Assessing erosion and flood risk in the coastal zone through the application of multilevel Monte Carlo methods

Mariana C. A. Clare[a,*], Matthew D. Piggott[a] and Colin J. Cotter[b]

[a]*Department of Earth Science and Engineering, Imperial College London, Exhibition Road, London SW7 2BU, UK.*
[b]*Department of Mathematics, Imperial College London, Exhibition Road, London SW7 2BU, UK.*

ARTICLE INFO

ABSTRACT

The risk from erosion and flooding in the coastal zone has the potential to increase in a changing climate. The development and use of coupled hydro-morphodynamic models is therefore becoming an ever higher priority. However, their use as decision support tools suffers from the high degree of uncertainty associated with them, due to incomplete knowledge as well as natural variability in the system. Here we show for the first time how the multilevel Monte Carlo method (MLMC) can be applied to hydro-morphodynamic models, in this case XBeach, to quantify uncertainty by computing statistics of output variables given uncertain input parameters. MLMC accelerates the Monte Carlo approach through the use of a hierarchy of models with different levels of resolution. A variety of theoretical and real-world coastal zone case studies are considered, for which output variables that are important to the assessment of flood and erosion risk are estimated, such as wave run-up height and total eroded volume. We show that MLMC can significantly reduce computational cost, resulting in speed up factors of 40 or greater compared to a simple Monte Carlo approach, whilst still maintaining the same level of accuracy. Furthermore, MLMC is used to estimate the cumulative distribution of these output variables for given uncertain parameters. This allows the risk of a variable exceeding a certain value to be calculated, for example the risk of the wave run-up height exceeding the height of a physical structure such as a seawall; this is a useful capability to inform decision-making processes.

This manuscript is a non-peer reviewed EarthArXiv preprint that has been submitted for publication in Coastal Engineering

---

1. Introduction

For millennia, the world's coastal zones have been at risk from erosion and flooding. However, climate change will in turn lead to hydrodynamic changes in the coastal zone resulting from location-dependent sea level rise and changes in the strength and frequency of storms, which may severely exacerbate these risks and the damage associated with them. Estimates suggest the economic cost of coastal flooding damage in Europe will be two to three orders of magnitude greater than current levels by the end of the century (Vousdoukas et al., 2018). Thus the modelling of coastal zones is a vital endeavour.

For several decades, complex hydro-morphodynamic models have been used to simulate flooding and erosion. Whilst these models may produce accurate results for a given idealised scenario, there is a large degree of uncertainty associated with them when they are applied in the real world, coming from both the input data as well as the models themselves (see for example Unguendoli (2018) and Villaret et al. (2016)). This uncertainty can be due to either incomplete knowledge or natural variability (Apel et al., 2004) and can be quantified by combining numerical hydro-morphodynamic models with statistical frameworks.

In this work, the depth-averaged finite-volume based coastal ocean model XBeach (Roelvink et al., 2015) is used as our model because of its ability to simulate wave propagation, flow, sediment transport and morphodynamic changes. XBeach has already been used within a Monte Carlo framework, for example in Harris et al. (2018), Pender and Karunarathna (2013) and Simmons et al. (2017). However, like all complex hydro-morphodynamic models, XBeach is computationally relatively expensive and thus does not lend itself well to Monte Carlo simulations and hazard assessments, which require a large number of individual model runs. For example, in Harris et al. (2018), a Monte Carlo

✉ m.clare17@imperial.ac.uk (M.C.A. Clare)
ORCID(s): 0000-0002-5010-0363 (M.C.A. Clare); 0000-0002-7526-6853 (M.D. Piggott); 0000-0001-7962-8324 (C.J. Cotter)

---

simulation with 240,000 individual runs of the 1D version of XBeach was required to perform the desired study. This computational cost has limited the problems that Monte Carlo based studies can be used to address and has meant that for more complex and long-term test cases, researchers have had to use simplified models. For example, Callaghan et al. (2013) used the simpler 1D semi-empirical model SBeach in their Monte Carlo simulation which took 40 days; the authors of Callaghan et al. (2013) estimate that even using the 1D version of XBeach would have taken four and a half millennia.

Instead of reverting to simpler, less costly models, we use the relatively novel multilevel Monte Carlo method (MLMC), first presented in Giles (2008). MLMC seeks to reduce computational cost by accelerating the Monte Carlo approach through the use of a hierarchy of model configurations, each with a different level of resolution. It has been used in areas as diverse as mathematical finance (Giles, 2008), data assimilation (Gregory and Cotter, 2017a) and atmospheric dispersion modelling (Katsiolides et al., 2018). In Sánchez-Linares et al. (2016), MLMC was successfully applied to estimate uncertainty in landslide generated tsunamis using a simple shallow-water like model. MLMC quantifies uncertainty by computing estimators for the expectation of discretised random variables for uncertain parameters. In this work we present an initial exploration of how MLMC can be applied to the hydro-morphodynamic model XBeach, to investigate the impact of a variety of uncertain input parameters in both theoretical and real-world test cases. To the best of our knowledge this is the first application of MLMC with complex hydro-morphodynamic models.

In coastal problems, we are not only interested in the expected value of a variable but also in the probability of a variable exceeding a certain value. This quantity can be expressed as an expectation: $\mathbb{E}[\mathbb{1}_{X \geq x}]$, where, for example, X is the maximum horizontal inundation distance and x the location of a physical structure, and $\mathbb{1}$ represents the indicator function. Thus $\mathbb{E}[\mathbb{1}_{X \geq x}]$ gives the probability, or the *risk* that the physical structure at x will be flooded. However, as discussed in Giles (2015), MLMC struggles with binary output variables, because a large number of samples are required to ensure accurate variance estimates. A number of different methods have been proposed to deal with this problem including smoothing (Giles et al., 2015) and selective refinement (Elfverson et al., 2016). In this work, we use the ensemble generating method from Gregory and Cotter (2017b) which we discuss further in Section 4. Unlike other approaches, this method generates the entire cumulative distribution function for the output variable, which is of much greater value when assessing risk than a single statistic. Furthermore, it does so in a simple and computationally efficient manner.

The remainder of this paper is structured as follows: in Section 2, we briefly outline the relevant MLMC theory; in Section 3, we integrate MLMC with XBeach and run a series of test cases; in Section 4, we estimate the cumulative distribution function of the output variable using MLMC and finally in Section 5, we present conclusions from this work.

2. Multilevel Monte Carlo method

The simplest method to estimate the expectation of a random variable $X(\omega)$ (where ω is a single sample from the sample space Ω) is to use a Monte Carlo method and take the average of $X(\omega)$ for N independent samples from Ω. However, the Monte Carlo method has an order of convergence of $N^{-1/2}$ (Caflisch, 1998) meaning to achieve an accuracy of ϵ requires $O(\epsilon^{-2})$ samples. This makes the simulation very computationally expensive. The multilevel Monte Carlo method (MLMC) was first introduced in Giles (2008). In this section, following Giles (2008) and Giles (2015), we discuss MLMC and its use in computing estimators for the expectation of discretised random variables given uncertain input parameters. We refer the reader to Giles (2008) and Giles (2015) for more details.

Whereas the Monte Carlo method considers the model at just one resolution, MLMC accelerates the Monte Carlo method by using a hierarchy of models at different levels of resolution. The fundamental idea underlying MLMC comes from the linearity of expectations: for two variables X_l and X_{l-1},

$$\mathbb{E}[X_l] = \mathbb{E}[X_{l-1}] + \mathbb{E}[X_l - X_{l-1}], \tag{1}$$

where $\mathbb{E}(\cdot)$ denotes the expectation. Extending this idea

$$\mathbb{E}[X_L] = \mathbb{E}[X_0] + \sum_{l=1}^{L} \mathbb{E}[X_l - X_{l-1}]. \tag{2}$$

In the MLMC framework, X_l denotes the numerical approximation to X on level l of the multilevel environment and thus X_0 and X_L denote the approximation on the coarsest and finest level respectively. Within the framework, each

level l is defined by its numeric mesh element size

$$h_l = M^{-l}T, \tag{3}$$

where T is the total length of the domain and M the integer factor the mesh element size is refined by at each level (following standard practice, we use $M = 2$ throughout). Thus as l increases, the mesh becomes more refined. Note the domain can be multi-dimensional in which case there will be an h_l for every dimension of T.

In MLMC, the expectation estimator is defined as

$$\hat{Y} = \sum_{l=0}^{L} \hat{Y}_l, \tag{4}$$

where $\hat{Y}_l$ is the difference estimator for $\mathbb{E}[X_l - X_{l-1}]$ defined as

$$\hat{Y}_l = \begin{cases} N_0^{-1} \sum_{i=1}^{N_0} X_0^{(i)} & l = 0, \\ N_l^{-1} \sum_{i=1}^{N_l} \left(X_l^{(i)} - X_{l-1}^{(i)} \right) & l > 0, \end{cases} \tag{5}$$

where N_l is the number of samples at each level $(l, l-1)$ pair. Here the same random numbers are used to construct the variables X_l and X_{l-1} in order to minimize the variance of the difference estimator and hence the overall error. To ensure independence, different independent samples are used at each level, meaning $\mathrm{Cov}(\hat{Y}_i, \hat{Y}_j) = 0$ if $i \neq j$ and thus

$$\mathbb{V}(\hat{Y}) = \mathbb{V}\left(\sum_{l=0}^{L} \hat{Y}_l \right) = \sum_{l=0}^{L} N_l^{-1} V_l, \tag{6}$$

where $\mathbb{V}(\cdot)$ denotes the variance and V_l is the variance of the sample $X_l^{(i)} - X_{l-1}^{(i)}$. The overall cost of calculating the expectation estimator (4) is given by $\sum_{l=0}^{L} N_l C_l$ where C_l is the cost per level l.

To calculate the error of the estimator (4), we note that root mean square error (RMSE) is defined as

$$\mathrm{RMSE} = \sqrt{\mathbb{E}[(\hat{Y} - \mathbb{E}[X])^2]}. \tag{7}$$

This is simplified by noting $\hat{Y}$ is an unbiased estimator of $\mathbb{E}[X_L]$ and so

$$\mathrm{RMSE} = \sqrt{\mathbb{E}[(\hat{Y} - \mathbb{E}[X_L])^2] + (\mathbb{E}[X_L] - \mathbb{E}[X])^2}. \tag{8}$$

Here $\mathbb{E}[(\hat{Y} - \mathbb{E}[X_L])^2]$ is equivalent to (6) and is the Monte Carlo error; and $(\mathbb{E}[X_L] - \mathbb{E}[X])^2$ is the square of the bias and is the numerical discretisation error.

