

High Performance Uncertainty Quantification with Parallelized Multilevel Markov Chain Monte Carlo

Linus Seelinger
mail@linusseelinger.de
Institute for Scientific Computing
Heidelberg University
Heidelberg, Germany

Michael Bader
bader@in.tum.de
Department of Informatics
Technical University of Munich
Garching, Germany

Anne Reinarz
anne.k.reinarz@durham.ac.uk
Department of Computer Science
Durham University
Durham, United Kingdom

Peter Bastian
peter.bastian@iwr.uni-heidelberg.de
Institute for Scientific Computing
Heidelberg University
Heidelberg, Germany

Leonhard Rannabauer
rannabau@in.tum.de
Department of Informatics
Technical University of Munich
Garching, Germany

Robert Scheichl
r.scheichl@uni-heidelberg.de
Institute for Applied Mathematics
Heidelberg University
Heidelberg, Germany




Figure 1: Three levels of discretisation of a tsunami model used in a multilevel Markov chain Monte Carlo method.

ABSTRACT

Numerical models of complex real-world phenomena often necessitate High Performance Computing (HPC). Uncertainties increase problem dimensionality further and pose even greater challenges.

We present a parallelization strategy for multilevel Markov chain Monte Carlo, a state-of-the-art, algorithmically scalable Uncertainty Quantification (UQ) algorithm for Bayesian inverse problems, and a new software framework allowing for large-scale parallelism across forward model evaluations and the UQ algorithms themselves. The main scalability challenge presents itself in the form of strong data dependencies introduced by the MLMCMC method, prohibiting trivial parallelization.

Our software is released as part of the modular and open-source MIT Uncertainty Quantification Library (MUQ), and can easily be

coupled with arbitrary user codes. We demonstrate it using the Distributed and Unified Numerics Environment (DUNE) and the ExaHYPE Engine. The latter provides a realistic, large-scale tsunami model in which we identify the source of a tsunami from buoy-elevation data.

CCS CONCEPTS

• **Computing methodologies** → **Massively parallel algorithms; Uncertainty quantification**; • **Applied computing** → *Earth and atmospheric sciences; Mathematics and statistics.*

KEYWORDS

Multilevel Methods, ADER-DG, Bayesian inverse problems, Tsunami simulation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SC '21, November 14–19, 2021, St. Louis, MO, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8442-1/21/11...\$15.00

<https://doi.org/10.1145/3458817.3476150>

ACM Reference Format:

Linus Seelinger, Anne Reinarz, Leonhard Rannabauer, Michael Bader, Peter Bastian, and Robert Scheichl. 2021. High Performance Uncertainty Quantification with Parallelized Multilevel Markov Chain Monte Carlo. In *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC '21)*, November 14–19, 2021, St. Louis, MO, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3458817.3476150>

1 INTRODUCTION

Numerical simulation is an established tool driving innovation in many fields of science and engineering. Numerical solutions of mathematical models, for example in the form of partial differential equations (PDE), provide a prediction about a real-world process. More generally, deterministic models map a given set of parameters to a specific set of predicted values. There are, however, many applications where this approach is insufficient and a stochastic model is needed: The exact values of model parameters might be known only up to a certain accuracy, leading to a corresponding uncertainty in model predictions. The importance of capturing these uncertainties cannot be overstated, since overlooking an unlikely yet dangerous scenario could easily have fatal consequences in real-world applications.

The aforementioned uncertainty in parameters is typically expressed in terms of probability distributions. Finding the distribution of model predictions resulting from stochastic parameter distributions is referred to as a forward UQ problem. When in turn real-world measurements, also affected by uncertainty, are available and the underlying stochastic model parameters explaining those measurements are to be quantified, we speak of an inverse UQ problem. The resulting parameter distribution is called the posterior.

Here, we focus on inverse UQ problems. These are notoriously ill-posed, that is, the solution may lack existence, uniqueness, or continuous dependence on the data. As a result, specialized methods are required for inverse UQ problems. There is a wide range of methods available, and they vary significantly in how they balance expressivity (in the sense of how much information can be gained) against computational efficiency and against required knowledge about the model. For example, there are highly efficient optimization-based methods such as [19] for determining the maximum a posteriori probability (MAP) point when derivatives of the data misfit functional are available. On the other hand, stochastic collocation type methods can work with simple model evaluations and recover the posterior [24], but become inefficient for higher dimensional parameter spaces [33]. Markov Chain Monte Carlo (MCMC) type methods in turn also recover the full posterior distribution and in many cases only need simple forward model evaluations, but come at a very high computational cost in terms of numerous model evaluations.

In this paper, we choose MCMC as the “gold standard” type of method [20] in the sense of recovering the exact posterior with minimal assumptions on the model. In order to overcome the enormous computational burden incurred by MCMC on costly models, we employ a multilevel MCMC method and present a new parallelization strategy to employ multilevel MCMC on modern HPC systems. We then couple our UQ method to sophisticated PDE solvers in order to obtain efficient forward models. Since our multilevel MCMC method only assumes simple forward model evaluations, no modifications to the solver frameworks are required. By building efficient model hierarchies, we exploit properties of the forward solvers and show synergies between the PDE models and the multilevel UQ approach. Efficient parallelization is a particular focus of this work since, on the one hand, it is inevitable when solving large-scale PDE models, and on the other hand, multilevel MCMC introduces

data dependencies that make parallelization non-trivial compared to plain-vanilla Monte Carlo (MC) methods.

To demonstrate the efficiency of our method in practical applications we use buoy data from coastal Japan to infer the location of initial displacements that led to the Tohoku tsunami in 2011. This method of using DART buoy data to predict tsunamis is commonly used in many early warning systems, such as the one at the Pacific Tsunami Warning Center operated by NOAA in the United States and at the National Tsunami Warning Center. These methods are known to work well for tsunamis which are initiated more than 2-3 wavelengths away from the coast.

We model the propagation of the tsunami by solving the shallow water equations. Our model is capable of modeling wetting and drying so that coastal regions can be included in the simulation [27]. For the numerical solution of the PDE, we apply an Arbitrary-high-order-DErivative Discontinuous Galerkin Method (ADER-DG) implemented in the ExaHyPE framework [29]. The parallelized Multilevel Markov Chain Monte Carlo (MLMCMC) method was implemented in the MUQ library [26], is publicly available and fully model agnostic. To our knowledge, this is the first parallel MLMCMC implementation available; a sequential one [22] is available as part of PyMC3 [32].

We further show parallel scalability of our UQ algorithm and implementation by coupling to a model implemented in the DUNE framework [1], at the same time demonstrating how our UQ software framework can be coupled with arbitrary model codes.

2 EFFICIENT SOLUTION OF INVERSE PROBLEMS VIA MULTILEVEL MCMC

For a general model F mapping parameters θ to model predictions $F(\theta)$, we formulate a Bayesian inverse problem. Its goal is to use the forward model F in order to infer the distribution of an uncertain model parameter θ from measurement data y :

$$\nu := \mathcal{P}(\theta|y) = \frac{\mathcal{L}(y|\theta)\pi(\theta)}{\mathcal{P}(y)} \propto \mathcal{L}(y|\theta)\pi(\theta). \quad (1)$$

Here, $\mathcal{P}(\theta|y)$ is the inferred parameter distribution given measurements y . The likelihood \mathcal{L} is defined as the probability of observing the given measurements if θ were the true parameter. Further, π is the prior distribution encoding a priori knowledge about the inferred parameters. We treat the distribution of measurements $\mathcal{P}(y)$ as a practically unobtainable scaling factor, as it is irrelevant for MCMC methods. The final target is the mean of a Quantity of Interest (QOI) $\mathbb{E}_\nu[Q]$ with respect to the above posterior distribution.

