DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Reference herein to any social initiative (including but not limited to Diversity, Equity, and Inclusion (DEI); Community Benefits Plans (CBP); Justice 40; etc.) is made by the Author independent of any current requirement by the United States Government and does not constitute or imply endorsement, recommendation, or support by the United States Government or any agency thereof.

LLNL-TR-2005705

Lawrence Livermore National Laboratory

Scalable Algorithms for Inverse Problems With High-Dimensional Parameter Spaces

H Fairbanks

June 2025

Disclaimer

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Title. Scalable Algorithms for Inverse Problems With High-Dimensional Parameter Spaces

Author. Hillary R. Fairbanks, Lawrence Livermore National Laboratory (fairbanks5@llnl.gov)

Topic. Inverse problems, which involve inferring unknown parameters from observed data, present significant computational challenges, especially in large-scale settings with high-dimensional unknown parameters and nonlinear relationships between the unknowns and observations. Bayesian inference provides an approach for addressing these problems, often relying on sequential sampling methods like Markov chain Monte Carlo (MCMC) to approximate the posterior distribution of the parameters. However, MCMC methods become computationally demanding as the dimensionality of the problem increases, particularly in large-scale systems where likelihood evaluations rely on solving partial differential equations (PDEs) on large spatial domains with finely resolved meshes. To overcome these limitations, recent advancements have focused on designing scalable computational techniques – for both PDE simulations and sampling strategies – to make Bayesian methods feasible for high-dimensional problems.

Challenge. To achieve scalable approaches for Bayesian inference, it is essential to develop efficient algorithms for both the inner and outer loop components. In this context, the outer loop involves sampling the high-dimensional parameter space, while the inner loop focuses on solving the PDE using the sampled parameter as input. Depending on the applications, several algorithms are available to perform scalable inner loop PDE solves, most notably multigrid methods [2]; furthermore, high performance computing software can be leveraged for distributed computing, e.g., [1, 11]. The inner loop, on the other hand, requires intelligent sampling strategies to efficiently explore the parameter space. This becomes more challenging in high-dimensional settings, where the posterior becomes more concentrated in small regions, making it harder to sample from efficiently. To help alleviate this issue, sequential approaches such as MCMC can utilize two forms of parallelism to improve scaling: each PDE solve may be performed in parallel using distributed computing, and multiple MCMC chains may be performed in parallel. But given the sequential nature of MCMC, individual simulations within a Markov chain cannot be executed in parallel.

Opportunity. In this high-dimensional, nonlinear setting, effective sampling strategies are necessary to improve scaling of Bayesian inference. These outer loop methods must intelligently sample from a large parameter space to converge to regions of high density with respect to the posterior distribution. Furthermore, the approaches used to *intelligently sample* – or, in the context of MCMC, propose samples – must not incur significant computational overhead. Ideally, no additional solvers will be needed; however, if they are, the computational benefits should be significantly justified. Even better, if these sequential algorithms can integrate some form of parallel simulation without substantially increasing communication overhead, it would greatly enhance scalability.

Innovation. Over the past couple of decades, several methods have been developed to accelerate MCMC mixing by modifying the proposal distribution. Notable approaches include those which utilize approximations of the Hessian and gradient to modify the MCMC proposal [13, 14, 6], which significantly reduce the number of PDE simulations compared to MCMC; though, these approaches require solvers for forward and adjoint PDE models. These approaches become even more effective when dealing with complex posterior distributions. Another approach, delayed acceptance MCMC [3, 8], leverages cheaper model approximations to accelerate parameter search, often using coarser spatial discretizations in multilevel hierarchy. Building on these methods, multilevel MCMC techniques improve mixing and sampling efficiency through variance reduction (similar to multilevel Monte Carlo [10, 4]) and hierarchical spatial discretizations [7, 9], in which [7] provides a theoretical cost analysis and numerical results applied to a two-dimensional example, and [9] demonstrates a

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL-TR-2005705)

speedup compared to traditional MCMC methods when applied to a three-dimensional problem. Notably, [16] performs scaling studies of two-dimensional spatial problems utilizing load balancing to manage running parallel chains (including independent expectation approximations) and distributed PDE evaluations, which results in efficient scaling. More recently, these cheaper models have included offline trained neural networks [12, 15] to accelerate MCMC.

All these approaches greatly improve MCMC scaling and efficiency. Future work should utilize similar tools, potentially combining them as in [5], and identify and improve aspects that create bottlenecks in overall scalability. Furthermore, approaches that increase the level of parallelism for outer lop methods will greatly improve overall scaling.

References

[1] MFEM: Modular finite element methods library. mfem.org.

[2] W. L Briggs, V. E. Henson, and S. F McCormick. *A multigrid tutorial*. SIAM, 2000.

[3] J. A. Christen and C. Fox. Markov chain Monte Carlo using an approximation. *Journal of Computational and Graphical Statistics*, 14(4):795–810, 2005.

[4] K. A. Cliffe, M. B. Giles, R. Scheichl, and A. L. Teckentrup. Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Computing and Visualization in Science*, 14(1):3–15, 2011.

[5] T. Cui, G. Detommaso, and R. Scheichl. Multilevel dimension-independent likelihood-informed MCMC for large-scale inverse problems. *Inverse Problems*, 40(3):035005, 2024.

[6] T. Cui, J. Martin, Y.M. Marzouk, A. Solonen, and A. Spantini. Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems*, 30(11):114015, 2014.

[7] T. J. Dodwell, C. Ketelsen, R. Scheichl, and A. L. Teckentrup. A hierarchical multilevel Markov chain Monte Carlo algorithm with applications to uncertainty quantification in subsurface flow. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):1075–1108, 2015.

[8] Y. Efendiev, T. Hou, and W. Luo. Preconditioning Markov chain Monte Carlo simulations using coarse-scale models. *SIAM Journal on Scientific Computing*, 28(2):776–803, 2006.

[9] H. R. Fairbanks, U. Villa, and P. S. Vassilevski. A multilevel, hierarchical decomposition of finite element white noise with application to multilevel Markov chain Monte Carlo. *SIAM Journal of Scientific Computing*, 43(5):S293–S316, 2021.

[10] M.B. Giles. Multilevel monte carlo path simulation. *Operations Research*, 56(3):607–617, 2008.

[11] Lawrence Livermore National Laboratory. hypre: High performance preconditioners.

[12] M. B Lykkegaard, T. J. Dodwell, and D. Moxey. Accelerating uncertainty quantification of groundwater flow modelling using a deep neural network proxy. *Computer Methods in Applied Mechanics and Engineering*, 383:113895, 2021.

[13] J. Martin, L. C. Wilcox, C. Burstedde, and O. Ghattas. A stochastic newton MCMC method for large-scale statistical inverse problems with application to seismic inversion. *SIAM Journal on Scientific Computing*, 34(3):A1460–A1487, 2012.

[14] N. Petra, J. Martin, G. Stadler, and O. Ghattas. A computational framework for infinite-dimensional Bayesian inverse problems, part II: Stochastic newton MCMC with application to ice sheet flow inverse problems. *SIAM Journal on Scientific Computing*, 36(4):A1525–A1555, 2014.

[15] S. Reddy and H. Fairbanks. Accelerating multilevel Markov chain Monte Carlo using machine learning models. *Physica Scripta*, 2024.

[16] L. Seelinger, A. Reinarz, L. Rannabauer, M. Bader, P. Bastian, and R. Scheichl. High performance uncertainty quantification with parallelized multilevel Markov chain Monte Carlo. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–15, 2021.