A bound on the RMSE (8) is provided by the key complexity theorem for MLMC. We outline it briefly here – the full details can be found in Giles (2015) with a proof given in detail in Cliffe et al. (2011). Briefly the theorem states that if the following conditions are met

$$\mathbb{E}[X_l - X] \leq c_1 2^{-\alpha l}, \tag{9a}$$

$$\mathbb{E}[\hat{Y}_l] = \begin{cases} \mathbb{E}[X_0], & l = 0, \\ \mathbb{E}[X_l - X_{l-1}], & l > 0, \end{cases} \tag{9b}$$

$$V_l \leq c_2 2^{-\beta l}, \tag{9c}$$

$$C_l \leq c_3 2^{\gamma l}, \tag{9d}$$

where $\alpha \geq 1/2 \min(\beta, \gamma)$, and $\alpha, \beta, \gamma, c_1, c_2, c_3$ are positive constants, then there exists a finest level resolution L and a finite number of samples N_l such that for any $\epsilon < e^{-1}$, the estimator $\hat{Y} = \sum_{l=0}^{L} \hat{Y}_l$ has

$$\mathrm{RMSE} \leq \epsilon. \tag{10}$$

Furthermore, the total computational cost of calculating this estimator with this RMSE is

$$
C \leq \begin{cases} c_4 \epsilon^{-2}, & \beta > \gamma, \\ c_4 \epsilon^{-2} (\log \epsilon)^2, & \beta = \gamma, \\ c_4 \epsilon^{-2-(\gamma-\beta)/\alpha}, & 0 < \beta < \gamma. \end{cases} \tag{11}
$$

where c_4 is a positive constant.

Note (9a), (9c) and (9d) represent a bound on the bias, variance and cost respectively at each level and (9b) follows from the definition in (5). As the mesh is refined, β controls the decay in the variance and γ the growth of the cost. Thus if $\beta > \gamma$, most of the computational cost is at the coarser levels and if $\beta < \gamma$ it is mostly at the finest levels. This means if $\beta < \gamma$, the order of convergence is better than the standard Monte Carlo method and if $\beta > \gamma$, it is equivalent to the Monte Carlo method. The values of α, β and γ are found using convergence tests: the gradient of the line of best fit of $\mathbb{E}(X_l - X_{l-1})$ against l is α, $\mathbb{V}(X_l - X_{l-1})$ against l is β and timed cost against l is γ.

A key part of MLMC is determining the optimum N_l to achieve the convergence stated in (10). Following standard practice (for example Giles (2008)), the optimum value of N_l is determined by using the Euler-Lagrange method to minimize the overall cost with respect to the fixed overall variance $\epsilon^2/2$. It can be shown that the optimum number of samples at each level is

$$
N_l = \left\lceil 2\epsilon^{-2} \sqrt{\frac{V_l}{C_l}} \left(\sum_{l=0}^{L} \sqrt{V_l C_l} \right) \right\rceil. \tag{12}
$$

This formula requires a knowledge of the variance and cost at level l. Following the proof in Cliffe et al. (2011), we instead use the conditions in (9) to create an upper bound on N_l (for example (9c) and (9d) mean that $V_l C_l < c_2 c_3 2^{(\gamma-\beta)l}$) and set this to be the optimum number of samples. Thus

- if $\beta > \gamma$:

$$
N_l = \lceil 2\epsilon^{-2} c_2 \left(1 - 2^{-(\beta-\gamma)/2} \right)^{-1} 2^{-(\beta+\gamma)l/2} \rceil, \tag{13}
$$

- if $\beta < \gamma$:

$$
N_l = \lceil 2\epsilon^{-2} c_2 2^{(\gamma-\beta)L/2} \left(1 - 2^{-(\gamma-\beta)/2} \right)^{-1} 2^{-(\beta+\gamma)l/2} \rceil, \tag{14}
$$

- if $\beta = \gamma$:

$$
N_l = \lceil 2\epsilon^{-2} (L+1) c_2 2^{-\beta L} \rceil. \tag{15}
$$

2.1. MLMC verification checks

Accompanying MLMC are the following checks, detailed in Giles (2015), to ensure the method implementation is both correct and appropriate.

- **Convergence check:** The convergence condition (9a) states that $\mathbb{E}[X_l - X_{l-1}] \propto 2^{-\alpha l}$ and thus it follows that the bias $\mathbb{E}[X - X_L] = \frac{\mathbb{E}[X_L - X_{L-1}]}{2^\alpha - 1}$. Therefore we use the following as a convergence test

$$
\frac{\left| \mathbb{E}[X_L - X_{L-1}] \right|}{2^\alpha - 1} < \frac{\epsilon}{\sqrt{2}}, \tag{16}
$$

where α is determined by calculating the gradient of the line of best fit of $\mathbb{E}(X_l - X_{l-1})$ against l. If this check fails, this means MLMC has not converged to within the specified tolerance – this may mean there is a mathematical or programming error or simply that the value of ϵ chosen is too small for the number of levels in the MLMC algorithm.

- **Consistency check:** Suppose a, b and c are estimates for $\mathbb{E}[X_{l-1}]$, $\mathbb{E}[X_l]$ and $\mathbb{E}[Y_l]$ respectively. Then $a - b + c \approx 0$ up to the expected error from Monte Carlo sampling. Furthermore, $\sqrt{\mathbb{V}[a - b + c]} \leq \sqrt{\mathbb{V}(a)} + \sqrt{\mathbb{V}(b)} + \sqrt{\mathbb{V}(c)}$. Therefore if calculated correctly the following ratio

$$\frac{|\mathbb{E}[X_l - X_{l-1}] - \mathbb{E}[X_l] + \mathbb{E}[X_{l-1}]|}{3 \left(\sqrt{\mathbb{V}(X_l - X_{l-1})} + \sqrt{\mathbb{V}(X_l)} + \sqrt{\mathbb{V}(X_{l-1})} \right)}, \tag{17}$$

  has a probability of 0.3% of being greater than 1 (see Giles (2015)). If this ratio is greater than 1, this indicates that there is a difference between how X_l and X_{l-1} are computed and not just that they are on a different mesh.

- **Kurtosis check:** The kurtosis, κ, of $X_l - X_{l-1}$ determines the number of samples required for a good variance estimate. If κ is very large this indicates that the variance estimate is poor and that either the number of samples used is insufficient or the method chosen for sampling is incorrect. Thus following standard practice, if the kurtosis is greater than 100, an error is raised.

We use these checks to evaluate the coding and mathematical implementation of the test cases in this work, and to identify any inherent problems with applying MLMC to these test cases.

2.2. MLMC algorithm

We have now outlined the basic MLMC method and conclude this section with a statement of the MLMC algorithm used in this work, modified from Giles (2008):

Algorithm 1 Multilevel Monte Carlo Method

1: Start with $L = 0$
2: Estimate the variance V_L using an initial estimate for the number of samples N_L
3: Define optimal N_l for $l = 0, ..., L$ using (12)
4: If the optimal N_l is greater than the number of samples you already have, evaluate the extra samples needed
5: If $L \geq 2$ test for convergence i.e. check if (16) is satisfied
6: If $L < 2$ or the algorithm has not converged set $L := L + 1$ and return to Step 2

---

3. Applying MLMC to XBeach

In this section, we apply MLMC to the XBeach model in order to estimate the expectation of random variables in test cases given uncertain input data. We apply MLMC through a Python wrapper around the XBeach model. We have also parallelised the wrapper meaning that, within the same level, multiple XBeach runs can be performed at the same time, thus increasing the efficiency of the algorithm. Note our MLMC wrapper is written so that it can easily be applied to other models in further work. We begin this section by giving a brief overview of the XBeach model.

3.1. XBeach model

XBeach is a depth-averaged finite-volume coastal ocean model which simulates near-shore hydrodynamics and morphodynamics (Roelvink et al., 2015). Surface elevation and flow are modelled using the Generalised Lagrangian Mean formulation of the depth-averaged shallow water equations (Andrews and McIntyre, 1978), enabling the modelling of short wave-induced mass fluxes and return flows. These equations are:

$$\frac{\partial \mathbf{u}^L}{\partial t} + \mathbf{u}^L \cdot \nabla \mathbf{u}^L + \mathbf{f} \wedge \mathbf{u}^L - \nu \nabla^2 \mathbf{u}^L = \frac{\tau_s}{\rho h} - \frac{\tau_b^E}{\rho h} - g \nabla \eta + \frac{\mathbf{F}}{\rho h}, \tag{18}$$

$$\frac{\partial \eta}{\partial t} + \nabla \cdot (h \mathbf{u}^L) = 0, \tag{19}$$

where $\mathbf{u}^L$ is the Lagrangian velocity defined in Roelvink et al. (2015) as the distance travelled by an individual water particle during one wave period divided by the wave period, f the Coriolis vector, ν the viscosity, τ_s and τ_b^E the wind

---

and bed shear stresses respectively where E denotes Eulerian velocity defined in Roelvink et al. (2015) as the fixed point short-wave-averaged velocity, η the elevation, $\mathbf{F}$ the wave-induced stresses, ρ the density and h the water depth.

Waves can be modelled using several different options in XBeach. In this work, the stationary wave model is used in Section 3.2 and the surfbeat model in Section 3.3. Both methods solve the following wave action equation for short wave energy

$$
\frac{\partial A}{\partial t} + \frac{\partial c_x A}{\partial x} + \frac{\partial c_y A}{\partial y} + \frac{\partial c_\theta A}{\partial \theta} = -\frac{D_w + D_f + D_v}{\sigma},
\tag{20}
$$

but in the stationary mode, $\frac{\partial A}{\partial t} = 0$. Here c_x is the group velocity in the x-direction, c_y the group velocity in the y-direction, c_θ the refraction speed, D_w the dissipation by wave breaking, D_f, the dissipation by bottom friction, D_v the dissipation by vegetation and A is the wave energy. Note the formulations of D_w and D_f are different in the stationary and surfbeat mode (see Bart (2017) for further details).

Finally, to model morphological changes in Section 3.3.1, XBeach uses the following standard depth-averaged advection-diffusion equation to model sediment transport

$$
\frac{\partial hC}{\partial t} + \nabla \cdot (hC\mathbf{u^E}) + \nabla \cdot (D_h h \nabla C) = \frac{hC_{\text{eq}} - hC}{T_s},
\tag{21}
$$

where C is the depth-averaged sediment concentration, $\mathbf{u^E}$ the Eulerian velocity, D_h the diffusivity coefficient, C_{eq} the equilibrium sediment concentration and T_s the adaptation time; and the Exner equation to model bed changes

$$
(1 - p)\frac{\partial z_b}{\partial t} + f_{\text{mor}}\nabla \cdot \mathbf{q} = 0,
\tag{22}
$$

where z_b is the bed, p the bed porosity, f_{mor} the morphological acceleration factor used to artificially increase the rate at which the bed changes compared with the underlying hydrodynamics and $\mathbf{q}$ the sediment transport rates.

