
Parallel tempering on optimized paths

Saifuddin Syed^{*1} Vittorio Romaniello^{*1} Trevor Campbell¹ Alexandre Bouchard-Côté¹

Abstract

Parallel tempering (PT) is a class of Markov chain Monte Carlo algorithms that constructs a path of distributions annealing between a tractable reference and an intractable target, and then interchanges states along the path to improve mixing in the target. The performance of PT depends on how quickly a sample from the reference distribution makes its way to the target, which in turn depends on the particular path of annealing distributions. However, past work on PT has used only simple paths constructed from convex combinations of the reference and target log-densities. This paper begins by demonstrating that this path performs poorly in the setting where the reference and target are nearly mutually singular. To address this issue, we expand the framework of PT to general families of paths, formulate the choice of path as an optimization problem that admits tractable gradient estimates, and propose a flexible new family of spline interpolation paths for use in practice. Theoretical and empirical results both demonstrate that our proposed methodology breaks previously-established upper performance limits for traditional paths.

1. Introduction

Markov Chain Monte Carlo (MCMC) methods are widely used to approximate intractable expectations with respect to un-normalized probability distributions over general state spaces. For hard problems, MCMC can suffer from poor mixing. For example, faced with well-separated modes, MCMC methods often get trapped exploring local regions of high probability. Parallel tempering (PT) is a widely applicable methodology (Geyer, 1991) to tackle poor mixing of MCMC algorithms.

^{*}Equal contribution ¹Department of Statistics, University of British Columbia, Vancouver, Canada. Correspondence to: Saifuddin Syed <saif.syed@stat.ubc.ca>, Vittorio Romaniello <vittorio.romaniello@stat.ubc.ca>.

Suppose we seek to approximate an expectation with respect to an intractable *target density* π_1 . Denote by π_0 a *reference density* defined on the same space, which is assumed to be tractable in the sense of the availability of an efficient sampler. This work is motivated by the case where π_0 and π_1 are nearly mutually singular. A typical case is where the target is a Bayesian posterior distribution, the reference is the prior—for which i.i.d. sampling is typically possible—and the prior is misspecified.

PT methods are based on a specific continuum of densities $\pi_t \propto \pi_0^{1-t} \pi_1^t$, $t \in [0, 1]$, bridging π_0 and π_1 . This path of intermediate distributions is known as the power posterior path in the literature, but in our framework it will be more natural to think of these continua as a *linear paths*, as they linearly interpolate between log-densities. PT algorithms discretize the path at some $0 = t_0 < \dots < t_N = 1$ to obtain a sequence of densities $\pi_{t_0}, \pi_{t_1}, \dots, \pi_{t_N}$. See Figure 1 (top) for an example of a linear path for two nearly mutually singular Gaussian distributions.

Given the path discretization, PT involves running $N + 1$ MCMC chains that together target the product distribution $\pi_{t_0} \pi_{t_1} \dots \pi_{t_N}$. Based on the assumption that the chain π_0 can be sampled efficiently, PT uses swap-based interactions between neighbouring chains to propagate the exploration done in π_0 into improved exploration in the chain of interest π_1 . By designing these swaps as Metropolis–Hastings moves, PT guarantees that the marginal distribution of the N^{th} chain converges to π_1 ; and in practice, the rate of convergence is often much faster compared to running a single chain (Woodard et al., 2009). PT algorithms are extensively used in hard sampling problems arising in statistics, physics, computational chemistry, phylogenetics, and machine learning (Desjardins et al., 2014; Ballnus et al., 2017; Kamberaj, 2020; Müller & Bouckaert, 2020).

Notwithstanding empirical and theoretical successes, existing PT algorithms also have well-understood theoretical limitations. Earlier work focusing on the theoretical analysis of reversible variants of PT has shown that adding too many intermediate chains can actually deteriorate performance (Lingenheil et al., 2009; Atchadé et al., 2011). Recent work (Syed et al., 2019) has shown that a nonreversible variant of PT (Okabe et al., 2001) is guaranteed to dominate its classical reversible counterpart, and moreover that in the

nonreversible regime adding more chains does not lead to performance collapse. However, even with these more efficient non reversible PT algorithms, Syed et al. (2019) established that the improvement brought by higher parallelism will asymptote to a fundamental limit known as the *global communication barrier*.

In this work, we show that by generalizing the class of paths interpolating between π_0 and π_1 from linear to nonlinear, the global communication barrier can be broken, leading to substantial performance improvements. Importantly, the nonlinear path used to demonstrate this breakage is computed using a practical algorithm that can be used in any situation where PT is applicable. An example of a path optimized using our algorithm is shown in Figure 1 (bottom).

We also present a detailed theoretical analysis of parallel tempering algorithms based on nonlinear paths. Using this analysis we prove that the performance gains obtained by going from linear to nonlinear path PT algorithms can be arbitrarily large. Our theoretical analysis also motivates a principled objective function used to optimize over a parametric family of paths.

Literature review Beyond parallel tempering, several methods to approximate intractable integrals rely on a path of distributions from a reference to a target distribution, and there is a rich literature on the construction and optimization of nonlinear paths for annealed importance sampling type algorithms (Gelman & Meng, 1998; Rischard et al., 2018; Grosse et al., 2013; Brekelmans et al., 2020). These algorithms are highly parallel; however, for challenging problems, even when combined with adaptive step size procedures (Zhou et al., 2016) they typically suffer from particle degeneracy (Syed et al., 2019, Sec. 7.4). Moreover, these methods use different path optimization criteria which are not well motivated in the context of parallel tempering.

Some special cases of non-linear paths have been used in the PT literature (Whitfield et al., 2002; Tawn et al., 2020). Whitfield et al. (2002) construct a non-linear path inspired by the concept of Tsallis entropy, a generalization of Boltzmann-Gibbs entropy, but do not provide algorithms to optimize over this path family. The work of Tawn et al. (2020), also considers a specific example of a nonlinear path distinct from the ones explored in this paper. However, the construction of the nonlinear path in Tawn et al. (2020) requires knowledge of the location of the modes of π_1 and hence makes their algorithm less broadly applicable than standard PT.