Evaluating the likelihood typically means comparing the model prediction $F(\theta)$ to the measurements y . For example, assuming Gaussian measurement errors with covariance Σ_f , the likelihood will have the distribution $\mathcal{N}(F(\theta), \Sigma_f)$. In our applications, evaluating F means solving a PDE, so this is where the main cost of a Bayesian problem lies.

The main idea of MCMC is to generate a Markov chain, carefully designed to have a stationary distribution matching the posterior distribution we are looking for. Thus, for a sufficiently large number of steps, the chain will approximately draw samples from the otherwise inaccessible posterior. Those samples can in turn be used

in a MC-like fashion to estimate $\mathbb{E}[Q]$ or other statistics of the posterior.

Algorithm 1 shows one of the most common MCMC algorithms, namely Metropolis-Hastings Markov Chain Monte Carlo (MHMCMC) [25]. One of its main advantages is that it only requires a finite number of direct evaluations of the posterior and thereby of the forward model. No further information such as model derivatives or adjoints are needed.

Algorithm 1: Metropolis Hastings MCMC

Result: Markov chain $\{\theta^i\}_{i=0}^N$.
 Choose starting parameter $\theta^0 \in \mathbb{R}^m$;
for $i = 0, \dots, N - 1$ **do**
 Draw proposal θ' from proposal distribution $q(\theta'|\theta^i)$.
 Compute acceptance probability

$$\alpha(\theta'|\theta^i) = \min \left\{ 1, \frac{v(\theta')q(\theta^i|\theta')}{v(\theta^i)q(\theta'|\theta^i)} \right\}.$$

 Draw a random number $r \in [0, 1]$.
 if $r < \alpha(\theta'|\theta^i)$ **then**
 Accept proposal: $\theta^{i+1} = \theta'$.
 else
 Reject proposal: $\theta^{i+1} = \theta^i$.
end
end

Once we have these samples approximating the posterior, we can apply some post processing to compute $\mathbb{E}_v[Q]$. In practice, since in the MCMC algorithm we already compute the forward model for each sample, we can immediately compute QOI samples derived from the model evaluations.

In order to accurately represent a multi-dimensional probability distribution, a significant number of samples may be required. Just like basic MC methods, algorithm 1 requires a full model evaluation per sample. If the model itself is costly, the overall computational effort can easily become intractable.

Most effort to improve the efficiency of MHMCMC goes into finding good proposal distributions $q(\theta'|\theta^i)$: MHMCMC produces correlated samples and strongly correlated samples barely contribute information to the posterior approximation. Less correlated proposals and high acceptance rates allow to achieve a good approximation with fewer samples. A wide variety of improved proposals aims to achieve this. Examples include preconditioned Crank-Nicolson [3, 5, 30], Adaptive Metropolis [16, 17], Hamiltonian MCMC [11], Dimension-Independent Likelihood-Informed (DILI) MCMC [6, 7], and many others.

Rather than minimizing the number of model evaluations for a single model, MLMCMC [9, 10, 18] takes a different route: An entire hierarchy of models is defined, ranging from cheap-to-compute rough approximations to the most accurate, yet expensive, full model. MLMCMC makes no assumptions on what exactly those coarse models could be. For example, we could use suitable ordinary differential equations (ODE) as rough approximations of a more complex PDE model, as long as model evaluations are sufficiently close. A more obvious and theoretically supported choice for a level hierarchy is varying mesh width in numerical PDE solvers,

where theory typically guarantees that coarser meshes still deliver reasonable approximations. Note that the fundamental approach is very closely related to Multilevel Monte Carlo (MLMC) methods [15]. However, as we will see later, there are additional intricacies and benefits to be gained in the MLMCMC setting.

The basic idea of MLMCMC is to replace the estimation of the expected value of the QOI by a telescoping sum

$$\mathbb{E}_{v_L}[Q_L] = \underbrace{\mathbb{E}_{v_0}[Q_0]}_{\text{Coarse approx.}} + \sum_{l=1}^L \underbrace{(\mathbb{E}_{v_l}[Q_l] - \mathbb{E}_{v_{l-1}}[Q_{l-1}])}_{\text{Corrections}}, \quad (2)$$

where Q_l denotes the approximation of the QOI Q using model l of the model hierarchy. Working with a hierarchy of models not only affects the QOI: Each forward model F_l now implies its very own likelihood function \mathcal{L}_l and corresponding posterior density v_l according to the Bayesian inverse problem formulated above.

This reformulation is clearly equivalent, but offers opportunities for significantly reducing computational cost:

- Coarser chains can be used as cheap-to-compute, well informed and nearly uncorrelated proposals for finer chains.
- Variance reduction between levels can be exploited: The coarser corrections in eq. (2) need many samples, but those are cheap to compute; fine corrections are more expensive per sample, but the variance in those corrections is reduced and thus only few samples are needed, especially when the models are converging as $l \rightarrow \infty$.

An MLMCMC method implementing both ideas is shown in algorithm 2, as originally introduced by the authors of [9, 10].

Here we begin with a regular MCMC chain for the coarsest level. We then proceed to generate MCMC chains for finer levels, while using samples from coarser chains as proposals. Since we permit increasing parameter dimensions across levels, the coarse samples $\theta'_{l,C}$ from level $l - 1$ used as proposals for level l may need to be complemented with a fine proposal density q_l in order to form a proposal θ'_l of suitable dimension. Further, compared to MHMCMC, the acceptance probability also needs to be adapted to avoid bias from the coarse proposals.

The efficiency gain of this multilevel algorithm over MHMCMC can be analyzed theoretically for models where PDE theory provides approximation error bounds for given cost. In [9, Thm. 3.4], the following cost estimate was proven for a Poisson equation model in d dimensions similar to the Poisson model we use later. It is based on the mean square error of the estimate of $\mathbb{E}_v[Q]$, either directly computed via MHMCMC or from the telescoping sum eq. (2) in the MLMCMC case.

The computational cost C_ϵ required to achieve a mean square error below ϵ is then bounded by

$$C_\epsilon^{\text{MCMC}} \lesssim \epsilon^{-(d+2)-\delta} \quad \text{and} \quad C_\epsilon^{\text{MLMCMC}} \lesssim \epsilon^{-(d+1)-\delta}$$

respectively, where δ is a model specific constant. In this setting, which should be representative of many PDE models with model hierarchies based on mesh width, the cost for the multilevel method is therefore one order below the single-level one.

