

Contribution ID: 26

Type: **not specified**

# Accelerating Parallel Tempering with Neural Transports

*Thursday 4 December 2025 14:40 (40 minutes)*

Markov Chain Monte Carlo (MCMC) is a powerful algorithmic framework for sampling from complex probability distributions. Standard MCMC methods struggle with high-dimensional distributions containing well-separated modes, becoming trapped in local regions. Parallel tempering (PT) addresses this by using intermediate annealing distributions to bridge a tractable reference (e.g., Gaussian) and an intractable target distribution. However, classical PT is inflexible, fragile, challenging to tune, and suffers from performance collapse for challenging inference tasks.

This talk introduces non-reversible parallel tempering (NRPT), which provably dominates classical reversible PT algorithms. We show that NRPT undergoes a sharp algorithmic phase transition with increased parallelism, where it becomes robust, easy to tune, and scales efficiently to GPUs. I will then demonstrate how to further accelerate PT using neural transports such as normalising flows and diffusions. We demonstrate this framework across a variety of examples in Bayesian inference and inference-time control for diffusion models, and discuss recent applications to cancer genomics and the Event Horizon Telescope.

## Special requests

**Author:** SYED, Saifuddin (University of British Columbia)

**Presenter:** SYED, Saifuddin (University of British Columbia)

**Session Classification:** Talks