Numerical methods for random spatial permutations

Stochastic Processes and Their Applications, Berlin, 2009

John Kerl

Department of Mathematics, University of Arizona

July 30, 2009

The probability model

State space: $\Omega_{\Lambda,N} = \Lambda^N \times \mathcal{S}_N$, where $\Lambda = [0, L]^3$ with periodic boundary conditions.

Point positions: $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_N)$ for $\mathbf{x}_1, \ldots, \mathbf{x}_N \in \Lambda$.

Hamiltonian, where $T = 1/\beta$ and $r_\ell(\pi)$ is the number of ℓ-cycles in π:

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

- The first term discourages long permutation jumps, moreso for higher T.
- The temperature scale factor $T/4$, not $\beta/4$, is surprising but correct for the Bose-gas derivation of the Hamiltonian.
- The second term discourages cycles of length ℓ, moreso for higher α_ℓ. These interactions are not between points, but rather between permutation jumps.

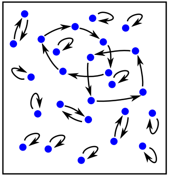

The probability model

Fixed point positions (quenched model — includes all simulations done up to the present on the lattice $N = L^3$):

$$P_{\mathbf{X}}(\pi) = \frac{1}{Y(\Lambda, \mathbf{X})} e^{-H(\mathbf{X},\pi)}, \quad Y(\Lambda, \mathbf{X}) = \sum_{\sigma \in \mathcal{S}_N} e^{-H(\mathbf{X},\sigma)}.$$

Varying positions (annealed model — many theoretical results are available):

$$P(\pi) = \frac{1}{Z(\Lambda, N)} e^{-H(\mathbf{X},\pi)}, \quad Z(\Lambda, N) = \frac{1}{N!} \int_{\Lambda^N} Y(\Lambda, \mathbf{X}) \, d\mathbf{X}.$$

In either case, we write the expectation of an RV as $\mathbb{E}_\pi[\theta(\pi)] = \sum_{\pi \in \mathcal{S}_N} P(\pi)\theta(\pi)$.

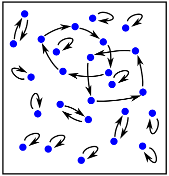

The probability model: intuition

What does a random spatial permutation actually look like? (Recall $H(\mathbf{X}, \pi) = \frac{T}{4} \sum_{i=1}^{N} \|\mathbf{x}_i - \mathbf{x}_{\pi(i)}\|^2 + \sum_{\ell=1}^{N} \alpha_\ell r_\ell(\pi)$.)

- As $T \to \infty$, the probability measure becomes supported only on the identity permutation. For large but finite T: there are tiny islands of 2-cycles, 3-cycles, etc.
- As $T \to 0$, length-dependent terms go to zero. The probability measure approaches the uniform distribution on $\mathcal{S}_N$: all π's are equally likely.

For intermediate T, things get more interesting:

- The length of each permutation jump, $\|\pi(\mathbf{x}) - \mathbf{x}\|$, remains small.
- For T above a critical temperature T_c, all cycles are short: 2-cycles, 3-cycles, etc. $T_c \approx 6.8$, and positive α terms increase T_c.
- Phase transition at T_c: for $T < T_c$ jump lengths remain short but *long cycles form*.
- Figures: high T, medium but subcritical T, and low T.

Quantifying the onset of long cycles

We observe the following:

- For $T > T_c$, $\mathbb{E}[\ell_{\max}]$ is constant as $N \to \infty$: cycles remain finite.
- For $T < T_c$, $\mathbb{E}[\ell_{\max}]$ scales with N: there are arbitrarily long cycles, or infinite cycles, in the infinite-volume limit. Feynman (1953) studied long cycles in the context of Bose-Einstein condensation for interacting systems. See also Sütő (1993, 2002).

Other random variables ("order parameters") besides $\mathbb{E}[\ell_{\max}/N]$:

- Fraction of sites in long cycles, f_I, goes to zero in L above T_c, non-zero below.
- Correlation lengths $\xi(T)$ which are (spatial or hop-count) length of the cycle containing the origin: for $T < T_c$, these blow up in L.
- Winding numbers: number of x, y, z wraps around the 3-torus (Λ with p.b.c.). Scaled winding number: $f_S = \frac{\langle \mathbf{W}^2 \rangle L^2}{3\beta N}$. This behaves much like f_I, but is easier to compute with. Also, f_W: fraction of sites which participate in winding cycles.

Central goal of my dissertation work: quantify the dependence of T_c on α, where $\Delta T_c(\alpha) = \frac{T_c(\alpha) - T_c(0)}{T_c(0)}$. Known results and conjectures are formulated quantitatively in terms of $\lim_{\alpha \to 0} \Delta T_c(\alpha)$.