In this work, we use version 1.23.5526 of the XBeachX release and unless explicitly stated, parameter values are left at the default for this version.

3.2. Uncertainty in bed slope angle

We now apply the MLMC framework from Section 2 to the XBeach model described above. In the following test cases, we evaluate the uncertainty associated with the beach bed profile by considering the bed slope angle θ to be uncertain. This represents a significant source of uncertainty, as discussed in Unguendoli (2018).

3.2.1. One-dimensional (1D) bed slope test case

For the first test case, we consider a simple 1D hydrodynamics-only problem of a stationary wave approaching the beach with a wave period of 10 s and a maximum wave amplitude of 2.5 m and each simulation is run for 200 s. The beach slope starts at 750 m and the bed set-up is depicted in Figure 1.

The distribution of θ is set to be $\theta \sim \mathcal{U}(\arctan(0.035), \arctan(0.5))$. The lower bound is chosen to ensure that there is always an area of the bed which is above the initial water level. We seek the expected value of the maximum horizontal inundation distance during the simulation. In order to satisfy the constraints on ϵ required to achieve the convergence stated in (10), the inundation distance is normalised by dividing it by the total x-length of the domain.

Before running the full MLMC algorithm (Algorithm 1), we first run Steps 1 and 2, in order to check that the MLMC approach is appropriate for this test case. We use 7 levels ($l = 2$ to 8) and 500 samples at each level. Figure 2a, a log plot of $\mathbb{V}(X_l - X_{l-1})$ against level, shows that the variance decreases as the level number l increases, i.e. as the mesh gets finer. This observation is important because, recalling (6), this means fewer samples are needed on the finer meshes. The figure also shows that $\mathbb{V}(X_l - X_{l-1})$ is lower than $\mathbb{V}(X_l)$ for all l. This is also important because otherwise a lower variance could be achieved by using the same number of samples with a simple Monte Carlo simulation. Thus, we expect the MLMC simulation to be less computationally expensive than a Monte Carlo simulation meaning the choice of the MLMC approach for this test case is justified.

Figure 2b shows as expected that the cost calculated by timing the code increases as the level number l increases (note here this is the time taken for the code to run in parallel). This is because for the preliminary MLMC run, a fixed number of sample points is being used and, at finer grid levels, XBeach is more computationally expensive. Using these figures, as well as the log plot of $\mathbb{E}(X_l - X_{l-1})$ against level l, convergence tests estimate the convergence values
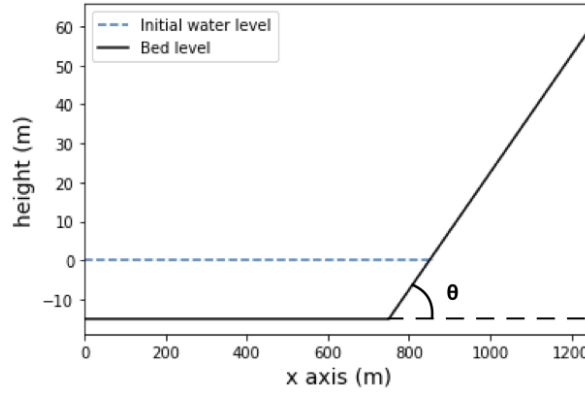
Figure 1: Schematic of bed with uncertain bed slope angle θ for 1D bed slope test case.
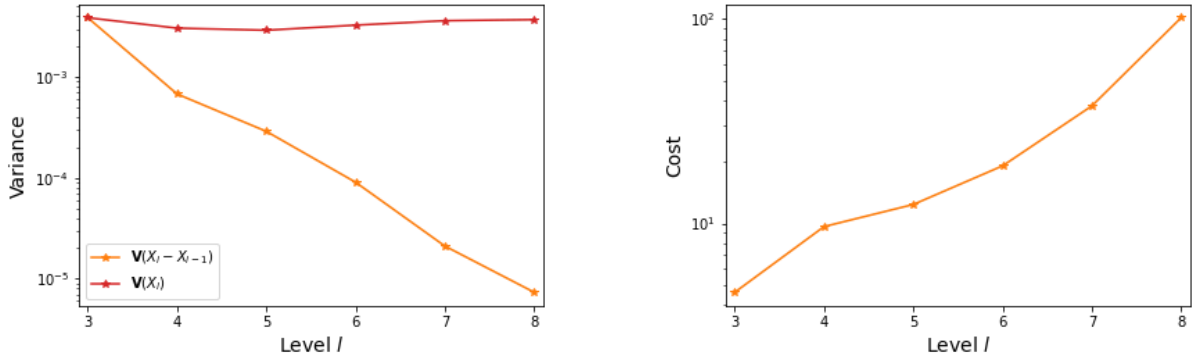
(a) Variance of difference, $X_l - X_{l-1}$, compared to variance of single variable, X_l at each level.

(b) Increasing computational cost (s) as the mesh is refined (i.e. l increases).

Figure 2: Preliminary MLMC results for 1D bed slope test case showing the use of MLMC is justified.

defined in (9) as $\alpha = 1.29; \beta = 1.89; \gamma = 1.10$ for this test case. Thus the numerical discretisation is approximately first order accurate (based upon the value of α) and as $\beta > \gamma$ most of the computational cost in the MLMC algorithm will come from the coarser levels. This means that we expect $C \propto \epsilon^{-2}$, which is the same as for Monte Carlo but we still expect MLMC to be computationally faster due to the improvements in the variance.

These observations from the preliminary study, as well as the fact that the consistency and kurtosis checks from Section 2.1 passed, means we can conclude that MLMC can be used successfully with XBeach and we can run the full algorithm. For this test case, the chosen range of tolerance values are $\epsilon = [0.0002, 0.0005, 0.0008, 0.0012, 0.002, 0.003]$ recalling that ϵ controls the bound on the convergence test (16). As $\beta > \gamma$, we use these ϵ values in the formula (13) to derive the optimum number of samples required at each level, N_l. We also set a maximum limit on the finest level that l can reach of $L_{\max} = 13$ (see Step 6 of Algorithm 1). Figure 3a shows the optimal N_l and most importantly that as the level number increases (*i.e.* as the grid becomes finer), fewer samples are needed, which makes MLMC more efficient than a simple Monte Carlo method.

From these N_l, we use the formula given in Giles (2008)

$$C = \left(1 + \frac{1}{2}\right) \sum_{l=0}^{L} N_l \, 2^{l+1}, \tag{23}$$

to calculate the theoretical cost C of the simulation for each ϵ. Figure 3b shows that both the gradient of the theoretical cost and the timed cost agree with the expected gradient $C \propto \epsilon^{-2}$ and thus the code has been correctly applied.
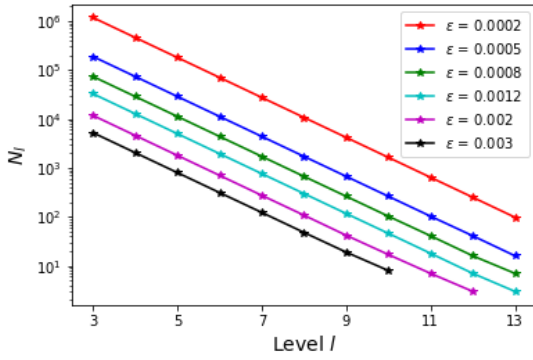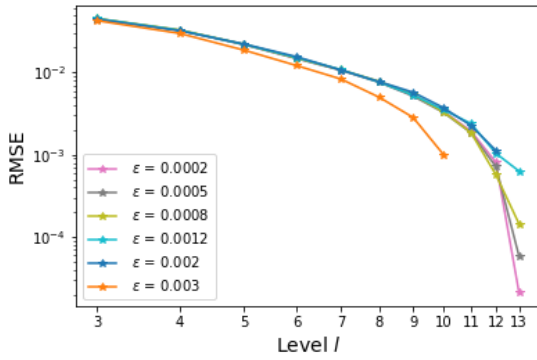
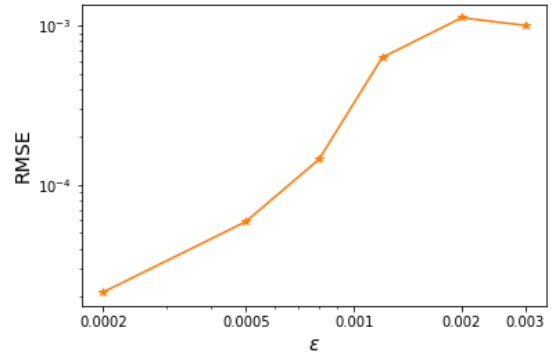*(a) Optimum number of samples required at each l for given tolerance value ε.*
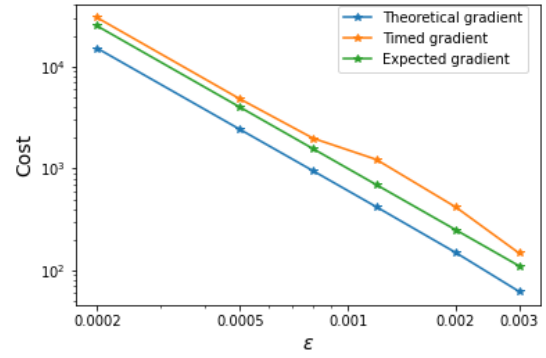
(b) Verification that both theoretical cost (23) and timed cost vary with ε at expected rate.

Figure 3: Effect of tolerance value ϵ on cost and sample size for 1D bed slope test case.

(a) RMSE of MLMC at each l for given tolerance value ε.

(b) RMSE of MLMC once it has converged for each tolerance value ε.

Figure 4: RMSE of MLMC (24) for 1D bed slope test case for varying tolerance value, ϵ.