The hidden constants in the bounds are proportional to the integrated autocorrelation times of the Markov chains. On the finest

Algorithm 2: Multilevel MCMC

Result: Markov chains $\{\theta_l^i\}_{i=0}^{N_l}$ for all levels $l \in \{0, \dots, L\}$.
 On level 0, run a conventional MCMC chain, delivering samples θ_0^i .

for level $l = 1, \dots, L - 1$ **do**
 Choose starting point θ_l^0 with coarse component from next coarser starting point θ_{l-1}^0 .
 for sample $j = 1, \dots, N_l$ **do**
 Given θ_l^j , generate proposal $\theta'_l = \begin{Bmatrix} \theta'_{l,C} \\ \theta'_{l,F} \end{Bmatrix}$ where
 • $\theta'_{l,C}$ is drawn from a level $l - 1$ chain with subsampling rate ρ_l and
 • $\theta'_{l,F}$ is drawn from a proposal density $q_l(\theta'_{l,F} | \theta_{l,F}^j)$.
 Compute acceptance probability

$$\alpha(\theta'_l, \theta_l^j) = \min \left\{ 1, \frac{v_l(\theta'_l) q_l(\theta_l^j | \theta'_{l,F})}{v_l(\theta_l^j) q_l(\theta'_{l,F} | \theta_l^j)} \cdot \frac{v_{l-1}(\theta_{l,C}^j)}{v_{l-1}(\theta'_{l,C})} \right\}.$$

 Draw a random number $r \in [0, 1]$.
 if $r < \alpha(\theta'_l, \theta_l^j)$ **then**
 Accept proposal: θ'_l as θ_l^{j+1}
 else
 Reject proposal: $\theta_l^{j+1} = \theta_l^j$
 end
 end
end

level in the multilevel approach this factor is essentially reduced to one [10], leading to significant additional computational gains.

It is beyond the scope of this work to also prove the assumptions underlying the theoretical error estimates for our more complex tsunami model. This would be a particularly challenging task since we do not only rely on mesh refinement, but change bathymetry data across levels for improved solver performance. However, since our numerical PDE solver should fundamentally exhibit comparable scaling behaviour, we can reasonably expect similar gains in efficiency by employing the multilevel method in that setting as well.

3 MODELS

We will apply parallelized MLMCMC to two different models: The widely used and cheap to compute Poisson equation modelling single-phase subsurface flow as well as a more expensive and realistic Tsunami model. The former serves as a well-known reference, and due to its low cost is also suitable to gauge the MLMCMC parallelization in relatively short scaling experiments. The latter, due to its more interesting properties, allows us to demonstrate the practical value of the method, as well as the opportunities in creating model-specific hierarchies.




Figure 2: Random field realization $\log(\kappa(\cdot, \hat{\theta}))$ and parameter field $\kappa(\cdot, \hat{\theta})$ from random field realization used for synthetic data.

3.1 Poisson Model

In this example, our forward model maps a parameter θ that models the uncertainty in the diffusion coefficient to the solution of the Poisson PDE evaluated at certain points. The inverse problem consists in estimating the underlying parameter θ from given synthetic data while taking uncertainty into account.

Specifically, we solve the PDE

$$\nabla \cdot (\kappa(x, \theta) \nabla u(x, \theta)) = 0 \quad \text{for } x \in \Omega \text{ and } \theta \in \Theta,$$

where we choose the domain $\Omega := [0, 1]^2$ and $\Theta := \mathbb{R}^m$. As boundary conditions, we apply $u(x) = 0$ on the left, $u(x) = 1$ on the right and natural Neumann boundary conditions at the remainder of the boundary. We model $\log(\kappa)$, the logarithm of the diffusion coefficient, as a zero-mean Gaussian random field with correlation length 0.15 and variance 1. In order to arrive at a finite dimensional representation of the field, we truncate use a Karhunen Loève (KL) expansion which we truncate after m terms, i.e.

$$\log(\kappa(x, \theta)) \approx \sum_{k=1}^m \phi_k(x) \theta_k,$$

where ϕ_1, \dots, ϕ_m are the KL modes of largest wave length. Consequently, θ is a vector of KL coefficients. We implement the model in the DUNE framework [1] with a Q_1 Finite Element discretization on simple structured grids. To form a three-level model hierarchy for our MLMCMC method, we choose mesh widths of $\frac{1}{16}$, $\frac{1}{64}$ and $\frac{1}{256}$. Across all three levels, we choose an identical parameter dimension $m = 113$; this specific number carries not much significance beyond being suitable for implementation in the circulant embedding [8] based random field generator dune-randomfield.

In order to form a Bayesian inverse problem, we generate synthetic data based on a random field $\kappa(x, \hat{\theta})$, where $\hat{\theta}$ is a fixed sample drawn from $\mathcal{N}(0, I)$ (shown in fig. 2).

The actual vector of measurements y is then defined by solving the above Poisson problem for $\theta = \hat{\theta}$, evaluating the solution u at a grid of points: $\{\frac{2}{32}, \frac{7}{32}, \frac{13}{32}, \frac{19}{32}, \frac{25}{32}, \frac{3}{32}\}^2$. Based on that, we define our likelihood $\mathcal{L}(y|\theta)$ to be a Gaussian $\mathcal{N}(F(\theta), \sigma_F^2 I)$ with $\sigma_F = 0.01$. Complementing it with a Gaussian prior $\pi(\theta) = \mathcal{N}(0, 4I)$, we complete our Bayesian inverse problem.

Note that, by generating synthetic measurements directly from our forward model, we commit an 'inverse crime' [4, p. 179]. In realistic applications a model error is inevitable, making it significantly harder to recover the underlying parameters from data accurately. However, for this problem we intentionally accept this simplification: Our focus in this case is algorithmic scalability and not a fully realistic setting. Further, verifying the correctness of the UQ method is somewhat easier without a model error.

As our QOI, we define $Q(\theta)_k = \kappa(x_k, \theta)$ where the x_k form a grid of width $\frac{1}{32}$. This captures the parameter field we seek and, as necessitated by the telescoping sum in eq. (2), allows for a consistent dimension in QOI even when varying the parameter dimension across levels.

3.2 Tsunami Model

Tsunami propagation is typically modelled by some variant of the shallow water equations [2, 21, e.g.], allowing for a simulation in only two dimensions. The shallow water equations are obtained via depth-averaging of quantities (esp. momentum) from the more complicated three-dimensional Navier-Stokes equations, based on the modeling assumption that horizontal length scales are considerably greater than the vertical length scales.

In this setting, we concentrate on the basic shallow water equations with bathymetry source terms (neglecting friction terms or more advanced models for with non-hydrostatic corrections). The resulting equations can be written in first-order hyperbolic form as

$$\frac{\partial}{\partial t} \begin{pmatrix} h \\ hu \\ hv \\ b \end{pmatrix} + \nabla \cdot \begin{pmatrix} hu & hv \\ hu^2 & huv \\ huv & hv^2 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} 0 \\ hg \partial_x(b+h) \\ hg \partial_y(b+h) \\ 0 \end{pmatrix} = 0, \quad (3)$$

where h denotes the height of the water column, (u, v) the horizontal flow velocity, g gravity and b denotes the bathymetry. This hyperbolic system of equations is supplemented by a set of suitable initial and boundary values.

We discretise with an ADER-DG method as proposed in [12]. It is essentially a predictor-corrector scheme. A high-order solution is found element-locally and then corrected to take into account neighbors by solving Riemann problems along element interfaces. To resolve known high-order issues such as the Gibbs phenomenon a corresponding a-posteriori finite volume sub-cell limiter is applied [13]. This limiter detects and revokes problematic ADER-DG solution candidates and recomputes them with a robust Finite-Volume scheme (cf. [27, 28] for details), following the approach by LeVeque et al. [21]. At coastlines the schemes relies entirely on the Finite-Volume limiter, to correctly treat inundation.