Behavior of order parameters as functions of L and T ($\alpha_\ell \equiv 0$)

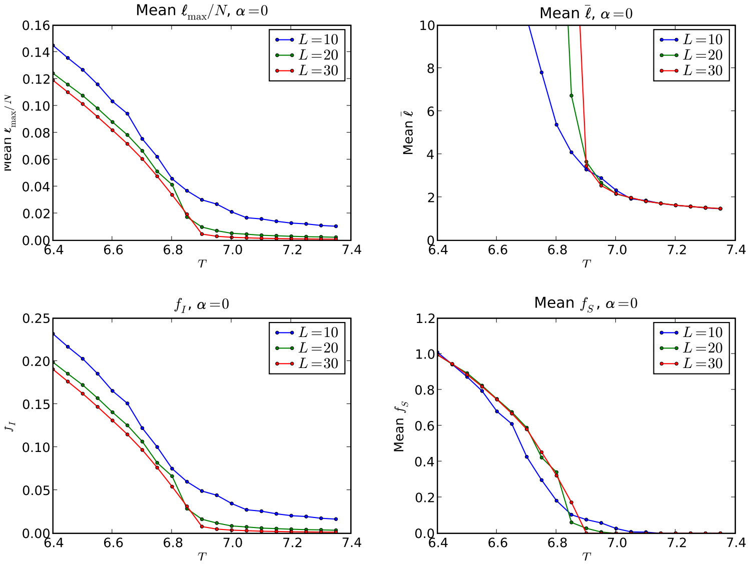

Known results and conjectures

Recall $H(\mathbf{X}, \pi) = \frac{T}{4} \sum_{i=1}^{N} \|\mathbf{x}_i - \mathbf{x}_{\pi(i)}\|^2 + \sum_{\ell=1}^{N} \alpha_\ell r_\ell(\pi)$. We have the following models:

- Non-interacting model: $\alpha_\ell \equiv 0$.
- Two-cycle model: $\alpha_2 = \alpha$ and other cycle weights are zero.
- Ewens model: α_ℓ is constant in ℓ.
- General-cycle model: No restrictions on α_ℓ.

Known results for the continuum (obtained largely using Fourier methods):

- $\Delta T_c(\alpha)$ is known (to first order in α) for two-cycle interactions (Betz and Ueltschi, CMP 2008) and small cycle weights (Betz and Ueltschi 2008). (This taps into a long and controversial history in the physics literature: see Baym et al., EJP B 2001, or Seiringer and Ueltschi, PRB 2009, for surveys.) The critical (ρ, T, α) manifold relates ρ_c to T_c.

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

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

Known results and conjectures

Known results (continued):

- $\langle \ell_{\max} \rangle / N f_I$ is constant for $T < T_c$ for $\alpha_\ell \equiv 0$. (That is, the two order parameters f_I and $\langle \ell_{\max} \rangle / N$ have the same critical exponent.) For uniform-random permutations (Shepp and Lloyd 1966 solved Golomb 1964), $\langle \ell_{\max} \rangle / N \approx 0.6243$; unpublished work of Betz and Ueltschi has found $\langle \ell_{\max} \rangle / N f_I$ is that same number for the non-interacting case $\alpha_\ell \equiv 0$. Intuition: long cycles are "uniformly distributed" within the zero Fourier mode.

Conjectures:

- $\langle \ell_{\max} \rangle / N f_I$ is constant for $T < T_c$ for all interaction models. Questions: Why should this be true on the lattice? How does that constant depend on α?
- $\xi(T)$ is monotone in T: currently unproved either for the continuum or the lattice.
- $\rho_c(\alpha)$ formula holds not only for small cycle weights ($\alpha_\ell \to 0$ faster than $1/\log \ell$).

Open questions:

- To what extent does the $\rho_c(\alpha)$ formula hold true on the lattice?
- $\Delta T_c(\alpha)$ on the lattice should be similar to that on the continuum.
- $\Delta T_c(\alpha)$ is theoretically unknown for Ewens interactions (continuum or lattice).

Metropolis sampling

The expectation of a random variable θ (e.g. $\ell_{\max}/N$, f_I, f_S, ξ) is

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

The number of permutations, $N!$, grows intractably in N. The expectation is instead estimated by summing over some number M (10^4 to 10^6) typical permutations.

The usual technical issues of Markov chain Monte Carlo (MCMC) methods are known and handled in my simulations and dissertation: thermalization time, proofs of detailed balance, autocorrelation, batched means, and quantification of variance of samples.

Metropolis step (analogue of single spin-flips for the Ising model): swap permutation arrows which end at nearest-neighbor lattice sites. This either splits a common cycle, or merges disjoint cycles:

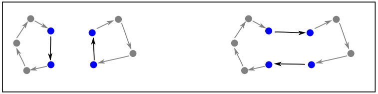

As usual, the proposed change is accepted with probability $\min\{1, e^{-\Delta H}\}$.