To test the accuracy of the MLMC estimate, we also calculate the expected value using a Monte Carlo simulation at the finest mesh considered by the MLMC simulation ($\Delta x = 0.153$ m). According to Monte Carlo the expected maximum horizontal inundation distance divided by the total domain length is 0.6785 (to four significant figures). As the total domain length is 1250 m, this in fact means the water inundates a horizontal distance of 848.1 m. This agrees up to four significant figures with the value found using the MLMC simulation using $\epsilon = 0.0002$. We can calculate the root mean square error of the MLMC simulation compared to the Monte Carlo simulation using

$$\text{RMSE of MLMC} = \sqrt{\left(\mathbb{E}[X_L] - \mathbb{E}[X_{MC}]\right)^2}. \tag{24}$$

Figure 4a shows that as the mesh becomes finer, the RMSE decreases for all ϵ values considered. In addition, Figure 4b shows that as ϵ decreases, in general the final expected value becomes closer to the real solution showing that the convergence condition (16) has worked as expected. Note the only exception to this trend is for $\epsilon = 0.002$ and $\epsilon = 0.003$. This is likely because at these values of ϵ, the optimal number of samples is quite low (see Figure 3a) and so the variance approximations are likely to not be completely accurate.

Finally Figure 5 compares the total cost and accuracy for MLMC versus Monte Carlo. It shows that, as expected, the Monte Carlo simulation does not converge uniformly to the final value. By contrast, the figure shows MLMC convergence is much more uniform, due to the use of the convergence test in Section 2.1 and optimum number of
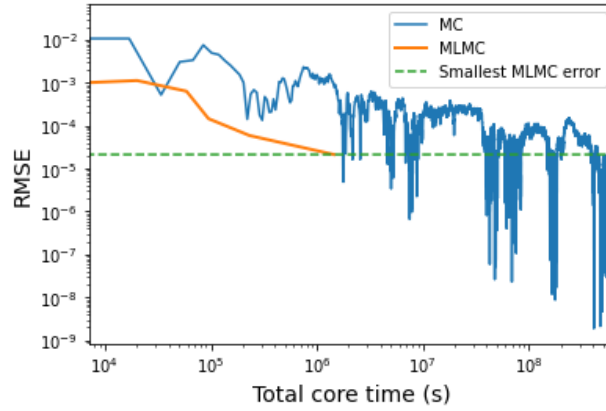
Figure 5: Comparison of total cost and accuracy when using MLMC and Monte Carlo (MC) for 1D bed slope test case.

samples calculation. Thus, unlike with Monte Carlo, for MLMC we do not need to conduct further samples to check whether the algorithm has really converged, thus decreasing the overall computational cost. More significantly, the figure also shows that it takes the Monte Carlo method more than 110,400 core hours (4.0×10^8 s) to consistently achieve an RMSE lower than that achieved by MLMC with $\epsilon = 0.0002$ in 768 core hours (2.8×10^6 s), a speed-up factor of over 140.

Although this test case is somewhat abstract, it demonstrates that MLMC can be used with XBeach for a wave with known properties approaching a 1D beach with unknown slope. We have shown that the key advantage of MLMC is that it can achieve results considerably faster when compared to using a simple Monte Carlo approach. This means we can confidently move to more realistic and complicated 2D problems.

3.2.2. Two-dimensional (2D) bed slope test case

In order to show the preceding promising results are not restricted to the simple 1D case, we briefly consider the more complex test case of a 2D beach. We use the same parameters as in Section 3.2.1 and modify the bed slope angle, θ, to apply in both the x and y direction, as shown in Figure 6. We again choose a distribution of $\theta \sim \mathcal{U}$ (arctan 0.035, arctan(0.5)) and seek to find the expected value of the maximum horizontal inundation distance along the x-axis that the water reaches during the simulation for all values of y. The results are again normalised by dividing by the total x-length of the domain, 1250 m.

To check whether MLMC can be used for this test case, we first run steps 1 and 2 of Algorithm 1. Figure 7 shows that $\mathbb{V}(X_l - X_{l-1})$ decreases as the mesh gets finer and that for all l it is lower than $\mathbb{V}(X_l)$, meaning fewer samples are needed on finer meshes. Using convergence tests, the values in (9) are found to be $\alpha = 1.18, \beta = 1.69, \gamma = 2.93$. Thus the numerical discretisation is approximately of the same order as the 1D test case. As $\gamma > \beta$, most of the computational cost will come from the finer levels and we expect the order of convergence to be better than the standard Monte Carlo method.

The simulation also passes the consistency and kurtosis checks from Section 2.1 and thus the preliminary application of MLMC to a 2D test case in XBeach is successful. To run the full MLMC algorithm, we choose tolerance values of $\epsilon = [0.001, 0.0017, 0.0025, 0.0035, 0.005, 0.006]$ and set the maximum limit on the finest level l can reach to be $L_{max} = 8$. Note that we choose larger values of ϵ here than in Section 3.2.1 so that we do not need to consider as fine a mesh here as in Section 3.2.1. Finer meshes are not an issue for MLMC but using a finer mesh than $L_{max} = 8$ would make the computational cost of the Monte Carlo simulation later in the section unmanageable. As $\gamma > \beta$, the optimum N_l are calculated using (14) and are shown in Figure 8a. Importantly the number of samples required decreases as the mesh becomes finer.

For this test case $\gamma > \beta$ and so following (11), the cost of the algorithm C should be such that $C \propto \epsilon^{-2-(\gamma-\beta)/\alpha}$. Figure 8b shows this is roughly true for both the gradient of the timed cost and the theoretical cost (calculated using (23)).
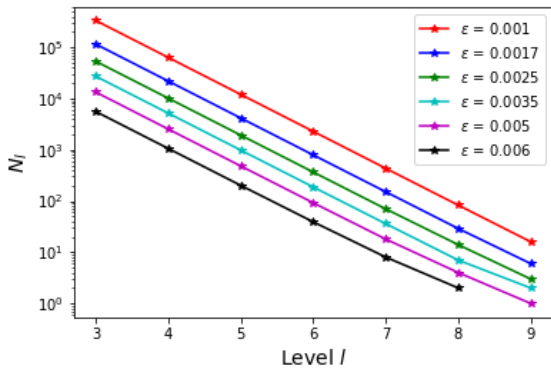
To calculate the RMSE of MLMC for this test case, we estimate the real solution $\mathbb{E}[X]$ using a Monte Carlo
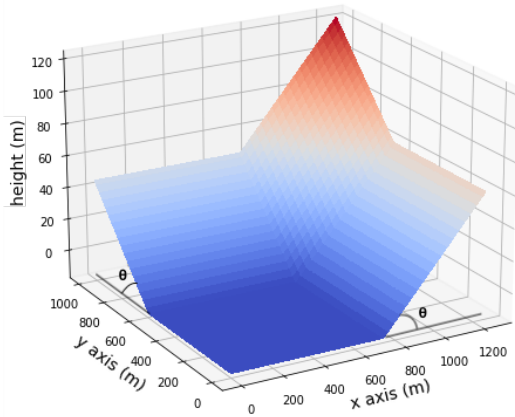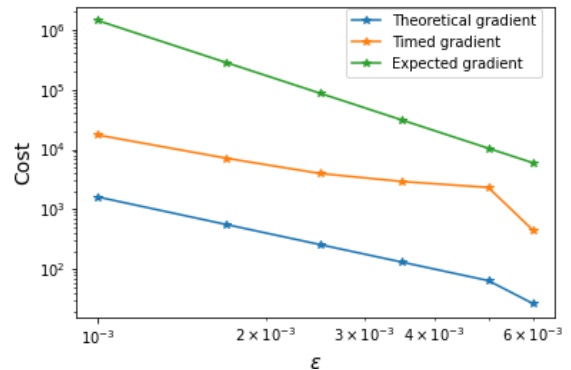
Figure 6: Schematic of 2D bed with uncertain bed slope angles θ for 2D bed slope test case.
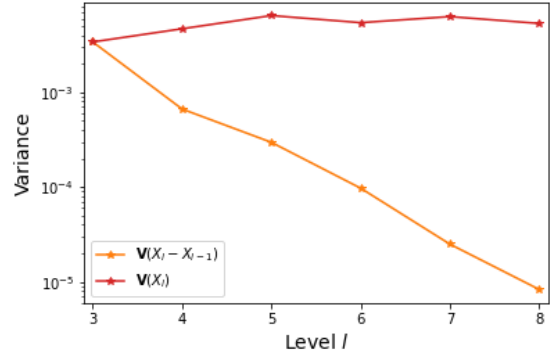
Figure 7: Variance of difference, $X_l - X_{l-1}$, compared to variance of single variable, X_l, at each level for 2D bed slope test case.

*(a) Optimum number of samples required at each l for given tolerance value, ϵ.*

(b) Verification that both theoretical cost (23) and timed cost vary with ϵ at expected rate.

Figure 8: Effect of tolerance value ϵ on cost and sample size for 2D bed slope test case.

simulation with $\Delta x = 1250/2^9 = 2.44$ m and $\Delta y = 1000/2^9 = 1.95$ m, which is the same as the finest mesh in the MLMC simulation. The maximum expected horizontal inundation distance along the x-axis over all time and all y divided by the total x-length of the domain is 0.6934 (to four significant figures). As the total domain length is 1250 m, this in fact means the water inundates a horizontal distance of 866.8 m. Figures 9a and 9b show that the RMSE of MLMC decreases both as the level number increases (*i.e.* the fineness of the mesh increases) and as ϵ decreases, showing that MLMC is working as expected.

Finally Figure 10 compares the cost and accuracy of the MLMC and Monte Carlo methods. As in Section 3.2.1, MLMC converges more uniformly than Monte Carlo and more importantly it takes the Monte Carlo method more than 222,240 core hours (8.0×10^8 s) to consistently achieve an RMSE that is lower than that achieved by MLMC with $\epsilon = 0.001$ in 160 core hours (5.7×10^5 s). (Note the time taken to achieve the best MLMC result is less than the time taken to achieve the best MLMC result in Section 3.2.1 despite the difference in dimensions because in this section we use larger ϵ values as discussed above). This is a speed up factor of approximately 1400, which is an order of magnitude larger than the speed-up factor achieved in the 1D version of this test case. We have thus again significantly
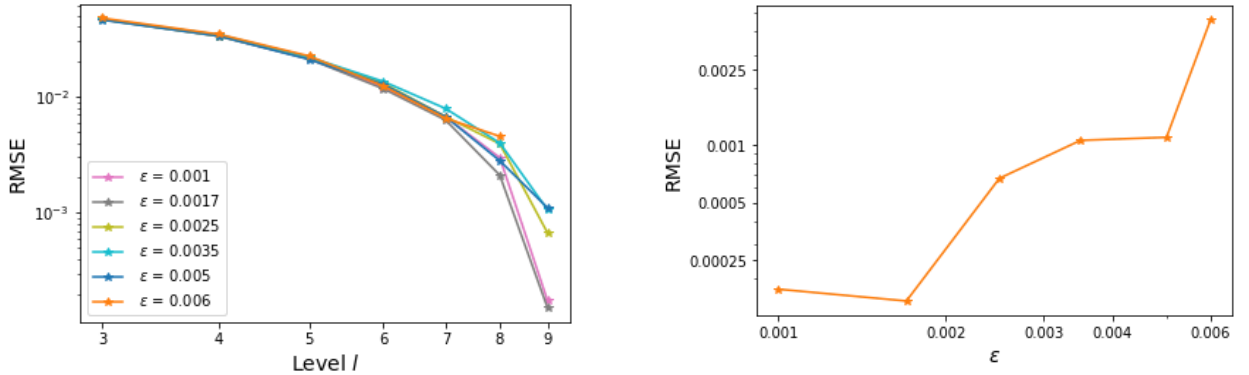
(a) RMSE of MLMC at each level l for given tolerance value ϵ.