As a large-scale example, we invert data from the Tohoku tsunami, which occurred subsequent to an earthquake in the Japan trench in 2011. We assume that the only significant sources of the tsunami are the displacements of the sea floor. In order to initialise the tsunami, we can impose the displacements as an instantaneous deformation of the bathymetry in the resting-lake case – compare respective modeling approaches by Saito et al. [31] or Madden et al. [23]. By keeping the water column constant, the change of the bathymetry is directly translated to the sea surface and generates the tsunami. Gravity, as the main acting force, initiates the propagation of the wave. The tsunami then evolves as a gravity wave. As reference




Figure 3: Bathymetry for the full computational domain with darker rectangle denoting parameter values in the prior.




Figure 4: Plot of Sea surface height anomaly [ssha] for samples taken from level 0 (top) and 1 (bottom) compared to NDBC data at buoy 21418.

solution for the initial displacements of the ocean floor we use respective simulation results provided by Galvez et al. [14]. The bathymetry data has been obtained from GEBCO ¹.

Our goal is to obtain the parameters describing the initial displacements from the data of two available buoys located near the Japanese coast. Some of these parameters are reasonably well-known already, these include location of the hypocenter, length, width, and to some extent depth. Other parameters such as uplift are more difficult to estimate. In these tests we estimate the location of the initial displacement. The prior cuts off all parameters which would lead to an initial displacement which is too close to the domain boundary. Figure 3 shows the cut-off values used. Some parameters may lead to unstable models, e.g. a parameter which initialises the tsunami on dry land, in this case we have treated the parameter as unphysical and assigned an almost zero likelihood.

In order to compute a likelihood of a given set of parameters given the simulation results we use a weighted average of the maximal wave height and the time at which it is reached. The likelihood is given by a normal distribution $\mathcal{N}(\mu, \Sigma)$ with mean μ given by maximum waveheight $\max\{h\}$ and the time t at which it is reached for the two DART buoys 21418 and 21419 ². The covariance matrix Σ depends on the level, but not the probe point. Figures 4 and 5 shows the data and samples of level 0 and 1. Table 1 gives the values of μ and the diagonal entries of Σ for all three

¹<https://www.gebco.net/>

²This data can be obtained from NDBC <https://www.ndbc.noaa.gov/>




Figure 5: Plot of Sea surface height anomaly [ssha] for samples taken from level 0 (top) and 1 (bottom) compared to NDBC data at buoy 21419.

μ	Σ		
	$l=0$	$l=1$	$l=2$
1.85232	0.15	0.1	0.1
0.6368	0.15	0.1	0.1
30.23	2.5	1.5	0.75
87.98	2.5	1.5	0.75

Table 1: Mean μ and covariance Σ for all three levels.

level	order	limiter	h	# timesteps	DOF updates
0	2	no	1/25	98	$2.4 \cdot 10^5$
1	2	yes	1/79	306	$9.4 \cdot 10^6$
2	2	yes	1/241	932	$2.7 \cdot 10^8$

Table 2: Polynomial order, limiter status, mesh width (h), number of timesteps and number of degree of freedom (DOF) updates for the three models.

levels. Alternative likelihood functions, such as a quadratic average of multiple buoys could also be used, see e.g. [2].

We set up a sequence of three models with increasing accuracy shown as in Figure 1. In the first model bathymetry is approximated only by a depth average over the entire domain. Since no calculations of wetting and drying are needed this model is computed purely with a DG method of order 2. The second and third models further include a finite volume subcell limiter allowing for wetting and drying. The second model uses smoothed bathymetry data and the third uses the full bathymetry data. The main advantage of using smoothed data is that the FV subcell limiter is needed in fewer cells. This model hierarchy demonstrates that not only mesh refinement and coarsening, but also model specific optimisations can be used to exploit multilevel MCMC.

It should be noted that the limiter, which works on a locally refined grid, leads to a varying computational load per sample. The number of degrees of freedom for each model are given in Table 2 for the sample calculated with the parameters (0, 0).

4 PARALLEL MLMCMC IMPLEMENTATION

In this section we describe the new, highly scalable parallel implementation of MLMCMC used for the experiments in this paper. While parallelization in classical MC is trivial due to the independence of the samples, MCMC introduces data dependency through proposals depending on the previous step. In the case of MLMCMC as in [9, 10], we use coarser chain samples as proposals, which also introduces data dependency between levels. There are, however, multiple opportunities for parallelizing MLMCMC:

- **Models:** The forward models themselves may be parallelized. In fact, for large models like the tsunami model we introduce in section 3.2, that is inevitable anyway due to memory constraints.
- **Chains:** Instead of running a single Markov chain, multiple chains can be run in parallel and their samples combined. It is beneficial to not purely rely on this approach though since each chain requires a burn-in phase.
- **Levels:** Contributions to the multilevel telescoping sum in eq. (2) can be evaluated in parallel.

Exploiting all those clearly introduces significant technical complexity. Therefore we provide our implementation as part of the MUQ C++ library [26]. The main goals of this implementation are:

- **Parallelism:** All of the above levels of parallelism are supported.
- **Simple user interface:** We hide the intricate details of communication from the user. The algorithm can be tweaked, but to get started the defaults suffice.
- **Model-agnosticity:** The MLMCMC algorithm as detailed in section 2 only requires simple forward evaluations of the model. This theoretically allows coupling to arbitrary forward models without any need for, e.g., derivatives of the model map. We retain this model agnosticity in the sense that any forward model that can be called (possibly through wrappers) from C++ can be used.
- **Modularity:** MUQ is, from the ground up, designed as a modular framework. Its modularity is closely modelled after the respective mathematical objects. We extend this concept to our parallel MLMCMC implementation by building on top of MUQ's existing MCMC stack and, as detailed in the following, constructing modular parallel units.

4.1 Model interface

MUQ provides an abstract interface for sampling algorithms, including MCMC type methods. In its most basic form, a model can be provided by implementing an `AbstractSamplingProblem` (see fig. 6). Our parallel MLMCMC implementation supports the same model interface, allowing to quickly use various UQ methods on a single model implementation. The interface requires the implementation of the density to sample from; in case of our Bayesian inverse problem, that is the posterior density in eq. (1). Further, a QOI may be provided. This is implemented as a separate function call, since discarded samples in MCMC do not contribute to the QOI computation, and potentially expensive evaluations can thus be skipped.

A single implementation of this interface is already sufficient to apply various MCMC type methods provided by MUQ. In the




Figure 6: AbstractSamplingProblem interface.

multilevel case, however, we clearly need to provide a hierarchy of posterior densities implying a hierarchy of models. Further, proposal densities and subsampling strategies for drawing from coarser chains can be chosen. Simply passing a fixed list of models may lead to various complications and performance issues in a parallel settings. Therefore, we provide a factory type interface as in fig. 7.

