

Transport and Multilevel Approaches for Large-Scale PDE-Constrained Bayesian Inference

Robert Scheichl, Heidelberg University

In this talk, I will give an introduction to multilevel MCMC (Dodwell et al, JUQ, 2015) for large-scale PDE-constrained Bayesian inference and then present a couple of approaches to "preconditioning" this sampler to improve its efficiency. In particular, I will present a non-trivial integration with dimension-independent likelihood-informed (DILI) MCMC (Cui, Law, Marzouk, JCP, 2016), which offers several advantages: First, DILI-MCMC employs an intrinsic likelihood-informed subspace (LIS) (Cui et al., Inverse Prob, 2014) -- which involves a number of forward and adjoint model simulations -- to design accelerated operator-weighted proposals. By exploiting the multilevel structure of the discretised parameters and discretised forward models, we design a Rayleigh-Ritz procedure to significantly reduce the computational effort in building the LIS and operating with DILI proposals. Second, the resulting DILI-MCMC can drastically improve the sampling efficiency of MCMC at each level, and hence reduce the integration error of the multilevel algorithm for fixed CPU time. To be able to fully exploit the power of multilevel MCMC and to reduce the dependencies of samples on different levels for a parallel implementation, we also suggest a new pooling strategy for allocating computational resources across different levels and constructing Markov chains at higher levels conditioned on those simulated on lower levels. Numerical results confirm the improved computational efficiency of the multilevel DILI approach. Time permitting, I will also briefly mention an alternative approach based on low-rank tensor approximation to construct an efficient, scalable surrogate of the posterior distribution in high dimensions, the Tensor Train Conditional Distribution (TT-CD) sampler (Dolgov et al., Stats & Comput, 2020).