2. Background

In this section, we provide a brief overview of parallel tempering (PT) (Geyer, 1991), as well as recent results on nonreversible communication (Okabe et al., 2001; Sakai

& Hukushima, 2016; Syed et al., 2019). Define a reference unnormalized density function π_0 for which sampling is tractable, and an unnormalized target density function π_1 for which sampling is intractable; the goal is to obtain samples from π_1 .¹ Define a path of distributions $\pi_t \propto \pi_0^{1-t} \pi_1^t$ for $t \in [0, 1]$ from the reference to the target. Finally, define the annealing schedule \mathcal{T}_N to be a monotone sequence in $[0, 1]$, satisfying

$$\begin{aligned} \mathcal{T}_N &= (t_n)_{n=0}^N, \quad 0 = t_0 \leq t_1 \leq \dots \leq t_N = 1 \\ \|\mathcal{T}_N\| &= \max_{n \in \{0, \dots, N-1\}} t_{n+1} - t_n. \end{aligned}$$

The core idea of parallel tempering is to construct a Markov chain (X_m^0, \dots, X_m^N) , $m = 1, 2, \dots$ that (1) has invariant distribution $\pi_{t_0} \cdot \pi_{t_1} \cdot \dots \cdot \pi_{t_N}$ —such that we can treat the marginal chain X_n^N as samples from the target π_1 —and (2) swaps components of the state vector such that independent samples from component 0 (i.e., the reference π_0) traverse along the annealing path and aid mixing in component N (i.e., the target π_1). This is possible to achieve by iteratively performing a *local exploration* move followed by a *communication* move as shown in Algorithm 1.

Local Exploration Given $(X_{m-1}^0, \dots, X_{m-1}^N)$, we obtain an intermediate state $(\tilde{X}_m^0, \dots, \tilde{X}_m^N)$ by updating the n^{th} component using any MCMC move targeting π_{t_n} , for $n = 0, \dots, N$. This move can be performed in parallel across components since each is updated independently.

Communication Given the intermediate state $(\tilde{X}_m^0, \dots, \tilde{X}_m^N)$, we apply pairwise swaps of components n and $n+1$, $n \in S_m \subset \{0, \dots, N-1\}$ for swapped index set S_m . Formally, a swap is a move from (x^0, \dots, x^N) to $(x^0, \dots, x^{n+1}, x^n, \dots, x^N)$, which is accepted with probability

$$\alpha_n = 1 \wedge \frac{\pi_{t_n}(x^{n+1})\pi_{t_{n+1}}(x^n)}{\pi_{t_n}(x^n)\pi_{t_{n+1}}(x^{n+1})}. \quad (1)$$

Since each swap only depends on components $n, n+1$, one can perform all of the swaps in S_m in parallel, as long as $n \in S_m$ implies $(n+1) \notin S_m$. The largest collection of such non-interfering swaps is $S_m \in \{S_{\text{even}}, S_{\text{odd}}\}$, where $S_{\text{even}}, S_{\text{odd}}$ are the even and odd subsets of $\{0, \dots, N-1\}$ respectively. In *non-reversible PT* (NRPT) (Okabe et al., 2001), the swap set S_m at each step m is set to

$$S_m = \begin{cases} S_{\text{even}} & \text{if } m \text{ is even} \\ S_{\text{odd}} & \text{if } m \text{ is odd.} \end{cases}$$

Round trips The performance of PT is sensitive to both the local exploration and communication moves. The quantity commonly used to evaluate the performance of MCMC

¹We assume all distributions share a common state space \mathcal{X} throughout, and will often suppress the arguments of (log-)density functions—i.e., π_1 instead of $\pi_1(x)$ —for notational brevity.

algorithms is the effective sample size (ESS); however, ESS measures the combined performance of local exploration and communication, and is not able to distinguish between the two. Since the major difference between PT and standard MCMC is the presence of a communication step, we require a way to measure communication performance in isolation such that we can compare PT methods without dependence on the details of the local exploration move. The *round trip rate* is a performance measure from the PT literature (Katzgraber et al., 2006; Lingenheil et al., 2009) that is designed to assess communication efficiency alone. We say a *round trip* has occurred when a new sample from the reference π_0 travels to the target π_1 and then back to π_0 ; the round trip rate $\tau(\mathcal{T}_N)$ is the frequency at which round trips occur. Based on simplifying assumptions on the local exploration moves, the round trip rate $\tau(\mathcal{T}_N)$ may be expressed as (Syed et al., 2019, Section 3.5)

$$\tau(\mathcal{T}_N) = \left(2 + 2 \sum_{n=0}^{N-1} \frac{r(t_n, t_{n+1})}{1 - r(t_n, t_{n+1})} \right)^{-1}, \quad (2)$$

where $r(t, t')$ is the expected probability of rejection between chains $t, t' \in [0, 1]$,

$$r(t, t') = \mathbb{E} \left[1 \wedge \frac{\pi_t(X') \pi_{t'}(X)}{\pi_t(X) \pi_{t'}(X')} \right],$$

and X, X' have distributions $X \sim \pi_t, X' \sim \pi_{t'}$. Further, if \mathcal{T}_N is refined so that $\|\mathcal{T}_N\| \rightarrow 0$ as $N \rightarrow \infty$, we find that the asymptotic (in N) round trip rate is

$$\tau_\infty = \lim_{N \rightarrow \infty} \tau(\mathcal{T}_N) = (2 + 2\Lambda)^{-1}, \quad (3)$$

where $\Lambda \geq 0$ is a constant associated with the pair π_0, π_1 called the *global communication barrier* (Syed et al., 2019). Note that Λ does not depend on the number of chains N or discretization schedule \mathcal{T}_N .

3. General annealing paths

In the following, we use the terminology *annealing path* to describe a continuum of distributions interpolating between π_0 and π_1 ; this definition will be formalized shortly. The previous work reviewed in the last section assumes that the annealing path has the form $\pi_t \propto \pi_0^{1-t} \pi_1^t$, i.e., that the annealing path linearly interpolates between the log densities. A natural question is whether using other paths could lead to an improved round trip rate.

In this work we show that the answer to this question is positive. The following proposition demonstrates that the traditional path $\pi_t \propto \pi_0^{1-t} \pi_1^t$ suffers from an arbitrarily suboptimal global communication barrier even in simple examples with Gaussian reference and target distributions.