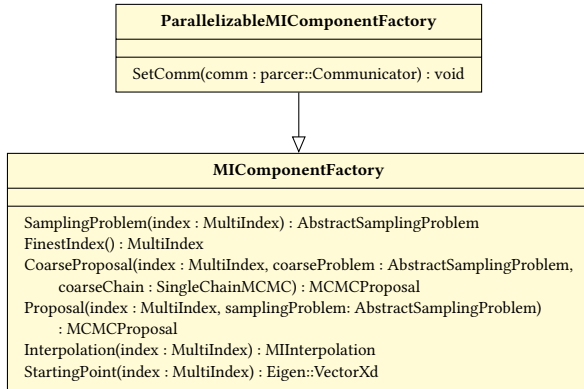


Figure 7: ParallelizableMComponentFactory interface.

When implementing the MComponentFactory interface, the SamplingProblem method is to return a model for a given model index, while FinestIndex specifies the index of the finest model the user provides, corresponding to L in algorithm 2. CoarseProposal specifies how proposals are being drawn from coarser chains, and Interpolation determines how these are combined with a finer chain proposal as defined in Proposal.

In order to support parallel models, the user may modify ParallelizableMComponentFactory, which receives an MPI communicator. The user can pass a parcer communicator to the model, this is essentially just an MPI C++ wrapper used in MUQ.

Note that, while in this paper we entirely focus on the multi-level case, the implementation actually supports a generalization to a multiindex method. Therefore, components are named MI for multiindex rather than ML, and the model hierarchy is indexed by a more general multiindex structure rather than an integer index.

4.2 Internal architecture

Our parallel process layout is shown in fig. 8. We define the following roles:

- **Fixed roles:** These are assigned to specific processes at the start of the parallel method. All other processes wait to be assigned a dynamic role.
 - **Root:** This process is responsible for launching the parallel method, assigning tasks to other processes and requesting collectors to begin collecting a certain number of MCMC

samples. It is also the best place for users to implement custom (possibly adaptive) sampling strategies.

- **Phonebook:** The phonebook tracks what dynamic roles processes are currently assigned to. Most importantly, it tracks which chains are currently sampling and which ones hold new samples ready to be picked up by other processes. Further, the phonebook can infer from that the computational load on a given level, since the relation between requested samples and assigned resources is available here. Therefore, it is also the key component in dynamic load balancing across levels and chains.
- **Dynamic roles:** These roles may be assigned and reassigned at any time. In particular, this permits dynamic load balancing and possibly more advanced sampling strategies. Workers and controllers solve forward models, where workers share the load of running a single model evaluation and controllers additionally run the inherently sequential MCMC chains; consequently, they are assigned (and, in case of dynamic load balancing, reassigned) synchronously. This is facilitated by Message Passing Interface (MPI) subcommunicators, which are passed through to and should be used by the user's model.
 - **Worker:** Workers are responsible for running the user's forward model, specifically the user's implementations of AbstractSamplingProblem. Mathematically, they provide parallelized evaluations of the posterior and quantity of interest on a given level for a given parameter. They listen to their respective controller's ParallelAbstractSamplingProblem to signal the beginning of an evaluation for a specific parameter θ . Once the signal arrives, they execute the user-implemented LogDensity method. Since all workers of a work group are called synchronously, the user can easily supply models assuming lock step parallelism.
 - **Controller:** Each controller is responsible for running a multilevel MCMC chain according to algorithm 2. Specifically, a controller on level l contains a chain on level l and one on $l - 1$, to compute the level l correction of the telescoping sum eq. (2) (with the obvious exception of the coarsest level 0). An instance of ParallelAbstractSamplingProblem is set up for each model needed. It provides an intermediate layer between instances of the user-implemented AbstractSamplingProblem instances running on multiple workers and the controller, allowing to transparently distribute model executions. As a result, neither the user nor the inherently sequential MLMCMC chains need to concern themselves with synchronizing worker processes to run forward models in lock step. The chains themselves are implemented using existing MUQ components: They are SingleChainMCMC instances with MCMCKernel implementations matching the acceptance probability of algorithm 2. Drawing samples from coarser chains as proposals is implemented with an MCMCProposal requesting coarser samples from other controllers via the phonebook process.

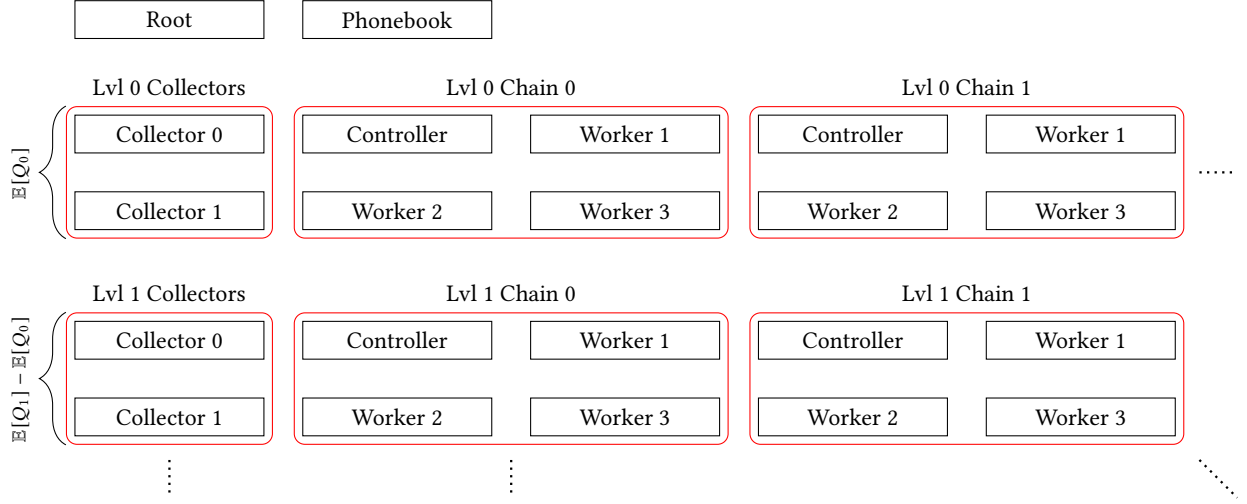


Figure 8: Parallel process layout.

- **Collector:** Collectors request samples from controllers via the phonebook in order to compute terms of the telescoping sum eq. (2). Multiple collectors may be responsible for a single level. Together, they hold a `DistributedCollection`, an existing class in MUQ for storing and computing statistics on samples in a parallel system.

Note that this architecture is defined on process level (more specifically, in terms of MPI ranks). Thread-level parallelism can easily be exploited by worker processes. In fact we make use of this through Intel Threading Building Blocks (TBB) in our Exascale Hyperbolic PDE Engine (ExaHyPE) Tsunami model, since ExaHyPE exhibits better performance characteristics with few MPI ranks per node each making use of several threads.

In order to make the parallel architecture as modular as sequential MUQ code, each of the above roles provide an MPI interface based on requests mimicking function calls. This allows recombining the parallel components in analogy to object orientation in order to implement other algorithms as well. For example, work groups as introduced above are based on a `ParallelAbstractSamplingProblem`, which in turn can be used to easily employ any sequential sampling algorithm on a parallelized model. Likewise, the phonebook could be swapped out for an alternative implementation with identical MPI interface, allowing for alternative load balancing strategies.

4.3 Load balancing

Data dependencies in MLMCMC (see algorithm 2) introduce a load balancing problem, since coarser chains need to provide proposals to finer ones only until the desired number of fine samples is computed. Estimating the ideal distribution of computational resources across levels is far from trivial or outright impossible in realistic applications, especially when adaptively determining the number of samples per level.

