EPSRC Maths Foresees Network Feasibility study report

Project title: Multilevel Monte Carlo Methods for Flood Risk Assessment

Investigators (names and affiliations): Alastair Gregory (Imperial), Colin J. Cotter (Imperial) and David Ham (Imperial)

Other participants: The Firedrake team at Imperial College London helped in the study, mostly by supporting code review for software packages generated for the use within the feasibility study. Rob Lamb of JBA Consulting also assisted in the project.

Project background: The goal of this project was to investigate the feasibility of using multilevel Monte Carlo (MLMC) methods for assessing flood risk based on shallow water inundation models. The focus was on the feasibility of the underlying MLMC methodology for flood risk assessment rather than on a specific discretisation approach for the shallow water equations; our conclusions are applicable to any suitable discretisation approach. The basic flood risk assessment problem can be defined as follows. Given a random model for rainfall in a modelled domain (such as an estuary), what is the probability distribution for flooding and how does it vary spatially over the domain. In a Monte Carlo method, we simply repeatedly sample rainfall patterns at random from the rain model, use them as sources in a shallow water model, and collect flooding statistics as Monte Carlo estimators that sum over all of the random samples. This approach is very challenging numerically. Monte Carlo statistics converge at a rate $O(1/\sqrt{N})$ (where N is the number of independent simulations); this can result in very demanding numerical computations since the shallow water model must be run many times.

This large cost can be addressed by the multilevel Monte Carlo (MLMC) method. Originally introduced by Giles (2008) [3] in the context of mathematical finance, MLMC has led to an explosion of applications in biology, physics, chemistry etc. (639 citations on Google Scholar). Indeed, the participants of this study, just last year, have had a paper accepted into SIAM SISC, applying MLMC to data assimilation [1]. The basic idea is to build a hierarchy of models on different resolutions. The finest resolution is our target model which is accurate but expensive, the coarser resolutions are inaccurate but cheap. The MLMC estimate combines a small number of samples from the finest level, with increasingly more samples at the coarser levels, trading off the bias due to larger model error with the variance error due to having less samples. This means that statistics can be estimated with much smaller computational effort. For example, Dodwell et al. (2015) [4] used a Multilevel MCMC algorithm to quantify risk in nuclear waste storage at the same cost as $O(10)$ fine model runs.

In this project, we have investigated the feasibility of applying MLMC to quantifying flood risk using a shallow water model. The goal was to provide a testbed where flood risk modellers can understand what speed-ups would be available using this multilevel framework. From this proof of concept, further projects could be started to implement this technology using operational flood risk tools. For the development of the model and testbed in this project, we have used an automated code generation library called Firedrake, developed at Imperial College, where multilevel infrastructure is seamlessly integrated. Firedrake allows for rapid development of finite element based models, by expressing the algorithm in a high level language. We expect that there is enough similarity between the finite element schemes used in this project and the finite volume methods in use by most operational codes so that our conclusions will be useful to operational code developers and their users.

Project objectives:

- To build a discontinuous Galerkin (DG) finite element discretized wetting and drying shallow-water model, using Firedrake, along with test cases. The model is based on the existing literature of Ern et al. (2007) [5] and Kuzmin (2013) [6]; the latter being using for the slope-limiting scheme in the model. As our project was aimed at the feasibility of applying the multilevel framework to our testbed problem, the choice of using algorithms in existing literature was made merely to find the quickest path to a working testbed code.

- To formulate a two-dimensional testbed problem with input (rainfall) given by some spatially correlated random field with magnitude given by a log-normal probabilistic distribution. This testbed represents a flooding bumpy river bank, and thus initially contains dry regions of domain.

- To evaluate the speed-ups, with respect to standard Monte Carlo, of multilevel Monte Carlo estimates of flood statistics, such as the mean flood depth at a certain coordinate on the river bank.

.

Key project outcomes: There have been two notable Python software packages, built on top of Firedrake, that have come from this project:

- **A Firedrake wetting/drying shallow-water model**: This package can run wetting/drying shallow-water testbed problems, with a wide-range of boundary condition and rainfall / source terms available. The package comes with a range of flooding test cases, some of which are commonplace in related literature.

Available at `https://github.com/firedrakeproject/flooddrake`

- **A multilevel Monte Carlo framework for Firedrake**: This package allows one to compute multilevel Monte Carlo approximations of certain probabilistic quantities (such as expectations and cumulative distribution functions) of solutions to systems expressed in Firedrake, including those simulated using the flood model listed above.

  Available at `https://github.com/firedrakeproject/firedrake-mlmc`

Together, these packages have been used to show the speed-ups in approximating flood statistics one can gain from using multilevel Monte Carlo, as opposed to standard Monte Carlo, for a suitable testbed problem. The uncertainty in this testbed problem comes from a source term (representing rainfall) given by a probability distribution. Thus different realisations of the testbed can be simulated by drawing a rainfall event from the distribution and running the shallow-water model over a fixed time period. Multilevel Monte Carlo approximations have been computed for this testbed problems by using a hierarchy of meshes, with the coarsest being 12×12 cells, and the finest being anything up to 96×96 (depending on the desired accuracy of the approximation). Each level of the hierarchy has four times the number of cells than that of the previous level. The standard Monte Carlo counterparts have been computed using just the corresponding finest mesh in the aforementioned hierarchy.