(b) RMSE of MLMC once it has converged for each tolerance value ϵ.

Figure 9: RMSE of MLMC (24) for 2D bed slope test case for varying tolerance value ϵ.
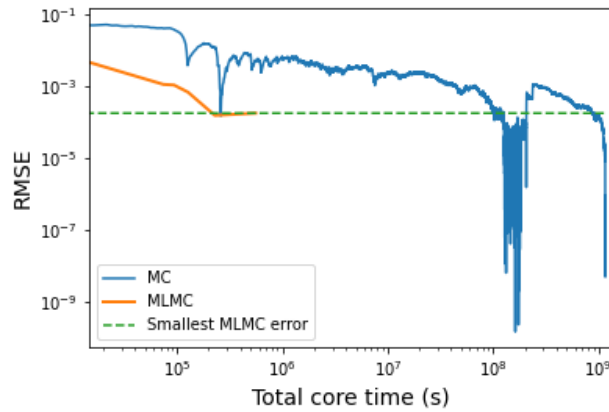

Figure 10: Comparison of total cost and accuracy when using MLMC and Monte Carlo (MC) for 2D bed slope test case.

reduced the computational time by using MLMC rather than Monte Carlo.

3.3. Uncertainty in wave height

So far in this work, we have only considered the uncertainty associated with a simple beach bed profile. Another significant source of uncertainty is the structure of the waves approaching the beach. Whilst the previous test cases are purely theoretical, in this section we consider a lab test case and a test case using real world data, with the uncertain parameter being the wave height in the JONSWAP wave spectrum, h (see Hasselmann et al. (1973) for more details on this spectrum).

A vital condition of MLMC theory is that for each X_l^i, X_{l-1}^i pair calculation, the only difference in the model must be the mesh element size. Thus, when calculating X_l^i, we use the XBeach option `random = 0` which means that the same random seed is used to initialise the wave boundary conditions for both X_{l-1}^i and X_l^i. Both test cases below pass the consistency check in Section 2.1 indicating that the formulation is both mathematically correct and implemented correctly.
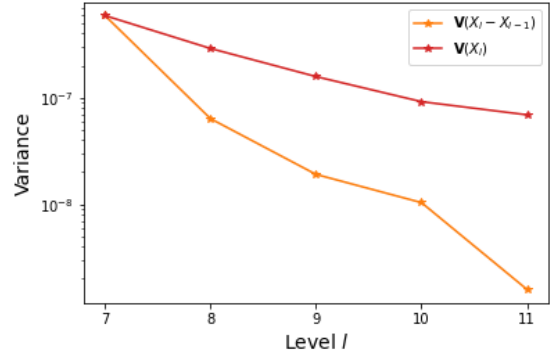
3.3.1. Morphology test case

An important variable to consider when assessing the risk to coastal areas is the volume of material eroded. To do this we consider the standard DelflandStorm test case, available as part of the XBeach documentation (Roelvink et al., 2015). Note this is a 1D test case so we only alter the mesh element size in the x-direction.

In this test case, we seek the expected total volume change over the simulation, found by using the trapezium rule
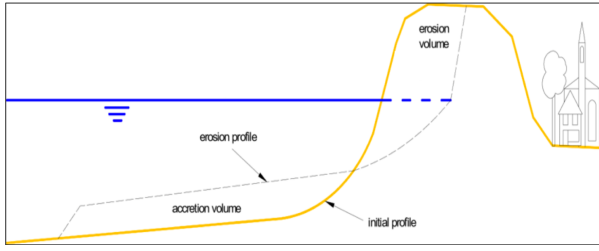
Figure 11: Schematic showing the initial and final bed profiles after a simulation. If the volume change in the morphology test case is positive, this means that the erosion volume is greater than the accretion volume and vice versa. Source: Van Gent et al. (2007).

Figure 12: Variance of difference, $X_l - X_{l-1}$, compared to variance of single variable, X_l, at each level for morphology test case.
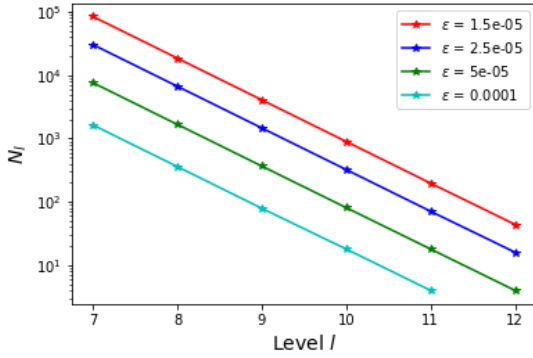
(a) Optimum number of samples required at each l for given tolerance value ϵ.
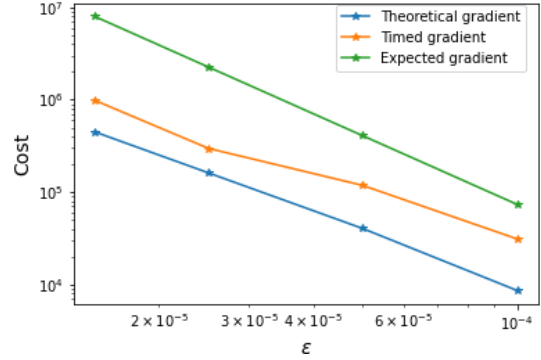

(b) Verification that both theoretical cost (23) and timed cost vary with ϵ at expected rate.

Figure 13: Effect of tolerance value ϵ on cost and sample size for the morphology test case.

to approximate the integral of the difference between the initial bed profile and the final bed profile, as shown in Figure 11. Thus if the number is positive, the erosion volume is greater than the accretion volume and the beach has lost material and been net eroded, and if the inverse is true the beach has accumulated material and grown. Note that as this is a test case in one horizontal dimension this volume is in fact an area but to avoid confusion with surface area we shall continue referring to it as a volume. We normalise the volume change using the total initial volume of the beach, *i.e.* the volume underneath the slope with the minimum value of z being the lower vertical limit. The distribution of the wave height is set to be $h \sim \mathcal{U}(0m, 5m)$.

To check whether MLMC can be used for this test case, we run steps 1 and 2 of Algorithm 1. Figure 12 shows that $\mathbb{V}(X_l - X_{l-1})$ decreases as the mesh becomes finer and that for all l it is lower than $\mathbb{V}(X_l)$. Using convergence tests, the convergence values in (9) are $\alpha = 0.921, \beta = 1.77, \gamma = 2.40$. Thus the numerical discretisation is approximately first order accurate and as $\gamma > \beta$, we expect the order of convergence to be better than the standard Monte Carlo method. In addition, the test case also passes the standard checks in Section 2.1.

Given the success of the preliminary study, we run the full algorithm and consider tolerance values of $\epsilon = [1.5 \times 10^{-5}, 2 \times 10^{-5}, 5 \times 10^{-5}, 10^{-4}]$ with a maximum limit for the finest level of $L_{\max} = 12$. As $\gamma < \beta$, the optimum N_l are calculated using (14) and are shown in Figure 13a. Importantly the number of samples required decreases as the mesh becomes finer. Furthermore, for this test case $\gamma > \beta$ and so following (11), the cost of the algorithm C should be such that $C \propto \epsilon^{-2-(\gamma-\beta)/\alpha}$. Figure 13b shows that both the timed cost and the theoretical cost (using (23)) follow
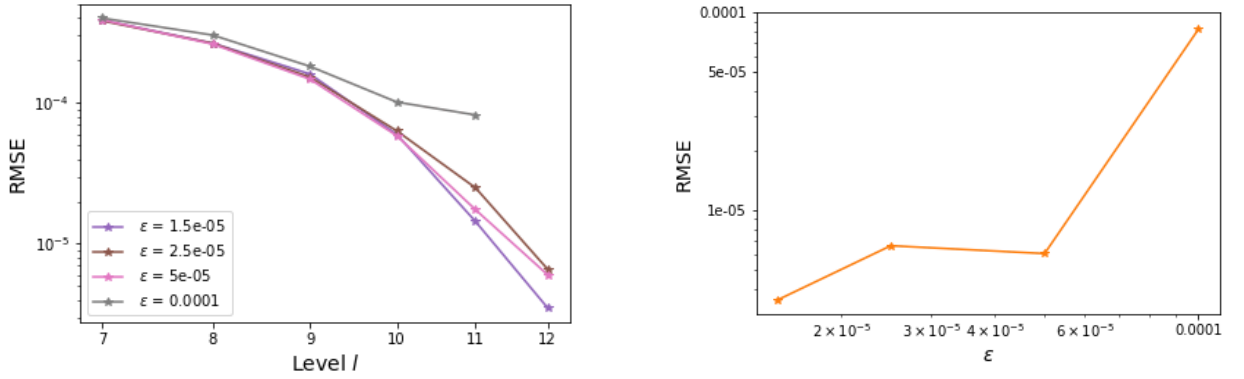
(a) RMSE of MLMC at each level l for given tolerance value ϵ.

(b) RMSE of MLMC once it has converged for each tolerance value ϵ.

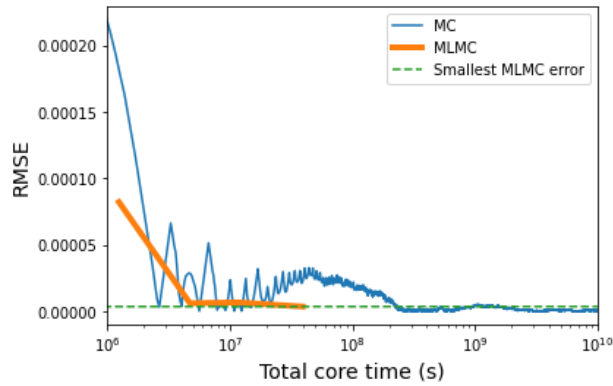**Figure 14:** RMSE of MLMC (24) for morphology test case for varying tolerance value ϵ.

