' = O(\pi,\mathbf{x}): & \mathbf{x} & \mapsto & w, & \mathbf{y} & \mapsto & \pi(\mathbf{y}), & w & \mapsto & \pi(\mathbf{x}); \\
\pi'' = T(\pi',\mathbf{y}): & \mathbf{x} & \mapsto & w, & \mathbf{y} & \mapsto & \pi(\mathbf{x}) & w & \mapsto & \pi(\mathbf{y}); \\
\pi''' = C(\pi''): & \mathbf{x} & \mapsto & \pi(\mathbf{y}), & \mathbf{y} & \mapsto & \pi(\mathbf{x}), & w & \mapsto & w.
\end{array}$$

Proposition 7.7.4 (Aperiodicity)**.** *The worm algorithm's Markov chain is aperiodic.*

Proof. The proof is the same as in the SO case, proposition 5.3.5. □

Proposition 7.7.5 (Detailed balance)**.** *The Markov chain of the worm algorithm satisfies detailed balance with $b = c$.*

Proof. We need

$$P'_{\text{Gibbs}}(\pi)A'(\pi,\pi') = P'_{\text{Gibbs}}(\pi')A'(\pi',\pi).$$

For closed π to closed π': If $\pi = \pi'$ then we have detailed balance trivially. If $\pi \neq \pi'$ then $A'(\pi,\pi') = A'(\pi',\pi) = 0$ since there are no transitions between distinct closed permutations.

For closed π to open π': If π' is not in the fiber above π, then $A'(\pi,\pi') = A'(\pi',\pi) = 0$ since opens and closes respect fibers (lemma 7.5.5). Now suppose π' is in the fiber above π. As in the SO algorithm (proposition 5.3.6), do cases on ΔH positive or negative. If $H'(\pi') \leq H'(\pi)$, then

$$e^{-H'(\pi)}c = e^{-H'(\pi')}be^{-H'(\pi)}e^{H'(\pi')}.$$

Choose

$$b = c \tag{7.7.6}$$

to satisfy detailed balance. The case $H'(\pi') > H'(\pi)$ results in the same $b = c$ condition.

For open π to closed π': If π is not in the fiber above π', then $A'(\pi, \pi') = A'(\pi', \pi) = 0$ (lemma 7.5.5). If π is in the fiber above π', then we recover the $b = c$ condition.

It now remains to consider open π transitioning to open π'. We assume this to be the case for the rest of the proof.

If $A'(\pi, \pi') = 0$ then we claim $A'(\pi', \pi) = 0$, as in lemma 5.3.8. We have $\pi' \neq \pi$, $\pi' \notin R_S(\pi)$, and $\pi' \notin R_T(\pi)$. We need to show $\pi \neq \pi'$ (which is obvious), $\pi \notin R_S(\pi')$, and $\pi \notin R_T(\pi')$. We prove the contrapositive:

$$\pi \in \{\pi'\} \cup R_S(\pi') \cup R_T(\pi') \quad \Longrightarrow \quad \pi' \in \{\pi\} \cup R_S(\pi) \cup R_T(\pi).$$

If $\pi = \pi'$ then detailed balance is trivially satisfied. Suppose $\pi \in R_S(\pi')$. Then for some $\mathbf{x}_i$, $i = 1, \ldots, 6$, π' and π have

$$\begin{array}{llllllll} \pi': & \mathbf{x}_i & \mapsto & \pi'(\mathbf{x}_i), & \pi^{-1}(w) & \mapsto & w & \mapsto & w, \\ \pi: & \mathbf{x}_i & \mapsto & w, & \pi^{-1}(w) & \mapsto & w & \mapsto & \pi'(\mathbf{x}_i). \end{array}$$

The lattice sites $\mathbf{x}_i$ and w are nearest neighbors and π', π agree at all other sites, so there is a head swap sending π to π'. The case $\pi \in \mathbb{R}_T(\pi')$ is completely analogous. This completes the proof of the claim that $A'(\pi, \pi') = 0 \implies A'(\pi', \pi) = 0$.

If $A'(\pi, \pi') \neq 0$ then we claim $A'(\pi', \pi) \neq 0$, again as in lemma 5.3.8. The logic is the same as in the contrapositive argument which was just completed.

The last step is to show detailed balance for open π, π' where $A'(\pi, \pi') \neq 0$. Again we do cases on whether the energy decreases or increases. If $H'(\pi') \leq H'(\pi)$, then equation (7.7.6) is

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

If $H'(\pi') > H'(\pi)$, then we have

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

In either case, detailed balance holds. ∎

Remark. Note that for closed π, there are N choices of open π'; for open π, there is one choice of closed π'. In the software implementation, the $1/N$ for opens comes in through uniform-random choice of $\mathbf{x} \in \Lambda$. The result is that, for closed π, one may only attempt an open. For open π, one attempts a close $1/N$ of the time, and head or tail swaps each half the rest of the time, respectively.