Therefore, our parallel MLMCMC implementation provides a load balancing mechanism to reassign worker processes to different tasks once samples on another level are more critical to runtime.




Figure 9: Dynamic load balancing in parallel MLMCMC. Run time is on the horizontal axis, while process indices are on the vertical. Green boxes indicate model evaluations, while yellow boxes indicate chains' burnin phases. In this example, model evaluations clearly differ strongly in run time.

Figure 9 illustrates this load balancing mechanism for a small test run.

Load balancing is implemented as part of the phonebook rank, since it keeps track of how samples are passed around. Levels with low load are detected when samples on that level are provided but not quickly picked up, while a high load is in turn detected when sample requests remain queued. Unanswered sample requests originating from other chains are given a higher impact than requests originating from collector processes, since the first case implies chains waiting and therefore bad machine utilization.

Models may have strongly varying run times. A new group of processes assigned to a certain level only reduces that level's load once it actually provides its first sample. This implies the danger of reassigning tasks too frequently or too infrequently. In order to stabilize the load balancer, the respective model run times are inferred by the phonebook process by the frequency of samples




Figure 10: Synthetic “true” field (left) and expected value of multilevel estimator (right).

level l	h_l	DOFs	t_l [ms]	ρ_l	τ_l	$\mathbb{V}[Q_0]$ or $\mathbb{V}[Q_l - Q_{l-1}]$
0	$\frac{1}{16}$	289	3.35	206	137.3	1.501×10^{-1}
1	$\frac{1}{64}$	4225	45.64	17	11.2	1.121×10^{-3}
2	$\frac{1}{256}$	66 049	931.81	0	1.05	4.165×10^{-5}

Table 3: Multilevel properties of Poisson application. For each level l of mesh width h_l with associated degrees of freedom (DOFs), we show the computational cost t_l and chosen subsampling rate ρ_l . Due to high dimension of QOI in this setting, we only show integrated autocorrelation time τ_l and variance for an single representative component of Q .

provided. Based on that, scheduling will only take place at the time scale of the respective model evaluations.

Note that this load balancer is unaware of the specific types of proposals or MCMC kernels being executed. As a result, it can, for example, also be applied in the MLMC setting. Also, all model specific tuning parameters are determined dynamically, so no user intervention is required.

5 NUMERICAL RESULTS

5.1 Poisson application and scalability

In order to fully specify the MLMCMC algorithm for the given problem, it is enough to set a Gaussian proposal on the coarsest level. We choose $\mathcal{N}(0, 3I)$ in order to roughly match the prior. Since we have identical parameter dimensions across levels, no fine level proposals are needed.

The MLMCMC method run with 10^4 , 10^3 and 10^2 samples on levels 0, 1 and 2 exhibits properties detailed in in table 3 and captures the main features of the parameter field underlying our synthetic data (see fig. 10). Clearly some higher frequency detail is not recovered. This, however, is expected due to the limited number of KL modes we include in our parameter space. Note that we can only recover this up to a scaling factor, since the solution is only determined by the parameter field up to a factor. Here, the choice of prior essentially determines the scaling of the solution we observe.

For a more detailed analysis of Bayesian inverse problems based on Poisson equation in MLMCMC, we refer to the original MLMCMC publication [9, 10].




Figure 11: Scalability of the Poisson model problem for 10^4 , 10^3 and 10^2 samples on levels 0, 1 and 2 respectively. Subsampling rates etc. are chosen according to table 3. The problem setup remains constant as the number of processors is increased.

In order to investigate parallel scalability of our MLMCMC implementation, we conduct weak and strong scaling experiments on the BwForCluster MLS&WISO Production HPC system. The partition we used consists of nodes with two 16-core Intel Xeon E5-2630v3 CPUs and 64 gigabytes of memory.

As forward model, we use the Poisson model, since its low computational demand allows us to stress the parallelized MLMCMC algorithm itself by running a large number of chains and samples. We use the same inverse problem detailed above, even though the particular inverse problem does not affect the algorithm’s communication patterns and therefore parallel scalability.

For the strong scaling setup, we draw 10^4 , 10^3 and 10^2 samples on levels 0, 1 and 2 respectively. We further set subsampling rates according to table 3, and enable dynamic load balancing. As the timing results in fig. 11 show, we achieve linear speedup until relatively large burnin phases and suboptimal load balancing due to few samples per chain occur.

Technically the observed speedup even slightly exceeds linear. That is simply due to the fact that a fixed number of the processes is reserved for book keeping tasks (i.e. the root, phonebook and collector processes). As a result, for increased number of processes, a larger fraction contributes to parallel speedup by generating samples.

In our weak scaling test, we begin with the same setup as in the strong scaling setting. In particular, we again choose to compute 10^4 , 10^3 and 10^2 samples on levels 0, 1 and 2, and solve this problem using 64 processes. We then expand to a range from 32 to 1024 processes while scaling the number of samples on each level linearly in accordance with the number of processes.

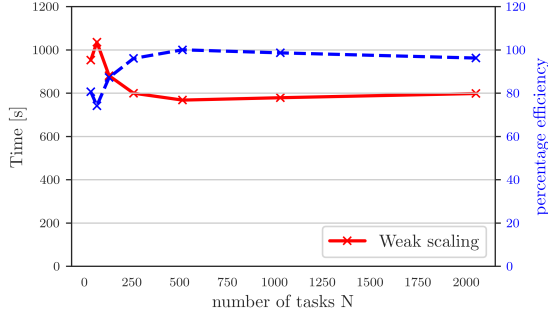


Figure 12: Weak scalability and parallel efficiency of the Poisson model problem. At 64 cores 10^4 , 10^3 and 10^2 samples are computed on levels 0, 1 and 2 respectively. The number of samples is modified linearly with the number of processors.

The parallel efficiency (given in blue in Figure 12) is measured here as $\frac{t_{\text{ref}}}{t_N} \cdot 100\%$, where t_{ref} is the quickest time taken over all runs and t_N is the time taken on N ranks. The initial increase to over 100% efficiency is due to the overhead of phonebook and collector ranks. We achieve fairly consistent results of up to 80 seconds of total run time except for the largest run. The latter is a very extreme scenario though, since the extremely short run time of the coarsest model leads to a significant load on the communication infrastructure. We therefore consider it reasonable to assume that exceeding the ideal range should only occur for significantly larger numbers of processes in more realistic applications.

5.2 Tsunami application

As a second example of the MLMCMC method we invert data from the Tohoku tsunami. We find the parameter distribution describing the initial displacements from the data of two available buoys located near the Japanese coast. The scenario is described in more detail in Section 3.2. The MLMCMC method with three levels computes 800 samples on level 0, 450 on level 1 and 240 on level 2 with a subsampling rate of 25 on level 0 and 5 on level 1.

These tests were run on up to 72 Intel Skylake Xeon Platinum 8174 nodes of SuperMUC-ng consisting of 48 cores each. The tests were run using Intel's TBB for parallelisation over the 48 cores of each node and MPI for parallelisation across nodes. Each ExaHyPE run used exactly one full node. The runtime for each forward model evaluation was on average 7.38 seconds on level 0, 97.3 seconds on level 1 and 438.1 seconds on level 2. These runtimes have a large variability as the model's timestep depends on the uncertain parameters, making it a challenging test for the scheduling infrastructure. In total 61,250 level 0, 34500 level 1 and 240 level 2 forward model evaluations were required for this test.