Figure 15: Comparison of total cost and accuracy when using MLMC and Monte Carlo (MC) for morphology test case.

the expected gradient.

In order to calculate the RMSE of MLMC for this test case, we estimate the real solution $\mathbb{E}[X]$ through a Monte Carlo simulation using $\Delta x = 1421.562/2^{12} = 0.347$ m, which is the same as the finest mesh utilised in the MLMC simulation. The expected volume change as a proportion of the total initial volume is 0.0006150 (to four significant figures). This is positive and thus means that the beach has been net eroded during the simulation. As the total volume is 16 333 m^2, this in fact means that 10 m^2 of beach material is eroded during the simulation (recall here that the test case is in one horizontal dimension so the volume is in fact an area). Figures 14a and 14b confirm that the RMSE decreases as the level number increases and as ϵ decreases respectively, thus confirming MLMC is working as expected. Unlike with the other test cases considered, Figure 14b shows that the relationship between the RMSE and ϵ is not uniform and there is a clear elbow in the plot at $\epsilon = 5 \times 10^{-5}$. This is a discrete effect caused by the fact that for $\epsilon = 10^{-4}$, the convergence test (16) passes without requiring any samples on the finest level ($l = 12$) (see Figure 13a).

Finally Figure 15 compares the cost and accuracy for MLMC versus Monte Carlo. It shows that in this test case the convergence of the Monte Carlo simulation is notably less uniform than that from MLMC, meaning the Monte Carlo simulation must be run for a long time to ensure convergence. (Note that the elbow present in Figure 14b is again present in the MLMC convergence for the same reason). Importantly, Figure 15 also shows that the Monte Carlo method requires over 400,000 core hours (1.4×10^9 s) to consistently achieve an RMSE lower than that achieved by MLMC with $\epsilon = 1.5 \times 10^{-5}$ in 10,911 core hours (3.9×10^7 s). Thus MLMC results in a speed-up factor of approximately 40 and we have shown that MLMC can significantly decrease computational costs without compromising on accuracy for a full hydro-morphodynamic model with uncertainties in the wave height.
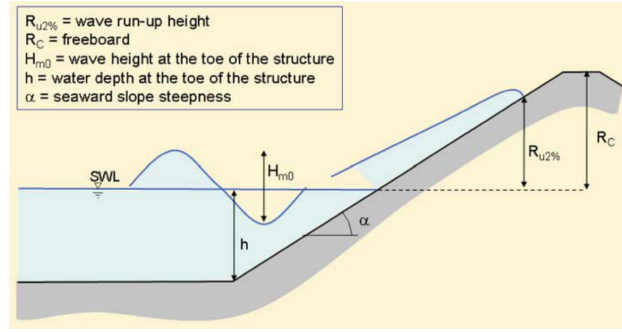
Figure 16: Schematic showing $R_{u2\%}$ wave run-up height at a beach (the quantity of interest in the Boscombe Beach test case) and key quantities needed to calculate it. Source: Allsop et al. (2007).

3.3.2. Boscombe Beach test case

For the final test case, we consider the real world case of Boscombe Beach in Dorset, UK, which is a standard 2D test case in the XBeach documentation (Roelvink et al., 2015). The distribution of the wave height is set to be $h \sim \mathcal{U}(0m, 3m)$. To make the problem simpler, the morphological and tidal component present in the standard test case are switched off, but all remaining values are kept the same as those in the standard test case.

When assessing flood risk, an important quantity not yet considered in this work is the wave run-up height. Traditionally, many coastal structures have been designed to withstand the wave run-up height exceeded by 2% of incoming waves, $R_{u2\%}$ (shown in Figure 16), with the idea being that soft engineering measures such as clay and grass are sufficient to withstand the 2% of waves that exceed this height (see Allsop et al. (2007)). Thus in this section, we seek the expected value of the maximum run-up height $R_{u2\%}$ attained during the simulation over all values of y. Note that $R_{u2\%}$ is by definition a probabilistic measure, but there exist a variety of standard formulae that can be used to estimate this value deterministically for an individual simulation (see Melby et al. (2012)). We use the industry standard EurOtop formula given in Allsop et al. (2007) as

$$R_{u2\%} = 1.65\gamma_b\gamma_f\gamma_\beta H_{m0}\xi_{m-1,0}, \tag{25}$$

where γ_b, γ_f and γ_β are the influence factors for the berm (raised bank), slope roughness and oblique wave respectively, H_{m0} is the wave height at the beach-toe (*i.e.* where the water level (z_s) and the bed (z_b) are equal to each other) and $\xi_{m-1,0}$ the breaker parameter. Note as the specific values of γ_b, γ_f and γ_β for Boscombe Beach are unavailable, they are assumed equal to 1 but given the linear relationship between these values and $R_{u2\%}$ it is simple to change these values after the simulations have finished. The breaker parameter relates the slope steepness $\tan\phi$ to the wave steepness $s_{m-1,0}$ by

$$\xi_{m-1,0} = \frac{\tan\phi}{\sqrt{s_{m-1,0}}}, \tag{26}$$

where ϕ is the beach slope angle at the beach-toe and $s_{m-1,0}$ is the ratio of the wave height to the wavelength (H_{m0}/L_0). The wavelength is calculated as

$$L_0 = \frac{gT_{m-1,0}^2}{2\pi}, \tag{27}$$

and we use the standard approximation (see Melby et al. (2012)) that $T_{m-1,0} = T_p/1.1$ where T_p is the peak wave period.

Convolution filter

To set the bed in the XBeach model, we use the bed data provided in the XBeach documentation (Roelvink et al., 2015). Clearly in real world environments, the beach approach is not smooth as has been the case in the test cases
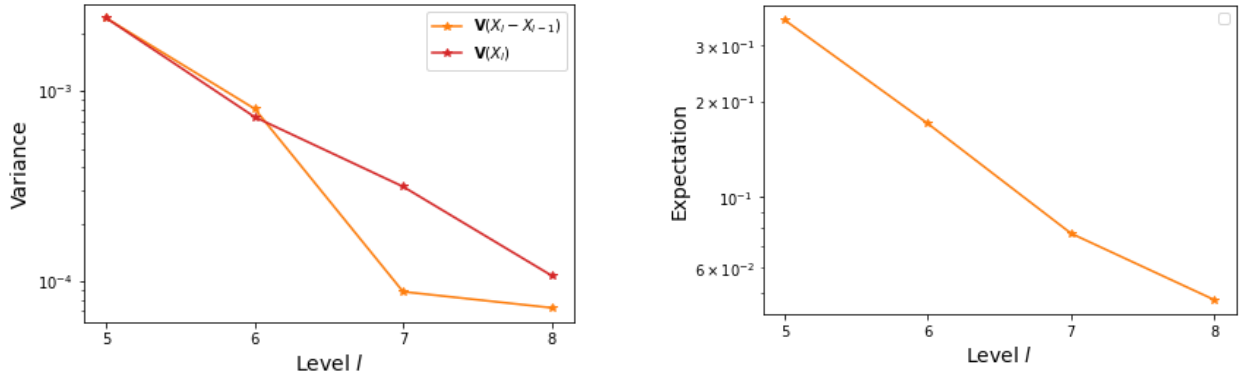
(a) Variance of difference, $X_l - X_{l-1}$, compared to variance of single variable, X_l, at each level, l.

(b) Expectation of difference, $X_l - X_{l-1}$, at each level, l.

Figure 17: Preliminary MLMC results for Boscombe Beach test case.

considered so far in this work. Instead there may be sandbars or other similar physical features present, either natural or anthropogenic. This could potentially cause issues especially if a feature is not represented in coarser meshes and suddenly appears in finer meshes. However, with MLMC it is not necessary that the coarsest mesh is a good approximation to the finest mesh, merely that each mesh l is a good enough approximation to the next finest mesh $l + 1$.

In order to mitigate the issue of features suddenly appearing, we use a convolution matrix filter which has the general expression

$$(f * g)(i) = \sum_{k=-\lambda}^{\lambda} g(k) f(i - k), \tag{28}$$

where $\lambda = \lfloor \frac{m}{2} \rfloor$ and m is the length of g. An appropriate filter kernel to use in this case is the normalised arithmetic mean

$$g(j) = \frac{1}{\sum_{i=1}^{m} g(i)}. \tag{29}$$

At the edges, we use the 'nearest' method meaning that the edge value is repeated outwards as many times as necessary for the multiplication of the weights.

For this test case, f is the underlying bed and the convolution filter is applied in the direction perpendicular to the coastline. The convolution filter length, m_l at each level l follows $m_l = \max((L_{max} - l), 1)$ where L_{max} is the finest mesh level considered. At level L_{max}, $m = 1$ which is the identity filter and thus the bed is unchanged. Hence, as the mesh becomes finer, the convolved bed slowly converges to the original bed and thus there are no extra errors associated with using a convolved bed.

To check whether it is valid to use MLMC for this test case, we run steps 1 and 2 of Algorithm 1. Figure 17a shows that $\mathbb{V}\left(X_l - X_{l-1}\right)$ decreases as the mesh gets finer. At $l = 5$, $\mathbb{V}\left(X_l - X_{l-1}\right)$ is approximately equal to $\mathbb{V}\left(X_l\right)$, probably due to the fact that, even with the convolution, the mesh at $l = 5$ does not give the best approximation to $l = 6$ thus creating a relatively large variance. To further check the effects of using a convolution, Figure 17b shows how $\mathbb{E}\left(X_l - X_{l-1}\right)$ varies; unlike with the variance, the decrease in the quantity is smooth and shows that the numerical discretisation is approximately first order accurate, as with the test case in Section 3.3.1. In addition, the simulation also passes the consistency and kurtosis checks from Section 2.1 and thus we can be confident that the MLMC algorithm is appropriate. Using convergence tests, the convergence values in (9) are $\alpha = 0.996$, $\beta = 1.84$ and $\gamma = 2.52$. As $\gamma > \beta$, we expect the order of convergence to be better than the standard Monte Carlo method.