Algorithm 1 NRPT

Require: state \mathbf{x}_0 , path π_t , schedule \mathcal{T}_N , # iterations M
 $r_n \leftarrow 0$ for all $n \in \{0, \dots, N-1\}$
for $m = 1$ **to** M **do**
 $\tilde{\mathbf{x}}_m \leftarrow \text{LocalExploration}(\mathbf{x}_{m-1})$
 $S_m \leftarrow S_{\text{even}}$ if m is even, otherwise $S_m \leftarrow S_{\text{odd}}$
for $n = 0$ **to** $N-1$ **do**
 $\alpha_n \leftarrow 1 \wedge \frac{\pi_{t_n}(x^{n+1}) \pi_{t_{n+1}}(x^n)}{\pi_{t_n}(x^n) \pi_{t_{n+1}}(x^{n+1})}$
 $r_n \leftarrow r_n + (1 - \alpha_n)$
 $U_n \sim \text{Unif}(0, 1)$
if $n \in S_m$ **and** $U_n \leq \alpha_n$ **then**
 $(\tilde{x}_m^n, \tilde{x}_m^{n+1}) \leftarrow (\tilde{x}_m^{n+1}, \tilde{x}_m^n)$
end if
 $\mathbf{x}_m \leftarrow \tilde{\mathbf{x}}_m$
end for
end for
 $r_n \leftarrow r_n/M$ for $n \in \{0, \dots, N-1\}$
Return: $\{\mathbf{x}_m\}_{m=1}^M, \{r_n\}_{n=0}^{N-1}$




Figure 1. Two annealing paths between a $\pi_0 = \mathcal{N}(-2, 0.2^2)$ (light blue) and $\pi_1 = \mathcal{N}(2, 0.2^2)$ (dark blue): the traditional linear path (top) and an optimized nonlinear path (bottom). While the distributions in the linear path are nearly mutually singular, those in the optimized path overlap substantially, leading to faster round trips.

Proposition 1. *Suppose the reference and target distributions are $\pi_0 = \mathcal{N}(\mu_0, \sigma^2)$ and $\pi_1 = \mathcal{N}(\mu_1, \sigma^2)$, and define $z = |\mu_1 - \mu_0|/\sigma$. Then as $z \rightarrow \infty$,*

1. *the path $\pi_t \propto \pi_0^{1-t} \pi_1^t$ has $\tau_\infty = \Theta(1/z)$, and*
2. *there exists a path of Gaussians distributions with $\tau_\infty = \Omega(1/\log z)$.*

Therefore, upon decreasing the variance of the reference and target while holding their means fixed, the traditional linear annealing path obtains an exponentially smaller asymptotic

round trip rate than the optimal path of Gaussian distributions. Figure 1 provides an intuitive explanation. The standard path (top) corresponds to a set of Gaussian distributions with mean interpolated between the reference and target. If one reduces the variance of the reference and target, so does the variance of the distributions along the path. For any fixed N , these distributions become nearly mutually singular, leading to arbitrarily low round trip rates. The solution to this issue (bottom) is to allow the distributions along the path to have increased variances, thereby maintaining mutual overlap and the ability to swap components with a reasonable probability. This motivates the need to design more general annealing paths. In the following, we introduce the precise general definition of an annealing path, an analysis of path communication efficiency in parallel tempering, and a rigorous formulation of—and solution to—the problem of tuning path parameters to maximize communication efficiency.

3.1. Assumptions

Let $\mathcal{P}(\mathcal{X})$ be the set of probability densities with full support on a state space \mathcal{X} . For any collection of densities $\pi_t \in \mathcal{P}(\mathcal{X})$ with index $t \in [0, 1]$, associate to each a log-density function W_t such that

$$\pi_t(x) = \frac{1}{Z_t} \exp(W_t(x)), \quad x \in \mathcal{X}, \quad (4)$$

where $Z_t = \int_{\mathcal{X}} \exp(W_t(x)) dx$ is the normalizing constant. Definition 1 provides the conditions necessary to form a path of distributions from a reference π_0 to a target π_1 that are well-behaved for use in parallel tempering.

Definition 1. *An annealing path is a map $\pi_{(\cdot)} : [0, 1] \rightarrow \mathcal{P}(\mathcal{X})$, denoted $t \mapsto \pi_t$, such that for all $x \in \mathcal{X}$, $\pi_t(x)$ is continuous in t .*

There are many ways to move beyond the standard linear path $\pi_t \propto \pi_0^{1-t} \pi_1^t$. For example, consider a nonlinear path $\pi_t \propto \pi_0^{\eta_0(t)} \pi_1^{\eta_1(t)}$ where $\eta_i : [0, 1] \rightarrow \mathbb{R}$ are continuous functions such that $\eta_0(0) = \eta_1(1) = 1$ and $\eta_0(1) = \eta_1(0) = 0$. As long as for all $t \in [0, 1]$, π_t is a normalizable density this is a valid annealing path between π_0 and π_1 . Further, note that the path parameter does not necessarily have to appear as an exponent: consider for example the mixture path $\pi_t \propto (1-t)\pi_0 + t\pi_1$. Section 4 provides a more detailed example based on linear splines.

3.2. Communication efficiency analysis

Given a particular annealing path satisfying Definition 1, we require a method to characterize the round trip rate performance of parallel tempering based on that path. The results presented in this section form the basis of the objective function used to optimize over paths, as well as the foundation for the proof of Proposition 1.

We start with some notation for the rejection rates involved when Algorithm 1 is used with nonlinear paths (Equation 4). For $t, t' \in [0, 1]$, $x \in \mathcal{X}$, define the rejection rate function $r : [0, 1]^2 \rightarrow [0, 1]$ to be

$$r(t, t') = 1 - \mathbb{E}[\exp(\min\{0, A_{t,t'}(X, X')\})] \\ A_{t,t'}(x, x') = (W_{t'}(x) - W_t(x)) - (W_{t'}(x') - W_t(x')),$$

where $X \sim \pi_t$ and $X' \sim \pi_{t'}$ are independent. Assuming that all chains have reached stationarity, and all chains undergo *efficient local exploration*—i.e., $W_t(X), W_t(\tilde{X})$ are independent when $X \sim \pi_t$ and \tilde{X} is generated by local exploration from X —then the round trip rate for a particular schedule \mathcal{T}_N has the form given earlier in Equation (2). This statement follows from Syed et al. (2019, Cor. 1) without modification, because the proof of that result does not depend on the form of the acceptance ratio.

Our next objective is to characterize the asymptotic communication efficiency of a nonlinear path in the regime where $N \rightarrow \infty$ and $\|\mathcal{T}_N\| \rightarrow 0$ —which establishes its fundamental ability to take advantage of parallel computation. In other words, we require a generalization of the asymptotic result in Equation (3). In this case, previous theory relies on the particular form of the acceptance ratio for linear paths (Syed et al., 2019); in the remainder of this section, we provide a generalization of the asymptotic result for nonlinear paths by piecewise approximation by a linear spline.