Again, due to constant parameter dimension across levels, we only need to choose a proposal density for the coarsest level. Like before, we choose Adaptive Metropolis [16, 17] provided by MUQ. As initial prior we set $\mathcal{N}(0, 10I)$ and update every 100 steps.

Figure 13 shows the resulting samples on level 0 and 1. The expected values $\mathbb{E}(Q_0)$, and $\mathbb{E}(Q_l - Q_{l-1})$ are shown as dashed




Figure 13: The three-level tsunami test case, each point represents an accepted sample at level $l = 0, 1, 2$. The dashed lines show the expected value $\mathbb{E}(Q_0)$ or $\mathbb{E}(Q_0) + \sum_l \mathbb{E}[Q_l - Q_{l-1}]$ with the point $(0, 0)$ in red for reference.




Figure 14: Visualization of samples used to estimate the corrections between levels in eq. (2). Left: Correction between levels zero and one. Right: Correction between levels one and two. Dots and origins of arrows indicate coarse samples. The arrows point towards their corresponding fine samples. Note that the arrow's length is scaled down to a factor of .15 for clarity. Coarse proposals that were accepted by the fine chain, leading to identical coarse and fine samples here, appear as simple dots instead of arrows.

line. The red marker shows the point $(0, 0)$. The point $(0, 0)$ is the position of the initial displacements as estimated in [14].

In order to illustrate how MLMCMC links coarser and finer level posteriors in order to obtain fine level corrections, fig. 14 shows how samples on level l relate to the level $l - 1$ samples that served as their respective coarse proposals. The result can be thought of as a transformation between corresponding coarser and finer distributions, even though only in a non-deterministic sense. Since in this application we choose the QOI to be the uncertain parameter itself, the estimate of the terms in the telescoping sum (eq. (2)) actually correspond to the mean of the corrections displayed here.

Variances and expected values are given in Table 4. The relatively cheap samples on level 0 provide a good initial estimate of the posterior, which are improved by the more expensive models utilising the full bathymetric data. In contrast to the Poisson case, we do not observe variance reduction across levels. This is somewhat expected, as the modified bathymetry does not permit the construction of a level hierarchy fulfilling the theoretical assumptions made by MLMCMC based on a priori error estimates. We do, however, still have the benefit of well-informed proposals on finer levels driven by coarser level chains.

6 CONCLUSION

In this paper we have presented a new parallelization strategy for MLMCMC as well as a model-agnostic implementation as part of the open-source and modular MUQ library. We have demonstrated the effectiveness of algorithm and implementation at solving Bayesian inverse UQ problems on complex and large-scale PDE models, highlighting opportunities in model-specific coarsening strategies in the process.

In order to verify our parallelization approach, we presented both strong and weak scaling results. An important feature of the work

lvl l	$t_l[s]$	ρ_l	$\mathbb{V}[Q_0]$ or		$\mathbb{E}[Q_0] +$	
			$\mathbb{V}[Q_l - Q_{l-1}]$	$\mathbb{V}[Q_l - Q_{l-1}]$	$\sum_{k=1}^l \mathbb{E}[Q_k - Q_{k-1}]$	$\sum_{k=1}^l \mathbb{E}[Q_k - Q_{k-1}]$
0	7.38	25	1984.09	1337.42	3.61	27.96
1	97.3	5	1592.17	1523.18	-12.29	23.39
2	438.1	0	340.56	938.53	-5.46	0.12

Table 4: Multilevel properties of the tsunami model. For each level l of mesh, we show the computational cost t_l and chosen subsampling rate ρ_l . We show variance and expected values for both components of Q .

is the applicability of the provided MUQ interface, which allows easy coupling to other models and software packages. This we have shown by applying it to a Poisson model problem implemented in the DUNE framework and a tsunami model implemented in the ExaHyPE-Engine.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the compute and data resources provided by the Leibniz Supercomputing Centre (www.lrz.de) under project number pr83no. We acknowledge funding from the European Union's Horizon 2020 Programme under the ENERXICO Project, grant agreement No. 828947.

The authors would also like to thank Jörn Behrens for his extremely useful advice.

We further acknowledge GEBCO Compilation Group (2020) GEBCO 2020 Grid for providing bathymetry data used in this work (doi:10.5285/a29c5465-b138-234d-e053-6c86abc040b9).

The authors acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through grant INST 35/1134-1 FUGG.