Proceeding, we consider tolerance values of $\epsilon = [0.002, 0.004, 0.006, 0.008, 0.01]$ and set a maximum limit of the finest mesh of $L_{max} = 8$ because this corresponds to the mesh that the Boscombe Beach bed data is defined on. As $\gamma > \beta$, the optimum N_l are calculated using (14) and are shown in Figure 18a. Importantly the number of samples decreases as the mesh becomes finer.
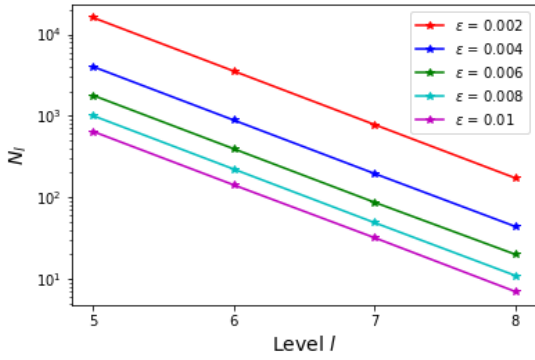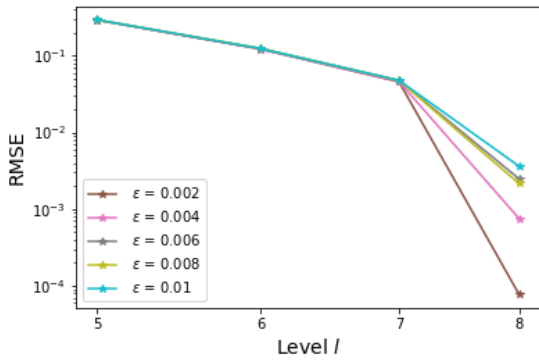
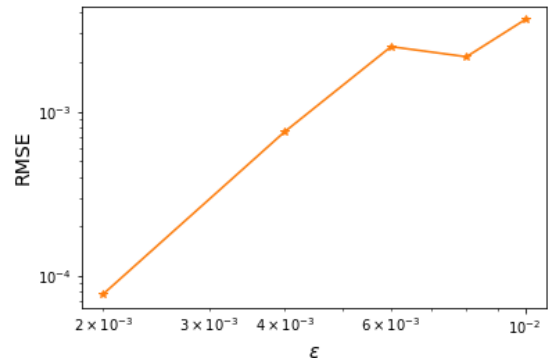*(a) Optimum number of samples required at each l for given tolerance value ϵ.*
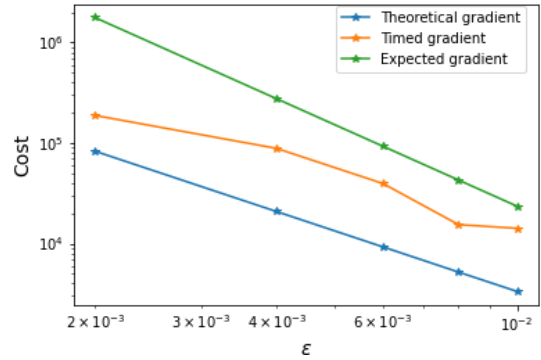
(b) Verification that both theoretical cost (23) and timed cost vary with ϵ at expected rate.

Figure 18: Effect of tolerance value ϵ on cost and sample size for Boscombe Beach test case.

(a) RMSE of MLMC at each level l for given tolerance value ϵ.

*(b) RMSE of MLMC once it has converged for each tolerance value ϵ.*

Figure 19: RMSE of MLMC (24) for Boscombe Beach test case for varying tolerance value ϵ.

For this test case $\gamma > \beta$ and so following (11), the cost of the algorithm C should be such that $C \propto \epsilon^{-2-(\gamma-\beta)/\alpha}$. Figure 18b shows that this is roughly true for both the timed cost and the theoretical cost (calculated using (23)). Any discrepancy is likely due to the fact that at coarser grids the convolution simplifies the bed meaning the XBeach model runs more quickly than it would if it were using the same coarse grid on the more complex original bed. This effect is a more significant proportion of the time at larger ϵ where there are very few samples at finer levels.

As before we calculate the RMSE of the MLMC simulation by estimating the real solution $\mathbb{E}[X]$ through an Monte Carlo simulation with $\Delta x = 1123.8461/2^8 = 4.39\,\mathrm{m}$ and $\Delta y = 1604/2^8 = 6.27\,\mathrm{m}$ which is the same as the finest mesh in the MLMC simulation. The expected maximum value of $R_{u2\%}$ is found to be $0.067\,64\,\mathrm{m}$ (to four significant figures). Using this value, Figures 19a and 19b confirm that the RMSE decreases as the level number increases and as ϵ decreases respectively, meaning the MLMC algorithm is working as expected. Interestingly, Figure 19a shows that ϵ only has an effect on the final error when the finest level is included. This makes sense because at previous levels, the convolved version of the bed is being used rather than the real data.

Finally Figure 20 compares the cost and accuracy of using MLMC versus Monte Carlo for this test case. Unlike with previous test cases, the figure shows that for large values of ϵ, it may be quicker to use Monte Carlo instead of MLMC. This may be because due to the use of convolutions, a certain number of samples are required on the finest levels to obtain an accurate result. We also note that, as has been the case for all the test cases considered in this work, with Monte Carlo the simulation must be run for longer than strictly necessary to ensure convergence, which is not the
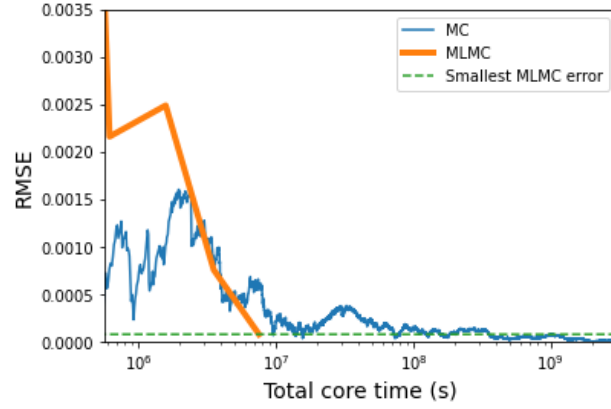
Figure 20: Comparison of total cost and accuracy when using MLMC and Monte Carlo (MC) for Boscombe Beach test case.

case for MLMC. The choice of Monte Carlo versus MLMC for this test case therefore depends on how accurate and how certain the required result needs to be. Furthermore, for the smallest ϵ and longest simulation considered here, the error using MLMC is much smaller than the error using Monte Carlo for the same computational cost. In order to achieve the same accuracy, the Monte Carlo simulation requires over 100,000 core hours (3.6×10^8 s) compared to 2092 core hours (7.5×10^6 s) for MLMC with $\epsilon = 0.002$. Thus using MLMC results in a speed-up factor of 50 showing that MLMC can significantly decrease computational cost for a real world test case with uncertain wave height.

4. Cumulative distribution functions

In an MLMC framework, the objective is normally to find the expectation of an output variable. However, to analyse risk in coastal problems, the quantity of interest is the probability of an output variable exceeding a certain value. As discussed in Section 1, this is a complicated output to compute because MLMC provides very few values on the same level from which to build that distribution.

To resolve this issue, we follow Gregory and Cotter (2017b) and use the inverse transform sampling method to evaluate the inverse cumulative distribution function (CDF), $F^{-1}(u)$, where $u \sim \mathcal{U}[0, 1]$. If F is strictly increasing and absolutely continuous then $x \equiv F^{-1}(u)$ is unique. A simple consistent estimate for x is found by sorting the samples $X^i_{i=1,...,N} \sim F$ such that $X^1 < X^2 < ... < X^N$ and then

$$\hat{F}^{-1}(u) = X^{\lceil N \times u \rceil}, \tag{30}$$

which, as argued in Gregory and Cotter (2017b), is consistent because it converges in probability to x as $N \to \infty$. For an MLMC approximation, the inverse CDF is then given by

$$F_L^{-1}(u) = R(X)_0^{\lceil N_0 \times u \rceil} + \sum_{l=1}^{L} \left(R(X)_l^{\lceil N_l \times u \rceil} - R * (X)_{l-1}^{\lceil N_l \times u \rceil} \right), \tag{31}$$

where $R(X)_l^i$ represents the i^{th} ordered statistic of X_l on each level l and $R * (X)_l^i$ represents the same but from a different run of the algorithm than $R(X)_l^i$. Note this means the MLMC algorithm must be run twice but given MLMC is significantly faster than Monte Carlo, even running it twice is much faster than running Monte Carlo once. Unlike with (2) there is no exact cancellation here as the approximations on each level are not unbiased.
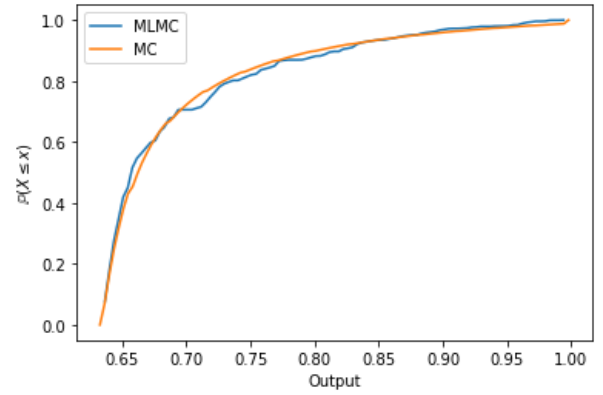
Thus an ensemble of X_F^i can be generated, which are not samples from X_L, but consistent approximations to $F_L^{-1}(u)$. Thus, as $N_l \to \infty$, then $x \xrightarrow{p} F_L^{-1}(u)$ and the CDF of the MLMC approximation is given by

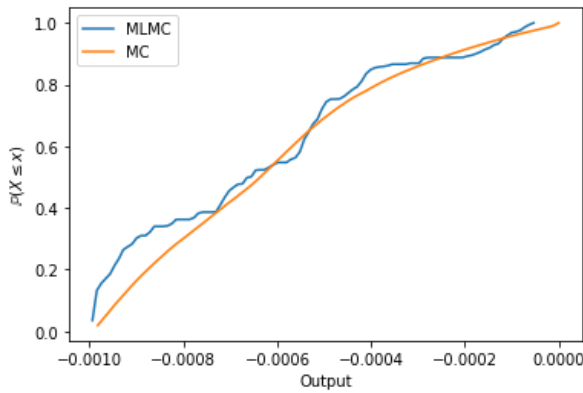$$F(x) = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}_{X_F^i \leq x}. \tag{32}$$