Profiles of the free-surface height (land bed height plus water depth) at the start and finish of one realisation of the testbed problem are given below.

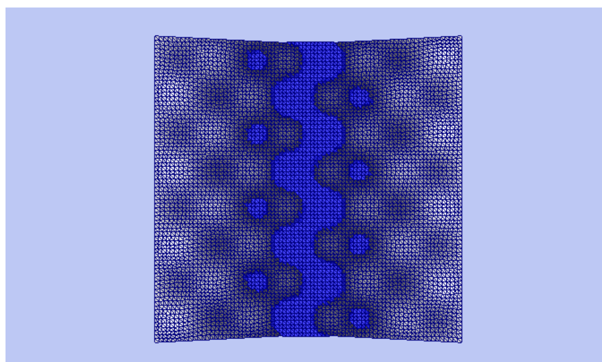

Figure 1: Profile of a realisation of the modelled testbed problem at the start of the time period.

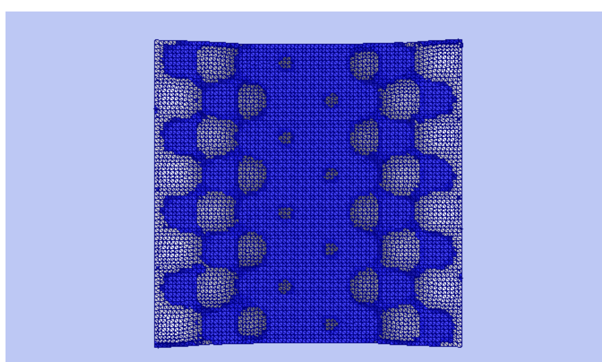

Figure 2: Profile of a realisation of the modelled testbed problem at the end of the time period.

The key results from this project are now summarised in Figure 3. The speed-up of the multilevel Monte Carlo approximation, as opposed to the standard Monte Carlo approximation with the same Root-Mean-Square-Error (RMSE), to the mean flooding depth at a certain coordinate in the domain is shown. The RMSE is estimated by evaluating the error against a very accurate standard Monte Carlo approximation of the mean flooding depth, with very large sample size and fine resolution (a mesh with 128×128 cells). The multilevel Monte Carlo method shows speed-ups of an order of magnitude of the RMSE of the approximation, conforming with theory in Cliffe (2011) [7].

There is a Wiki page at `https://github.com/firedrakeproject/flooddrake/wiki`, summarising how a user can use Flooddrake (the Firedrake wetting/drying shallow-water model) to simulate their own testbed problems and find

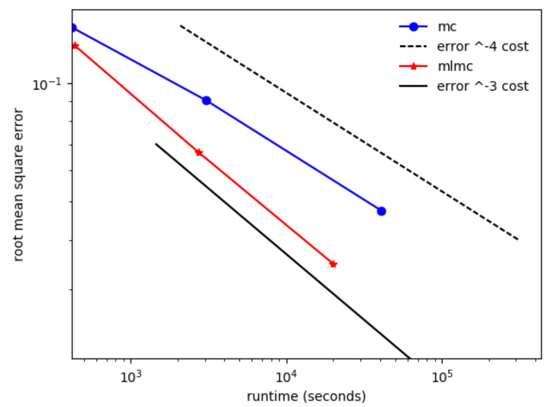

Figure 3: Speed-up (root-mean-square-error vs runtime in seconds) of approximating the mean flooding depth at a certain coordinate of the domain using the multilevel Monte Carlo method as opposed to standard Monte Carlo.

speed-ups in approximating probability distributions with respect to flood quantities using multilevel Monte Carlo. It also has a summary of the aforementioned main testbed problem designed for this project and the results of applying multilevel Monte Carlo to the uncertainty quantification of it.

Potential for initiating or developing future multidisciplinary collaborations: The Flooddrake tool has been mentioned throughout the recent EPSRC Programme grant proposal Maths Fights Floods, as well as the multilevel Monte Carlo technique. There is great potential for investigating this methodology in more challenging and realistic geometries in collaboration with the JBA Trust and EA.

References:

[1] A. Gregory, C. J. Cotter and S. Reich (2016). Multilevel Ensemble Transform Particle Filtering. *SIAM Journal of Scientific Computing.* **38**.

[2] A. Gregory and C. J. Cotter (2016). A Seamless Multilevel Ensemble Transform Particle Filter. *arXiv:1611.00266*.

[3] M. Giles (2008). Multilevel Monte Carlo Path Simulation. *Operations Research.* **56**, 607–617.

[4] T. J. Dodwell, C. Ketelsen, R. Scheichl and A. L. Teckentrup (2015) A Hierarchical Multilevel Markov Chain Monte Carlo Algorithm with Applications to Uncertainty Quantification in Subsurface Flow. *SIAM Journal on Uncertainty Quantification.* **3**, 1075–1108.

[5] A. Ern, S. Piperno and K. Djadel (2007). A well-balanced Runge–Kutta Discontinuous Galerkin method for the Shallow-Water Equations with flooding and drying. *Internation Journal for Numerical Methods in Fluids.* **58**, 1–25.

[6] D. Kuzmin (2013). Slope limiting for discontinuous Galerkin approximations with a possibly non-orthogonal Taylor basis. *Internation Journal for Numerical Methods in Fluids.* **71**, 1178–1190.

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

.