For $t, t' \in [0, 1]$, define $\Lambda(t, t')$ to be the global communication barrier for the linear secant connecting π_t and $\pi_{t'}$,

$$\Lambda(t, t') = \frac{1}{2} \int_0^1 \mathbb{E}[|A_{t,t'}(X_s, X'_s)|] ds, \quad (5) \\ X_s, X'_s \stackrel{\text{i.i.d.}}{\sim} \frac{1}{Z_s(t, t')} \exp((1-s)W_t + sW_{t'}).$$

Lemma 1 shows that the global communication barrier $\Lambda(t, t')$ along the secant of the path connecting t to t' is a good approximation of the true rejection rate $r(t, t')$ with $O(|t - t'|^3)$ error as $t \rightarrow t'$.

Lemma 1. *Suppose that for all $x \in \mathcal{X}$, $W_t(x)$ is piecewise continuously differentiable in t , and that there exists $V_1 : \mathcal{X} \rightarrow [0, \infty)$ such that*

$$\forall x \in \mathcal{X}, \quad \sup_{t \in [0, 1]} \left| \frac{dW_t}{dt}(x) \right| \leq V_1(x), \quad (6)$$

and

$$\sup_{t \in [0, 1]} \mathbb{E}_{\pi_t}[V_1^3] < \infty. \quad (7)$$

Then there exists a constant $C < \infty$ independent of t, t' such that for all $t, t' \in [0, 1]$,

$$|r(t, t') - \Lambda(t, t')| \leq C|t - t'|^3. \quad (8)$$

A direct consequence of Lemma 1 is that for any fixed schedule \mathcal{T}_N ,

$$\left| \sum_{n=0}^{N-1} r(t_n, t_{n+1}) - \Lambda(\mathcal{T}_N) \right| \leq C \|\mathcal{T}_N\|^2, \quad (9)$$

where $\Lambda(\mathcal{T}_N) = \sum_{n=0}^{N-1} \Lambda(t_n, t_{n+1})$. Intuitively, in the $\|\mathcal{T}_N\| \approx 0$ regime where rejection rates are low,

$$\frac{r(t_n, t_{n+1})}{1 - r(t_n, t_{n+1})} \approx r(t_n, t_{n+1}),$$

and we have that $\tau(\mathcal{T}_N) \approx (2 + 2\Lambda(\mathcal{T}_N))^{-1}$. Therefore, $\Lambda(\mathcal{T}_N)$ characterizes the communication efficiency of the path in a way that naturally extends the global communication barrier from the linear path case. Theorem 2 provides the precise statement: the convergence is uniform in \mathcal{T}_N and depends only on $\|\mathcal{T}_N\|$, and $\Lambda(\mathcal{T}_N)$ itself converges to a constant Λ in the asymptotic regime. We refer to Λ , defined below in Equation (13) as the global communication barrier for the general annealing path.

Theorem 2. *Suppose that for all $x \in \mathcal{X}$, $W_t(x)$ is piecewise twice continuously differentiable in t , that there exists $V_1 : \mathcal{X} \rightarrow [0, \infty)$ satisfying (6) and (7), and that there exists $V_2 : \mathcal{X} \rightarrow [0, \infty)$, $\epsilon > 0$ satisfying*

$$\forall x \in \mathcal{X}, \sup_{t \in [0, 1]} \left| \frac{d^2 W_t}{dt^2}(x) \right| \leq V_2(x), \quad (10)$$

and

$$\sup_{t \in [0, 1]} \mathbb{E}_{\pi_t} [(1 + V_1) \exp(\epsilon V_2)] < \infty. \quad (11)$$

Then,

$$\lim_{\delta \rightarrow 0} \sup_{\mathcal{T}_N : \|\mathcal{T}_N\| \leq \delta} |(2 + 2\Lambda(\mathcal{T}_N))^{-1} - \tau(\mathcal{T}_N)| = 0 \quad (12)$$

$$\text{and } \lim_{\delta \rightarrow 0} \sup_{\mathcal{T}_N : \|\mathcal{T}_N\| \leq \delta} |\Lambda(\mathcal{T}_N) - \Lambda| = 0, \quad (13)$$

where $\Lambda = \int_0^1 \lambda(t) dt$ for an instantaneous rejection rate function $\lambda : [0, 1] \rightarrow [0, \infty)$ given by

$$\begin{aligned} \lambda(t) &= \lim_{\Delta t \rightarrow 0} \frac{r(t + \Delta t, t)}{|\Delta t|} \\ &= \frac{1}{2} \mathbb{E} \left[\left| \frac{dW_t}{dt}(X_t) - \frac{dW_t}{dt}(X'_t) \right| \right], \quad X_t, X'_t \stackrel{i.i.d.}{\sim} \pi_t. \end{aligned}$$

The integrability condition is required to control the tail behaviour of distributions formed by linearized approximations to the path π_t . This condition is satisfied by a wide range of annealing paths, e.g., the linear spline paths proposed in this work—since in that case $V_2 = 0$.

3.3. Annealing path families and optimization

It is often the case that there are a set of candidate annealing paths in consideration for a particular target π . For example, if a path has tunable parameters $\phi \in \Phi$ that govern its shape, we can generate a collection of annealing paths that all target π by varying the parameter ϕ . We call such collections an annealing path family.

Definition 3. *An annealing path family for target π_1 is a collection of annealing paths $\{\pi_t^\phi\}_{\phi \in \Phi}$ such that for all parameters $\phi \in \Phi$, $\pi_1^\phi = \pi_1$.*

There are many ways to construct useful annealing path families. For example, if one is provided a parametric family of variational distributions $\{q_\phi : \phi \in \Phi\}$ for some parameter space Φ , one can construct the annealing path family of linear paths $\pi_t^\phi = q_\phi^{1-t} \pi_1^t$ from a variational reference q_ϕ to the target π_1 . More generally, given $\eta_i(t)$ satisfying the constraints in Section 3.1, $\pi_t^\phi = q_\phi^{\eta_0(t)} \pi_1^{\eta_1(t)}$ defines a nonlinear annealing path family. Another example of an annealing path family used in the context of PT are q -paths $\{\pi_t^q\}_{q \in [0, 1]}$ (Whitfield et al., 2002). Given a fixed reference and target π_0, π_1 , the path π_t^q interpolates between the mixture path ($q = 0$) and the linear path ($q = 1$) (Brekelmans et al., 2020). In Section 4, we provide a new flexible class of nonlinear paths based on splines that is designed specifically to enhance the performance of parallel tempering.