For more details on this method we refer the reader to Gregory and Cotter (2017b).
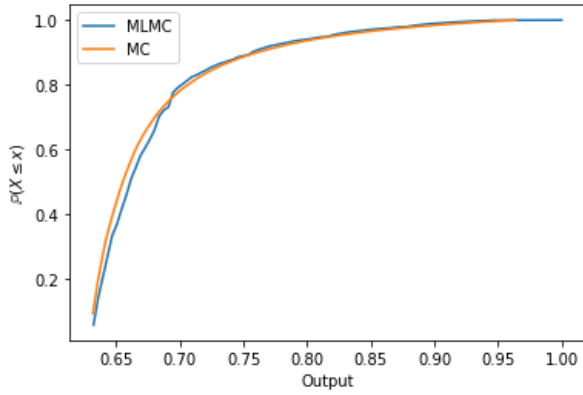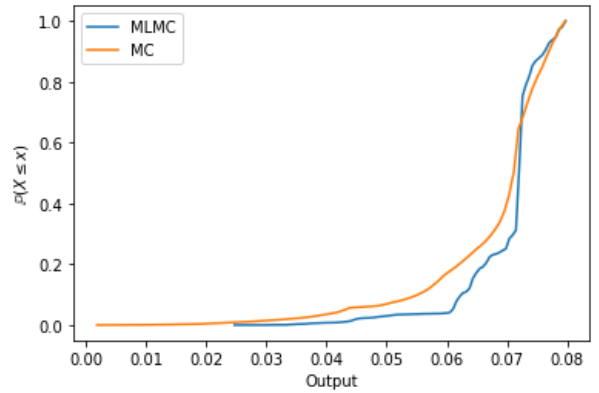
(a) 1D bed slope test case. Section 3.2.1

(b) 2D bed slope test case. Section 3.2.2

(c) Morphology test case. Section 3.3.1

(d) Boscombe Beach test case. Section 3.3.2

Figure 21: Cumulative distribution functions generated using inverse transform sampling method.

4.1. Applying inverse transform sampling to test cases

In this section, we apply the inverse transform sampling method to each of the four test cases discussed in Section 3. This method provides a CDF from which the probability that the output variable exceeds a certain value can be predicted. This is particularly useful for predicting the probability of extreme flooding/erosion events occurring, a useful tool in coastal risk assessment. To apply the method, we select the lowest ϵ value considered in each test case and run Algorithm 1 twice. The output is recorded at each level separately, ordered and then (31) is used to generate pseudo-samples on the finest level.

The results of applying this method to all four test cases are shown in Figure 21, which also compares the MLMC results with those obtained by the Monte Carlo method. In all cases the MLMC results predict a similar CDF to the Monte Carlo method. The good approximation between MLMC and Monte Carlo is further shown by the small error norms between the two methods, shown in Table 1. Note the error for the Boscombe Beach test case whilst still small is a little higher than for the other test cases. This is in part due to the fact that MLMC does not predict any $R_{u2\%}$ lower than 0.025 m – if values below this quantity are excluded from the error comparison, the normalised L2 error norm drops by 25% to 0.0334. Whilst in an ideal situation MLMC would predict these lower values too, the lower values of the $R_{u2\$}$ distribution are of the least interest because they pose the lowest risk to the beach and thus the fact that MLMC does not predict any values in this range is arguably not a significant issue. Thus we conclude that we can use MLMC not only to find the expected value of output variables but also to determine the probability of extreme flooding and erosion events occurring and assess whether certain areas are at risk.

Table 1
Error norms between MLMC and Monte Carlo cumulative distribution functions for each test case.

Test Case	Normalised L2 Error Norm (3sf)
1D bed slope	0.000612
2D bed slope	0.00127
Morphology	0.00519
Boscombe Beach	0.0442

5. Conclusion

In this work, we have applied MLMC to a complex hydro-morphodynamic model for the first time to evaluate risk in the coastal zone. We have shown that MLMC can be used successfully to estimate both the expected value of a hydro-morphodynamic output variable and its distribution for both 1D and 2D test cases and for both experimental and real-world test cases. The key advantage of MLMC over more traditional methods such as the Monte Carlo method is that we have shown that it is considerably faster whilst maintaining the same level of accuracy in the results. Another significant advantage of MLMC over Monte Carlo is that the solution is stable – in all test cases considered, decreasing ϵ (*i.e.* the accuracy tolerance) results in a more accurate solution. Thus there is no need to run the simulation for longer to ensure the solution has converged, as is the case with Monte Carlo. This again provides a computational cost advantage for MLMC compared to Monte Carlo. Furthermore, throughout this work, MLMC has no difficulties in dealing with the complex nature of the XBeach model nor with the fact that some parametrisations in XBeach only appear at finer meshes. This is mainly due to the fact that MLMC only requires that the result from one mesh is a good approximation to the result from the next finest mesh, which makes it very flexible.

In Section 1, we discussed the case in Callaghan et al. (2013), which would take approximately 4.5 millennia to run with Monte Carlo. This test case is similar to both the morphology and Boscombe Beach test case considered in this work, for which using MLMC resulted in a speed up of 40 and 50 times respectively. Therefore it is not unreasonable to suggest that using MLMC, the test case in Callaghan et al. (2013) would only take 100,000 years. Whilst this is still a large cost, because the code is parallelised, this is not outside the realm of certain supercomputers.

This work has shown the potential for MLMC to be used in the coastal engineering field. In future work we will consider more complex examples and investigate the effect of uncertainty in more than one parameter on more than one output variable.

Computer code availability

The relevant code for the MLMC framework presented in this work can be found at `https://github.com/mc4117/MLMC_XBeach`.

Acknowledgements

References

Allsop, N.W.H., Pullen, T., Bruce, T., van der Meer, J.W., Schüttrumpf, H., Kortenhaus, A., 2007. EurOtop Wave Overtopping of Sea Defences and Related Structures: Assessment Manual. Boyens Medien GmbH & Co. KG.

Andrews, D.G., McIntyre, M.E., 1978. An exact theory of nonlinear waves on a lagrangian-mean flow. Journal of fluid Mechanics 89, 609–646.

Apel, H., Thieken, A.H., Merz, B., Blöschl, G., 2004. Natural Hazards and Earth System Sciences Flood risk assessment and associated uncertainty. Natural Hazards and Earth System Sciences 4, 295–308.

Bart, L.J.C., 2017. Long-term modelling with XBeach : combining stationary and surfbeat mode in an integrated approach. Masters thesis. TU Delft, Delft, The Netherlands.

Caflisch, R.E., 1998. Monte Carlo and quasi-Monte Carlo methods. Acta Numerica 1998, 1–49.

Callaghan, D.P., Ranasinghe, R., Roelvink, D., 2013. Probabilistic estimation of storm erosion using analytical, semi-empirical, and process based storm erosion models. Coastal Engineering 82, 64–75.

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

Elfverson, D., Hellman, F., Målqvist, A., 2016. A Multilevel Monte Carlo Method for Computing Failure Probabilities. SIAM/ASA Journal on Uncertainty Quantification 4, 312–330.

Giles, M.B., 2008. Multilevel Monte Carlo Path Simulation. Operations Research 56, 607–617. doi:10.1287/opre.1070.0496.

Giles, M.B., 2015. Multilevel Monte Carlo methods. Acta Numerica 24, 259–328.

Giles, M.B., Nagapetyan, T., Ritter, K., 2015. Multilevel monte carlo approximation of distribution functions and densities. SIAM/ASA journal on Uncertainty Quantification 3, 267–295. doi:10.1137/140960086.

Gregory, A., Cotter, C.J., 2017a. A seamless multilevel ensemble transform particle filter. SIAM Journal on Scientific Computing 39, A2684–A2701.

Gregory, A., Cotter, C.J., 2017b. On the calibration of multilevel Monte Carlo ensemble forecasts. Quarterly Journal of the Royal Meteorological Society 143, 1929–1935. doi:10.1002/qj.3052.

Harris, D.L., Rovere, A., Casella, E., Power, H., Canavesio, R., Collin, A., Pomeroy, A., Webster, J.M., Parravicini, V., 2018. Coral reef structural complexity provides important coastal protection from waves under rising sea levels. Science Advances 4.

Hasselmann, K., Barnett, T., Bouws, E., Carlson, H., Cartwright, D., Enke, K., Ewing, J., Gienapp, H., Hasselmann, D., Kruseman, P., et al., 1973. Measurements of wind-wave growth and swell decay during the Joint North Sea Wave Project (JONSWAP). Ergänzungsheft 8-12 .

Katsiolides, G., Müller, E.H., Scheichl, R., Shardlow, T., Giles, M.B., Thomson, D.J., 2018. Multilevel Monte Carlo and improved timestepping methods in atmospheric dispersion modelling. Journal of Computational Physics 354, 320–343.

Melby, J., Caraballo-Nadal, N., Kobayashi, N., 2012. Wave runup prediction for flood mapping. Coastal Engineering Proceedings 1, management.79.

Pender, D., Karunarathna, H., 2013. A statistical-process based approach for modelling beach profile variability. Coastal Engineering 81, 19–29.

Roelvink, D., Van Dongeren, A., McCall, R., Hoonhout, B., Van Rooijen, A., Van Geer, P., De Vet, L., Nederhoff, K., Quataert, E., 2015. XBeach Technical Reference: Kingsday Release. Delft, The Netherlands: Deltares, Technical report .

Sánchez-Linares, C., de la Asunción, M., Castro, M.J., González-Vida, J.M., Macías, J., Mishra, S., 2016. Uncertainty quantification in tsunami modeling using multi-level Monte Carlo finite volume method. Journal of Mathematics in Industry 6, 1–26.

Simmons, J.A., Harley, M.D., Marshall, L.A., Turner, I.L., Splinter, K.D., Cox, R.J., 2017. Calibrating and assessing uncertainty in coastal numerical models. Coastal Engineering 125, 28–41.

Unguendoli, S., 2018. Propagation of uncertainty across modeling chains to evaluate hydraulic vulnerability in coastal areas. Ph.D. thesis. Università di Bologna, Bologna, Italy.

Van Gent, M.R., Coeveld, E.M., De Vroeg, H., Van De Graaff, J., 2007. Dune erosion prediction methods incorporating effects of wave periods, in: Coastal Sediments '07 - Proceedings of 6th International Symposium on Coastal Engineering and Science of Coastal Sediment Processes, pp. 612–625.

Villaret, C., Kopmann, R., Wyncoll, D., Riehme, J., Merkel, U., Naumann, U., 2016. First-order uncertainty analysis using Algorithmic Differentiation of morphodynamic models. Computers & geosciences 90, 144–151.

Vousdoukas, M.I., Mentaschi, L., Voukouvalas, E., Bianchi, A., Dottori, F., Feyen, L., 2018. Climatic and socioeconomic controls of future coastal flood risk in europe. Nature Climate Change 8, 776–780.