REFERENCES

- [1] P. Bastian, F. Heimann, and S. Marnach. 2010. Generic implementation of finite element methods in the Distributed and Unified Numerics Environment (Dune). *Kybernetika* 46, 2 (2010), 294–315.
- [2] J. Behrens, A. Androsov, A. Y. Babeyko, S. Harig, F. Klaschka, and L. Mentrup. 2010. A new multi-sensor approach to simulation assisted tsunami early warning. *Natural Hazards and Earth System Sciences* 10, 6 (2010), 1085–1100. <https://doi.org/10.5194/nhess-10-1085-2010>
- [3] Alexandros Beskos, Gareth Roberts, Andrew Stuart, and Jochen Voss. 2008. MCMC methods for diffusion bridges. *Stochastics and Dynamics* 08, 03 (2008), 319–350. <https://doi.org/10.1142/S0219493708002378> arXiv:<https://doi.org/10.1142/S0219493708002378>
- [4] David Colton and Rainer Kress. 2019. *Inverse Acoustic and Electromagnetic Scattering Theory* (4 ed.). Springer-Verlag, New York. <https://doi.org/10.1007/978-1-4614-4942-3>
- [5] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White. 2013. MCMC Methods for Functions: Modifying Old Algorithms to Make Them Faster. *Statist. Sci.* 28, 3 (08 2013), 424–446. <https://doi.org/10.1214/13-STS421>
- [6] Tiangang Cui, Gianluca Detommaso, and Robert Scheichl. 2019. Multilevel Dimension-Independent Likelihood-Informed MCMC for Large-Scale Inverse Problems. (10 2019).
- [7] Tiangang Cui, Kody J.H. Law, and Youssef M. Marzouk. 2016. Dimension-independent likelihood-informed MCMC. *J. Comput. Phys.* 304 (2016), 109 – 137. <https://doi.org/10.1016/j.jcp.2015.10.008>
- [8] C. Dietrich and G. Newsam. 1997. Fast and Exact Simulation of Stationary Gaussian Processes through Circulant Embedding of the Covariance Matrix. *SIAM J. Sci. Comput.* 18 (1997), 1088–1107.
- [9] Tim Dodwell, Chris Ketelsen, Robert Scheichl, and Aretha Teckentrup. 2015. A Hierarchical Multilevel Markov Chain Monte Carlo Algorithm with Applications to Uncertainty Quantification in Subsurface Flow. (08 2015). <https://doi.org/10.1137/130915005>
- [10] Tim Dodwell, Chris Ketelsen, Robert Scheichl, and Aretha Teckentrup. 2019. Multilevel Markov Chain Monte Carlo. *SIAM Rev.* 61 (01 2019), 509–545. <https://doi.org/10.1137/19M126966X>
- [11] Simon Duane, A.D. Kennedy, Brian J. Pendleton, and Duncan Roweth. 1987. Hybrid Monte Carlo. *Physics Letters B* 195, 2 (1987), 216 – 222. [https://doi.org/10.1016/0370-2693\(87\)91197-X](https://doi.org/10.1016/0370-2693(87)91197-X)
- [12] Michael Dumbser, Dinshaw S. Balsara, Eleuterio F. Toro, and Claus-Dieter Munz. 2008. A unified framework for the construction of one-step finite volume and discontinuous Galerkin schemes on unstructured meshes. *J. Comput. Phys.* 227, 18 (2008), 8209 – 8253. <https://doi.org/10.1016/j.jcp.2008.05.025>
- [13] Michael Dumbser and Raphaël Loubère. 2016. A simple robust and accurate a posteriori sub-cell finite volume limiter for the discontinuous Galerkin method on unstructured meshes. *J. Comput. Phys.* 319 (2016), 163 – 199. <https://doi.org/10.1016/j.jcp.2016.05.002>
- [14] P. Galvez, J.-P. Ampuero, L. A. Dalguer, S. N. Somala, and T. Nissen-Meyer. 2014. Dynamic earthquake rupture modelled with an unstructured 3-D spectral element method applied to the 2011 M9 Tohoku earthquake. *Geophysical Journal International* 198, 2 (06 2014), 1222–1240. <https://doi.org/10.1093/gji/ggu203> arXiv:<https://academic.oup.com/gji/article-pdf/198/2/1222/1651569/ggu203.pdf>
- [15] Michael B. Giles. 2008. Multilevel Monte Carlo Path Simulation. *Oper. Res.* 56, 3 (May 2008), 607–617. <https://doi.org/10.1287/opre.1070.0496>
- [16] Heikki Haario and Eero Saksman. 1998. Adaptive Proposal Distribution for Random Walk Metropolis Algorithm. *Computational Statistics* 14 (07 1998). <https://doi.org/10.1007/s001800050022>
- [17] Heikki Haario, Eero Saksman, and Johanna Tamminen. 2001. An Adaptive Metropolis Algorithm. *Bernoulli* 7 (04 2001). <https://doi.org/10.2307/3318737>
- [18] Viet Hoang, Christoph Schwab, and Andrew Stuart. 2012. Complexity Analysis of Accelerated MCMC Methods for Bayesian Inversion. *Inverse Problems* 29 (07 2012). <https://doi.org/10.1088/0266-5611/29/8/085010>
- [19] Ole Klein, Olaf A. Cirpka, Peter Bastian, and Olaf Ippisch. 2017. Efficient geostatistical inversion of transient groundwater flow using preconditioned non-linear conjugate gradients. *Advances in Water Resources* 102 (2017), 161–177. <https://doi.org/10.1016/j.advwatres.2016.12.006>
- [20] Kody Law, Andrew Stuart, and Konstantinos Zygalakis. 2015. *Data Assimilation* (1st ed. 2015 ed.). Springer, Cham. Online–Resource (XVIII, 242 p. 61 illus., 41 illus. in color, online resource) pages. <https://doi.org/10.1007/978-3-319-20325-6>
- [21] Randall J. LeVeque, David L. George, and Marsha J. Berger. 2011. Tsunami modelling with adaptively refined finite volume methods. *Acta Numerica* 20 (2011), 211–289. <https://doi.org/10.1017/S0962492911000043>
- [22] Mikkel B. Lykkegaard, Grigorios Mingas, Robert Scheichl, Colin Fox, and Tim J. Dodwell. 2020. Multilevel Delayed Acceptance MCMC with an Adaptive Error Model in PyMC3. arXiv:2012.05668 [stat.CO]
- [23] Elizabeth Madden, Michael Bader, Jörn Behrens, Ylona van Dinther, Alice-Agnes Gabriel, Leonhard Rannabauer, Thomas Ulrich, Charsten Uphoff, Stefan Vater, Stephanie Wollherr, and Iris van Zelst. 2021. Linked 3D modeling of megathrust earthquake–tsunami events: from subduction to tsunami run up. *Geophysical Journal International* 224, 1 (Oct 2021), 487–516. <https://doi.org/10.1093/gji/ggaa484>
- [24] Youssef Marzouk and Dongbin Xiu. 2009. A Stochastic Collocation Approach to Bayesian Inference in Inverse Problems. *PRISM: NNSA Center for Prediction of Reliability, Integrity and Survivability of Microsystems* 6 (10 2009). <https://doi.org/10.4208/cicp.2009.v6.p826>
- [25] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. 1953. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics* 21, 6 (1953), 1087–1092. <https://doi.org/10.1063/1.1699114>
- [26] Matthew Parno, Andrew Davis, Linus Seelinger, and Youssef Marzouk. 2014. MIT uncertainty quantification (MUQ) library.
- [27] Leonhard Rannabauer, Michael Dumbser, and Michael Bader. 2018. ADER-DG with a-posteriori finite-volume limiting to simulate tsunamis in a parallel adaptive mesh refinement framework. *Computers and Fluids* 173 (2018), 299 – 306. <https://doi.org/10.1016/j.compfluid.2018.01.031>
- [28] Leonhard Rannabauer, Stefan Haas, Dominic Etienne Charrier, Tobias Weinzierl, and Michael Bader. 2018. Simulation of tsunamis with the exascale hyperbolic PDE engine ExaHyPE. In *Environmental Informatics: Techniques and Trends. Adjunct Proceedings of the 32nd edition of the EnviroInfo*. Shaker Verlag.
- [29] Anne Reinartz, Dominic E. Charrier, Michael Bader, Luke Bovard, Michael Dumbser, Kenneth Duru, Francesco Fambri, Alice-Agnes Gabriel, Jean-Matthieu Gallard, Sven Köppel, Lukas Krenz, Leonhard Rannabauer, Luciano Rezzolla, Philipp Samfass, Maurizio Tavelli, and Tobias Weinzierl. 2020. ExaHyPE: An engine for parallel dynamically adaptive simulations of wave problems. *Computer Physics Communications* 254 (2020), 107251. <https://doi.org/10.1016/j.cpc.2020.107251>
- [30] Daniel Rudolf and Björn Sprungk. 2015. On a Generalization of the Preconditioned Crank–Nicolson Metropolis Algorithm. *Foundations of Computational Mathematics* (04 2015). <https://doi.org/10.1007/s10208-016-9340-x>
- [31] Tatsuhiko Saito and Takashi Furumura. 2009. Three-dimensional tsunami generation simulation due to sea-bottom deformation and its interpretation based on the linear theory. *Geophysical Journal International* 178, 2 (08 2009), 877–888. <https://doi.org/10.1111/j.1365-246X.2009.04206.x> arXiv:<https://academic.oup.com/gji/article-pdf/178/2/877/5917374/178-2-877.pdf>
- [32] John Salvatier, Thomas V Wiecki, and Christopher Fonnesbeck. 2016. Probabilistic programming in Python using PyMC3. *PeerJ. Computer science* 2 (2016), e55.
- [33] Dongbin Xiu. 2015. *Stochastic Collocation Methods: A Survey*. 1–18. https://doi.org/10.1007/978-3-319-11259-6_26-1