Since every path in an annealing path family has the desired target distribution π_1 , we are free to optimize the path over the tuning parameter space $\phi \in \Phi$ in addition to optimizing the schedule \mathcal{T}_N .² Motivated by the analysis of Section 3.2, a natural objective function for this optimization to consider is the non-asymptotic round trip rate

$$\phi^*, \mathcal{T}_N^* = \arg \max_{\phi \in \Phi, \mathcal{T}_N} \tau^\phi(\mathcal{T}_N) \quad (14)$$

$$= \arg \min_{\phi \in \Phi, \mathcal{T}_N} \sum_{n=0}^{N-1} \frac{r^\phi(t_n, t_{n+1})}{1 - r^\phi(t_n, t_{n+1})}, \quad (15)$$

where now the round trip rate and rejection rates depend both on the schedule and path parameter, denoted by superscript ϕ . We solve this optimization using an approximate coordinate-descent procedure, iterating between an update of the schedule \mathcal{T}_N for a fixed path parameter $\phi \in \Phi$, followed by a gradient step in ϕ based on a surrogate objective function and a fixed schedule. This is summarized in Algorithm 2. We outline the details of schedule and path tuning procedure in the following.

Tuning the schedule Fix the value of ϕ , which fixes the path. We adapt a schedule tuning algorithm from past work

²We assume that the optimization over ϕ ends after a finite number of iterations to sidestep the potential pitfalls of adaptive MCMC methods (Andrieu & Moulines, 2006).

Algorithm 2 PathOptNRPT

Require: state \mathbf{x} , path family π_t^ϕ , parameter ϕ , # chains N ,
 # PT iterations M , # tuning steps S , learning rate γ
 $\mathcal{T}_N \leftarrow (0, 1/N, 2/N, \dots, 1)$
for $s = 1$ **to** S **do**
 $\{\mathbf{x}_m\}_{m=1}^M, (r_n)_{n=0}^N \leftarrow \text{NRPT}(\mathbf{x}, \pi_t^\phi, \mathcal{T}_N, M)$
 $\lambda^\phi \leftarrow \text{CommunicationBarrier}(\mathcal{T}_N, \{r_n\})$
 $\mathcal{T}_N \leftarrow \text{UpdateSchedule}(\lambda^\phi, N)$
 $\phi \leftarrow \phi - \gamma \nabla_\phi \sum_{n=0}^{N-1} \text{SKL}(\pi_{t_n}^\phi, \pi_{t_{n+1}}^\phi)$
 $\mathbf{x} \leftarrow \mathbf{x}_M$
end for
Return: ϕ, \mathcal{T}_N

to update the schedule $\mathcal{T}_N = (t_n)_{n=0}^N$ (Syed et al., 2019, Section 5.1). Based on the same argument as this previous work, we obtain that when $\|\mathcal{T}_N\| \approx 0$, the non-asymptotic round trip rate is maximized when the rejection rates are all equal. The schedule that approximately achieves this satisfies

$$\forall n \in \{1, \dots, N-1\}, \quad \frac{1}{\Lambda^\phi} \int_0^{t_n} \lambda^\phi(s) ds = \frac{n}{N}. \quad (16)$$

Following Syed et al. (2019), we use Monte Carlo estimates of the rejection rates $r^\phi(t_n, t_{n+1})$ to approximate $t \mapsto \int_0^t \lambda^\phi(s) ds$, $s \in [0, 1]$ via a monotone cubic spline, and then use bisection search to solve for each t_n according to Equation (16).

Optimizing the path Fix the schedule \mathcal{T}_N ; we now want to improve the path itself by modifying ϕ . However, in challenging problems this is not as simple as taking a gradient step for the objective in Equation (14). In particular, in early iterations—when the path is near its oft-poor initialization—the rejection rates satisfy $r^\phi(t_n, t_{n+1}) \approx 1$. As demonstrated empirically in Appendix F, gradient estimates in this regime exhibit a low signal-to-noise ratio that precludes their use for optimization.

We propose a surrogate, the symmetric KL divergence, motivated as follows. Consider first the global communication barrier $\Lambda^\phi(\mathcal{T}_N)$ for the linear spline approximation to the path; Theorem 2 guarantees that as long as $\|\mathcal{T}_N\|$ is small enough, one can optimize $\Lambda^\phi(\mathcal{T}_N)$ in place of the round trip rate $\tau^\phi(\mathcal{T}_N)$. By Jensen’s inequality,

$$\frac{1}{N^2} \Lambda^\phi(\mathcal{T}_N)^2 \leq \frac{1}{N} \sum_{n=0}^{N-1} \Lambda^\phi(t_n, t_{n+1})^2.$$

Next, we apply Jensen’s inequality again to the definition of $\Lambda^\phi(t_n, t_{n+1})$ from (5), which shows that

$$\Lambda^\phi(t_n, t_{n+1})^2 \leq \frac{1}{4} \int_0^1 \mathbb{E}[A_{t_n, t_{n+1}}(X_s, X'_s)^2] ds,$$

where X_s are defined in (5) are drawn from the linear path between $\pi_{t_n}^\phi$ and $\pi_{t_{n+1}}^\phi$. Finally, we note that the inner expectation is the path integral of the Fisher information metric along the linear path and evaluates to the symmetric KL divergence (Dabak & Johnson, 2002, Result 4),

$$\int_0^1 \mathbb{E}[(A_{t_n, t_{n+1}}(X_s, X'_s))^2] ds = 2\text{SKL}(\pi_{t_n}^\phi, \pi_{t_{n+1}}^\phi).$$

Therefore we have

$$\frac{2\Lambda^\phi(\mathcal{T}_N)^2}{N} \leq \sum_{n=0}^{N-1} \text{SKL}(\pi_{t_n}^\phi, \pi_{t_{n+1}}^\phi). \quad (17)$$

The slack in the inequality in Equation (17) could potentially depend on ϕ even in the large N regime. Therefore, during optimization, we recommend monitoring the value of the original objective function (Equation (14)) to ensure that the optimization of the surrogate SKL objective indeed improves it, and hence the round trip rate performance of PT via Equation (2). In the experiments we display the values of both objective functions.

4. Spline annealing path family

In this section, we develop a family of annealing paths—the *spline annealing path family*—that offers a practical and flexible improvement upon the traditional linear paths considered in past work. We first define a general family of annealing paths based on the exponential family, and then provide the specific details of the spline family with a discussion of its properties. Empirical results in Section 5 demonstrate that the spline annealing path family resolves the problematic Gaussian annealing example in Figure 1.