As a sanity check, we point out that cycles may grow or shrink upon worm moves.

Proposition 7.7.7. *Non-trivial worm head swaps and tail swaps either split one cycle into two, or join two cycles into one.*

Proof. This is the same as for the SO case (proposition 7.7.7), which is strictly an algebraic result involving permutations: the non-spatiality of the w point plays no role. □

7.8 Stopping time and modification ideas

The essence of the winding-number problem, as discussed in section 5.4, is that the configuration space has multiple energy minima (which are equivalent to probability maxima), indexed by winding numbers W_x, W_y, and W_z. One might also say that the probability distribution for random spatial permutations is *multimodal*. The swap-only algorithm creates only permutations with winding numbers equal to 0. The swap-and-reverse algorithm creates permutations with even winding numbers: the cycle-reversal move has zero energy change and allows subsequent permutations to hop across a low double-winding-number barrier.

The worm algorithm was designed to permit cycles with winding numbers of both parities to be created: a cycle is opened, its tips wander around (perhaps around the torus), and then it recloses — all of these steps happening with low-energy changes afforded by worm tunneling through the energy barrier. The only problem is that the open worm tips wander around randomly within the L box, and fail to reconnect as L increases. This is the *stopping-time problem.* Specifically, histograms show that the distribution of the wormspan $\|\pi(w) - \pi^{-1}(w)\|_\Lambda$ peaks around $L/2$.

Recall from section 4.2 that the correctness proofs of sections 5.3 and 7.7 only address the limit $M \to \infty$; they do not address rate of convergence. The worm algorithm is correct, but we are not willing to wait long enough for it to produce its correct results. Figure 7.4 shows the problem. CPU time is plotted for 10^4 SAR sweeps, at $T = 6.0$, as a function of $N = L^3$ for $L = 5$ to 12. For the SAR algorithm, CPU time is nearly linear in N. (In fact, it has an N^2 dependency, but with a low constant of proportionality, as discussed in section 9.5.) For the worm algorithm, CPU time is not linear in N; we cannot complete a computation for L large enough to be interesting, namely, 40 to 80. Specifically, a log-log plot and regression on the data of figure 7.4 show that the worm algorithm is strongly quadratic in N. Fortunately, examination of random-variable plots for $L = 10$, comparing SAR to worm, show that similar results are produced — other than, of course, the winding-number histogram itself.

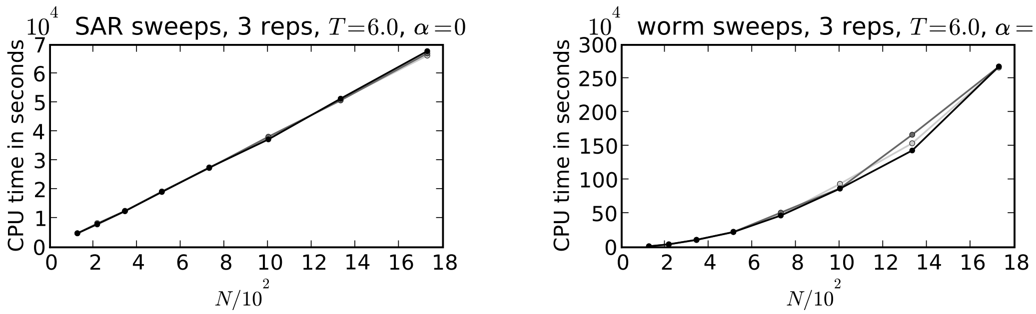

FIGURE 7.4. Scalability of SAR and worm algorithms. CPU times for 10^4 SAR sweeps and 10^3 worm sweeps are shown as a function of $N = L^3$ for $L = 5$ to 12. SAR time is nearly linear in N; worm time is strongly quadratic in N. Interesting L (40-80) are unattainable.

Other ideas for addressing the winding-number problem include the following:

- In analogy with cluster updates for the Ising model, form a band around the torus and do an L-cycle transformation. Couple the SAR algorithm with occasional band updates. However, band updates have a too-low acceptance rate, as was shown in chapter 6.
- Temporarily pinch the torus geometry somehow in the SAR algorithm, such that the distance penalty for wrapping around the torus is decreased.
- Temporarily reduce and restore the temperature T in the SAR algorithm — this is an annealing method. This approach brings with it a performance problem: re-thermalization (section 9.6) would need to be performed after each annealing step.
- Modify the worm algorithm to direct the worm somehow. At the time the worm is opened, add a distance weight of $\pm L$ in the x, y, or z direction which will be removed by a wrap around the torus, increasing or decreasing that winding-number component by 1. Our attempts to do this have not satisfied detailed balance.
- Review the PIMC literature again and seek other inspiration.

The worm algorithm, even though it is effectively unusable as currently designed, is the only way we currently have of sampling from the full winding-number configuration space, i.e. odd as well as even winding numbers. Thus, it will be worth the future effort to solve the stopping-time problem.