Metropolis sampling and winding numbers: the GKU algorithm

- Figure part 1: a long cycle on the torus almost meets itself in the x direction.
- Part 2: after a Metropolis step, one cycle winds by $+1$, and the other by -1. Metropolis steps create winding cycles only in opposite-direction pairs; total $W_x(\pi)$ is still zero.
- Part 3: if we reverse one cycle (zero-energy move), $W_x(\pi)$ is now 2.

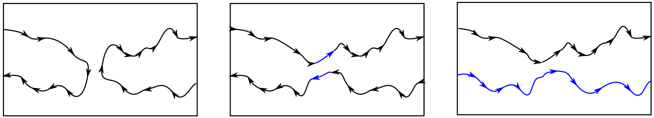

Our current best algorithm (GKU) has two types of sweeps: (1) For each lattice site, do a Metropolis step as above (Gandolfo, K). (2) For each cycle in the permutation, reverse the direction of the cycle with probability $1/2$ (Ueltschi). This permits winding numbers of even parity in each of the three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

Methods for obtaining winding numbers of all parities: try (so far with mixed success) to adapt non-local updates (e.g. Swendsen-Wang for Ising) and worm algorithm. Problems with low acceptance rate and stopping time for worm closure, respectively.

Computational results: ΔT_c

These results are preliminary. For fixed L, one may sandwich $T_c(L)$ between the vertical asymptotes of $1/f_S$ and ξ. From such graphs, we obtain, for $L = 40$, with points on the lattice,

- $\Delta T_c(L)/\alpha = 0.0759 \pm 15\%$ for the r_2 model (vs. 0.088 theoretically for the continuum), and
- $\Delta T_c(L)/\alpha = 0.483 \pm 10\%$ for the Ewens model (theoretical value is unknown, but small-cycle-weight prediction for the continuum is 0.66).

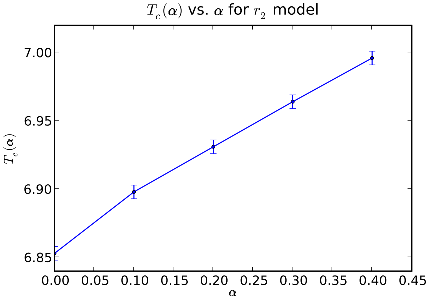

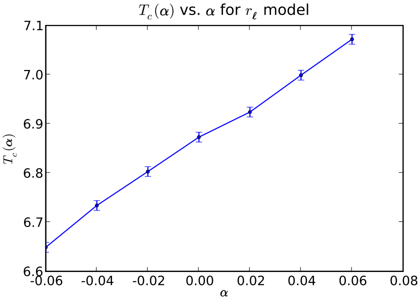

The GRU quotient varies with α in the Ewens model, but not in the r_2 model. For small L, it is non-constant for $T < T_c$; this bias seems to disappear as $L \to \infty$. (Needs a statistical confidence test.)

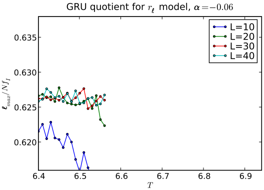

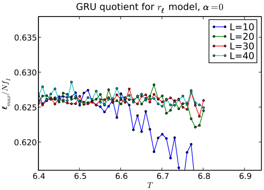

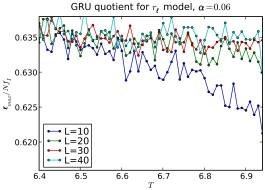

For r_2, GRU quotient is ≈ 0.626 regardless of α. For Ewens, averaging at all subcritical T's, we get the following dependence on α. This merits theoretical investigation.

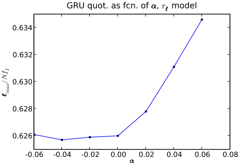

Future work

Theory:

- Prove monotonicity of $\xi(T)$ for points on the continuum.
- Find theoretical expectations for the GRU quotient $\langle \ell_{\max} \rangle / N f_I$, as a function of α, on the continuum. Empirically, we know that there are negative-α and positive-α regimes with different α-dependence.

Experiment:

- Apply more careful finite-size scaling to simulation data. (Hallway note: I would be delighted to discuss finite-size scaling with a practitioner.)
- Conduct simulations with off-lattice quenched positions (Poisson point process). Lebowitz, Lenci, and Spohn 2000 showed that the point distribution for the Bose gas is not Poisson. Yet, this is a step away from the lattice and toward the true point distribution.
- Conduct simulations with varying (annealed) point positions on the continuum. This samples from the true point distribution. Software efficiency (namely, finding which points are near to which) requires a hierarchical partitioning of Λ.
- Develop an algorithm to permit odd winding numbers.

Vielen Dank für Ihre Aufmerksamkeit!

Thank you for attending!