4.1. Exponential annealing path family

We begin with the practical desiderata for an annealing path family given a fixed reference π_0 and target π_1 distribution.³ First, the traditional linear path $\pi_t \propto \pi_0^{1-t} \pi_1^t$ should be a member of the family, so that one can achieve at least the round trip rate provided by that path. Second, the family should be broadly applicable and not depend on particular details of either π_0 or π_1 . Finally, using the Gaussian example from Figure 1 and Proposition 1 as insight, the family should enable the path to smoothly vary from π_0 to π_1 while inflating / deflating the variance as necessary.

These desiderata motivate the design of the *exponential annealing path family*, in which each annealing path takes the form

$$\pi_t \propto \pi_0^{\eta_0(t)} \pi_1^{\eta_1(t)} = \exp(\eta(t)^T W(x)),$$

³A natural extension of this discussion would include parametrized variational reference distribution families. For simplicity we restrict to a fixed reference.

for some function $\eta(t) = (\eta_0(t), \eta_1(t))$ and reference/target log densities $W(x) = (W_0(x), W_1(x))$. Intuitively, $\eta_0(t)$ and $\eta_1(t)$ represent the level of annealing for the reference and target respectively along the path. Proposition 2 shows that a broad collection of functions η indeed construct a valid annealing path family including the linear path.

Proposition 2. *Let $\Omega \subseteq \mathbb{R}^2$ be the set*

$$\Omega = \left\{ \xi \in \mathbb{R}^2 : \int \exp(\xi^T W(x)) dx < \infty \right\}.$$

Suppose $(0, 1) \in \Omega$ and \mathcal{A} is a set of piecewise twice continuously differentiable functions $\eta : [0, 1] \rightarrow \Omega$ such that $\eta(1) = (0, 1)$. Then

$$\{\pi_t(x) \propto \exp(\eta(t)^T W(x)) : \eta \in \mathcal{A}\}$$

is an annealing path family for target distribution π_1 . If additionally $(1, 0) \in \Omega$, then the linear path $\eta(t) = (1-t, t)$ may be included in \mathcal{A} . Finally, if for every $\eta \in \mathcal{A}$ there exists $M > 0$ such that $\sup_t \max\{\|\eta'(t)\|_2, \|\eta''(t)\|_2\} \leq M$ and Equation (11) holds with $V_1 = V_2 = M\|W\|_2$, then each path in the family satisfies the conditions of Theorem 2.

4.2. Spline annealing path family

Proposition 2 reduces the problem of designing a general family of paths of probability distributions to the much simpler task of designing paths in \mathbb{R}^2 . We argue that a good choice can be constructed using linear spline paths connecting K knots $\phi = (\phi_0, \dots, \phi_K) \in (\mathbb{R}^2)^{K+1}$, i.e., for all $k \in \{1, \dots, K\}$ and $t \in [\frac{k-1}{K}, \frac{k}{K}]$,

$$\eta^\phi(t) \mapsto (k - Kt)\phi_{k-1} + (Kt - k + 1)\phi_k.$$

Let Ω be defined as in Proposition 2. The K -knot *spline annealing path family* is defined as the set of K -knot linear spline paths such that

$$\phi_0 = (1, 0), \quad \phi_K = (0, 1), \quad \text{and} \quad \forall k, \phi_k \in \Omega.$$

Validity: Since Ω is a convex set per the proof of Proposition 2, we are guaranteed that $\eta^\phi([0, 1]) \subseteq \Omega$, and so this collection of annealing paths is a subset of the exponential annealing family and hence forms a valid annealing path family targeting π_1 .

Convexity: Furthermore, convexity of Ω implies that tuning the knots $\phi \in \Omega^{K+1}$ involves optimization within a convex constraint set. In practice, we enforce also that the knots are monotone in each component—i.e., the first component monotonically decreases, $1 = \phi_{0,0} \geq \phi_{1,0} \geq \dots \geq \phi_{K,0} = 0$ and the second increases, $0 = \phi_{0,1} \leq \phi_{1,1} \leq \dots \leq \phi_{K,1} = 1$ —such that the path of distributions always moves from the reference to the target. Because monotonicity constraint sets are linear and hence convex, the overall monotonicity-constrained optimization problem has a convex domain.




Figure 2. The spline path for $K = 1, 2, 3, 4, 5, 10$ knots for the family generated by $\pi_0 = N(-1, 0.5)$ and $\pi_1 = N(1, 0.5)$.

Flexibility: Assuming the family is nonempty, it trivially contains the linear path. Further, given a large enough number of knots K , the spline annealing family well-approximates subsets of the exponential annealing family for fixed $M > 0$. In particular,

$$\sup_{\eta \in \mathcal{A}_M} \inf_{\phi \in \Omega^{K+1}} \|\eta^\phi - \eta\|_\infty \leq \frac{M}{4K^2}. \quad (18)$$

Figure 2 provides an illustration of the behaviour of optimized spline paths for a Gaussian reference and target. The path takes a convex curved shape; starting at the bottom right point of the figure (reference), this path corresponds to increasing the variance of the reference, shifting the mean from reference to target, and finally decreasing the variance to match the target. With more knots, this process happens more smoothly.

Appendix D provides an explicit derivation of the stochastic gradient estimates we use to optimize the knots of the spline annealing path family. It is also worth making two practical notes. First, to enforce the monotonicity constraint, we developed an alternative to the usual projection approach, as projecting into the set of monotone splines can cause several knots to become superposed. Instead, we maintain monotonicity as follows: after each stochastic gradient step, we identify a monotone subsequence of knots containing the endpoints, remove the nonmonotone jumps, and linearly interpolate between the knot subsequence with an even spacing. Second, we take gradient steps in a log-transformed space so that knot components are always strictly positive.

5. Experiments

In this section, we study the empirical performance of non-reversible PT based on the spline annealing path family ($K \in \{2, 3, 4, 5, 10\}$) from Section 4, with knots and schedule optimized using the tuning method from Section 3. We compare this method to two PT methods

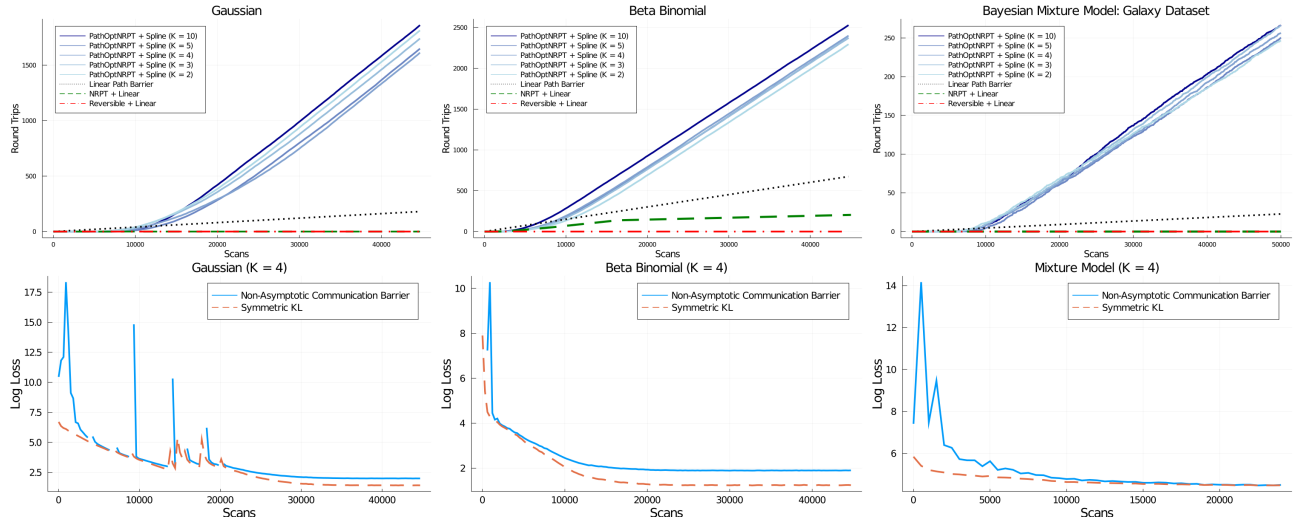


Figure 3. Top: Cumulative round trips averaged over 10 runs for the spline path with $K = 2, 3, 4, 5, 10$ (solid blue), NRPT using a linear path (dashed green), and reversible PT with linear path (dash/dot red). The slope of the lines represent the round trip rate. We observe large gains going from linear to non-linear paths ($K > 1$). For all values of $K > 1$, the optimized spline path substantially improves on the theoretical upper bound on round trip rate possible using linear path (dotted black). **Bottom:** Non-asymptotic communication barrier from Equation 15 (solid blue) and Symmetric KL (dash orange) as a function of iteration for one run of PathOptNRPT + Spline ($K = 4$ knots).

based on standard linear paths: non-reversible PT with adaptive schedule (“NRPT+Linear”) (Syed et al., 2019), and reversible PT (“Reversible+Linear”) (Atchadé et al., 2011). Code for the experiments is available at <https://github.com/vittrom/PT-pathoptim>.

We use the terminology “scan” to denote one iteration of the for loop in Algorithm 2. The computational cost of a scan is comparable for all the methods, since the bottleneck is the local exploration step shared by all methods. These experiments demonstrate two primary conclusions: (1) tuned nonlinear paths provide a substantially higher round trip rate compared to linear paths for the examples considered; and (2) the symmetric KL sum objective (Equation 17) is a good proxy for the round trip rate as a tuning objective.

We run the following benchmark problems; see the supplement for details. **Gaussian:** a synthetic setup in which the reference distribution is $\pi_0 = N(-1, 0.01^2)$ and the target is $\pi_1 = N(1, 0.01^2)$. For this example we used $N = 50$ parallel chains and fixed the computational budget to 45000 samples. For Algorithm 2, the computational budget was divided equally over 150 scans, meaning 300 samples were used for every gradient update. “Reversible+Linear” performed 45000 local exploration steps with a communication step after every iteration while for “NRPT+Linear” the computational budget was used to adapt the schedule. The gradient updates were performed using Adagrad (Duchi et al., 2011) with learn-

ing rate equal to 0.2. **Beta-binomial model:** a conjugate Bayesian model with prior $\pi_0(p) = \text{Beta}(180, 840)$ and likelihood $L(x|p) = \text{Binomial}(x|n, p)$. We simulated data $x_1, \dots, x_{2000} \sim \text{Binomial}(100, 0.7)$ resulting in a posterior $\pi_1(p) = \text{Beta}(140180, 60840)$. The prior and posterior are heavily concentrated at 0.2 and 0.7 respectively. We used the same settings as for the Gaussian example. **Galaxy data:** A Bayesian Gaussian mixture model applied to the galaxy dataset of (Roeder, 1990). We used six mixture components with mixture proportions w_0, \dots, w_5 , mixture component densities $N(\mu_i, 1)$ for mean parameters μ_0, \dots, μ_5 , and a cluster label categorical variable for each data point. We placed a $\text{Dir}(\mathbf{1}_6)$ prior on the proportions, where $\mathbf{1}_6 = (1, 1, 1, 1, 1, 1)$ and a $N(150, 1)$ prior on each of the mean parameters. We did not marginalize the cluster indicators, creating a multi-modal posterior inference problem over 94 latent variables. In this experiment we used $N = 35$ chains and fixed the computational budget to 50000 samples, divided into 500 scans using 100 samples each. We optimized the path using Adagrad with a learning rate of 0.3. Exploring the full posterior distribution is challenging in this context due to a combination of misspecification of the prior and label switching. Label switching refers to the invariance of the likelihood under relabelling of the cluster labels. In Bayesian problems label switching can lead to increased difficulty of the sampling problem as it generates symmetric multi-modal posterior distributions. **High dimensional Gaussian:** a similar setup to the




Figure 4. Round trips rate averaged over 5 runs for the spline path with $K = 4$ (solid blue), NRPT using a linear path (dashed green) and theoretical upper bound on round trip rate possible using linear path (dotted black) as a function of the dimensionality of the target distribution.

one-dimensional Gaussian experiment where the number of dimensions ranges from $d = 1$ to $d = 256$. The reference distribution is $\pi_0 = N(-\mathbf{1}_d, (0.1^2)I_d)$ and the target is $\pi_1 = N(\mathbf{1}_d, (0.1^2)I_d)$ where the subscript d indicates the dimensionality of the problem. The number of chains N is set to increase with dimension at the rate $N = \lceil 15\sqrt{d} \rceil$. We fixed the number of spline knots K to 4 and set the computational budget to 50000 samples divided into 500 scans with 100 samples per gradient update. The gradient updates were performed using Adagrad with learning rate equal to 0.2. For all the experiments we performed one local exploration step before each communication step.

The results of these experiments are shown in Figures 3 and 4. Examining the top row of Figure 3—which shows the number of round trips as a function of the number of scans—one can see that PT using the spline annealing family outperforms PT using the linear annealing path across all numbers of knots tested. Moreover, the slope of these curves demonstrates that PT with the spline annealing family exceeds the theoretical upper bound of round trip rate for the linear annealing path (from Equation (12)). The largest gain is obtained from going from $K = 1$ (linear) to $K = 2$. For all the examples, increasing the number of knots to more than $K > 2$ leads to marginal improvements. In the case of the Gaussian example, note that since the global communication barrier Λ for the linear path is much larger than N , algorithms based on linear paths incurred rejection rates of nearly one for most chains, resulting in no round trips.

The bottom row of Figure 3 shows the value of the surrogate SKL objective and non-asymptotic communication barrier from Equation (15). In particular, these figures demonstrate that the SKL provides a surrogate objective that is a reason-

able proxy for the non-asymptotic communication barrier, but does not exhibit as large estimation variance in early iterations when there are pairs of chains with rejection rates close to one.

Figure 4 shows the round trip rate as a function of the dimensionality of the problem for Gaussian target distributions. As the dimensionality increases, the sampling problem becomes fundamentally more difficult, explaining the decay in performance of both NRPT with a linear path and an optimized path. Both sampling methods are provided with a fixed computational budget in all runs. Due to this fixed budget, NRPT is unable for $d \geq 64$ to approach the optimal schedule in the allotted time. This leads to an increasing gap between the round trip rates of NRPT with linear path and the spline path for $d \geq 64$.

6. Discussion

In this work, we identified the use of linear paths of distributions as a major bottleneck in the performance of parallel tempering algorithms. To address this limitation, we have provided a theory of parallel tempering on nonlinear paths, a methodology to tune parametrized paths, and finally a practical, flexible family of paths based on linear splines. Future work in this line of research includes extensions to estimate normalization constants, as well as the development of techniques and theory surrounding the use of variational reference distributions.

References

- Andrieu, C. and Moulines, E. On the ergodicity properties of some adaptive MCMC algorithms. *Annals of Applied Probability*, 16(3):1462–1505, 2006.
- Atchadé, Y. F., Roberts, G. O., and Rosenthal, J. S. Towards optimal scaling of Metropolis-coupled Markov chain Monte Carlo. *Statistics and Computing*, 21(4): 555–568, 2011.
- Ballnus, B., Hug, S., Hatz, K., Görlitz, L., Hasenauer, J., and Theis, F. J. Comprehensive benchmarking of Markov chain Monte Carlo methods for dynamical systems. *BMC Systems Biology*, 11(1):63, 2017.
- Brekelmans, R., Masrani, V., Bui, T., Wood, F., Galstyan, A., Steeg, G. V., and Nielsen, F. Annealed importance sampling with q-paths. *arXiv:2012.07823*, 2020.
- Costa, S. I., Santos, S. A., and Strapasson, J. E. Fisher information distance: A geometrical reading. *Discrete Applied Mathematics*, 197:59–69, 2015.
- Dabak, A. G. and Johnson, D. H. Relations between Kullback-Leibler distance and Fisher information. *Journal of The Iranian Statistical Society*, 5:25–37, 2002.

- Desjardins, G., Luo, H., Courville, A., and Bengio, Y. Deep tempering. *arXiv:1410.0123*, 2014.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.
- Gelman, A. and Meng, X.-L. Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical science*, pp. 163–185, 1998.
- Geyer, C. J. Markov chain Monte Carlo maximum likelihood. *Interface Proceedings*, 1991.
- Grosse, R. B., Maddison, C. J., and Salakhutdinov, R. R. Annealing between distributions by averaging moments. In *Advances in Neural Information Processing Systems*, 2013.
- Kamberaj, H. *Molecular Dynamics Simulations in Statistical Physics: Theory and Applications*. Springer, 2020.
- Katzgraber, H. G., Trebst, S., Huse, D. A., and Troyer, M. Feedback-optimized parallel tempering Monte Carlo. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(03):P03018, 2006.
- Lingenheil, M., Denschlag, R., Mathias, G., and Tavan, P. Efficiency of exchange schemes in replica exchange. *Chemical Physics Letters*, 478(1-3):80–84, 2009.
- Müller, N. F. and Bouckaert, R. R. Adaptive parallel tempering for BEAST 2. *bioRxiv*, 2020. doi: 10.1101/603514.
- Okabe, T., Kawata, M., Okamoto, Y., and Mikami, M. Replica-exchange Monte Carlo method for the isobaric–isothermal ensemble. *Chemical Physics Letters*, 335(5-6): 435–439, 2001.
- Predescu, C., Predescu, M., and Ciobanu, C. V. The incomplete beta function law for parallel tempering sampling of classical canonical systems. *The Journal of Chemical Physics*, 120(9):4119–4128, 2004.
- Rainforth, T., Kosiorek, A. R., Le, T. A., Maddison, C. J., Igl, M., Wood, F., and Teh, Y. W. Tighter variational bounds are not necessarily better. In *International Conference on Machine Learning*, 2018.
- Rischar, M., Jacob, P. E., and Pillai, N. Unbiased estimation of log normalizing constants with applications to Bayesian cross-validation. *arXiv:1810.01382*, 2018.
- Roeder, K. Density estimation with confidence sets exemplified by superclusters and voids in the galaxies. *Journal of the American Statistical Association*, 85(411):617–624, 1990.
- Sakai, Y. and Hukushima, K. Irreversible simulated tempering. *Journal of the Physical Society of Japan*, 85(10): 104002, 2016.
- Syed, S., Bouchard-Côté, A., Deligiannidis, G., and Doucet, A. Non-reversible parallel tempering: an embarrassingly parallel MCMC scheme. *arXiv:1905.02939*, 2019.
- Tawn, N. G., Roberts, G. O., and Rosenthal, J. S. Weight-preserving simulated tempering. *Statistics and Computing*, 30(1):27–41, 2020.
- Whitfield, T., Bu, L., and Straub, J. Generalized parallel sampling. *Physica A: Statistical Mechanics and its Applications*, 305(1-2):157–171, 2002.
- Woodard, D. B., Schmidler, S. C., Huber, M., et al. Conditions for rapid mixing of parallel and simulated tempering on multimodal distributions. *The Annals of Applied Probability*, 19(2):617–640, 2009.
- Zhou, Y., Johansen, A. M., and Aston, J. A. Toward automatic model comparison: An adaptive sequential Monte Carlo approach. *Journal of Computational and Graphical Statistics*, 25(3):701–726, 2016.