Probabilistic Engineering Mechanics 63 (2021) 103116

Contents lists available at ScienceDirect

Probabilistic Engineering Mechanics

journal homepage: www.elsevier.com/locate/probengmech

Multilevel Monte Carlo simulations of composite structures with uncertain manufacturing defects

T.J. Dodwell [a,b,*], S. Kynaston [c], R. Butler [d], R.T. Haftka [e], Nam H. Kim [e], R. Scheichl [f,c]

[a] College of Engineering, Mathematics and Physical Sciences, Exeter EX4 4PY, UK
[b] The Alan Turing Institute, London NW1 2DB, UK
[c] Department of Mathematical Sciences, University of Bath, Bath BA2 7AY, UK
[d] Department of Mechanical Engineering, University of Bath, Bath BA2 7AY, UK
[e] Department of Mechanical and Aerospace Engineering, University of Florida, FL, USA
[f] Institute of Applied Mathematics & Interdisciplinary Center for Scientific Computing, University of Heidelberg, 69120 Heidelberg, Germany

ARTICLE INFO

The authors dedicate this paper to Prof. Raphael T. Haftka, a great collaborator, colleague, mentor, and friend

ABSTRACT

By adopting a Multilevel Monte Carlo (MLMC) framework, this paper shows that only a handful of costly fine scale computations are needed to accurately estimate statistics of the failure of a composite structure, as opposed to the many thousands typically needed in classical Monte Carlo analyses. The paper introduces the MLMC method and provides an extension called MLMC with selective refinement to efficiently calculated structural failure probabilities. Simple-to-implement, self-adaptive algorithms are given, and the results demonstrate huge computational gains for two novel, real world example problems in composites performance analysis: (i) the effects of fibre waviness on the compressive strength of a composite material and (ii) the uncertain buckling performance of a composite panel with random ply orientations. For the most challenging test case of estimating a 1/150 probability of buckling failure of a composite panel the results demonstrate a speed-up factor of > 1000 over classical Monte Carlo. In absolute terms, the computational time was reduced from 218 CPU days to just 4.4 CPU hours, making stochastic simulations that would otherwise be unthinkable now possible.

1. Introduction

Within the aerospace manufacturing sector, where safety is paramount, risk is quantified and reduced by heuristic safety factors and expensive programmes of empirical testing over a variety of length scales before new designs can enter production. The high cost of certification and the inefficiency of general safety factors has led to new initiatives admitting numerical simulation and stochastic methods for certification [1], as well as to an increasing interest in probabilistic design [2]. Both of which provided opportunities to demonstrate structural integrity even when experimental or statistical data is incomplete, offering scope to challenge conservative failure limits and reduce design-to-manufacture time.

In complex composite manufacturing processes, uncertainty arises from a number of different sources, for example material variability [3], machine tolerance [4] and process-induced defects [5–8]. Therefore, structural performance of a real composite component is uncertain. This uncertainty can be explored using stochastic simulation methods, where the effect of introducing a distribution of uncertain material properties and/or structural configurations can pass through a model to predict the distribution of part performance. However, such a statistical analysis typically requires a large number of complex model evaluations, and thus quickly becomes prohibitively computationally expensive. Various techniques have been developed to reduce the computational cost of such stochastic studies.

The traditional engineering approaches of First and Second Order Reliability Methods (FORM/SORM) [9, Sect. 4.4] can been applied to substantially reduce the number of simulations required in stochastic studies. Such methods have been applied to buckling of shells with random imperfections [10] and to composite laminates [11]. However, they do not capture the statistical correlation of multiple failure modes; and their accuracy degenerates significantly when random variables have complex non-Gaussian and/or multi-modal distributions [12]. In such cases, Monte Carlo simulation is often the only choice to guarantee unbiased estimates [13], and therefore numerous methods have been developed to reduce the cost of basic Monte Carlo sampling and averaging. Importance sampling, e.g. [9, Sect. 3.4], reduces the number of required simulations by preferentially focussing computations near the boundary of the failure domain. However, to identify this boundary using a prudent choice of the importance sampling density is often not

* Corresponding author at: College of Engineering, Mathematics and Physical Sciences, Exeter EX4 4PY, UK.
E-mail address: t.dodwell@exeter.ac.uk (T.J. Dodwell).

https://doi.org/10.1016/j.probengmech.2020.103116
Received 21 June 2020; Received in revised form 19 August 2020; Accepted 14 December 2020
Available online 17 December 2020

possible for complex engineering problems, or at least can be as difficult as the original Monte Carlo simulation itself. Similarly, separable Monte Carlo methods [14] take advantage of the independence of different sources of uncertainty, and the two approaches can be combined for additional savings [15]. Sampling points can be chosen in a structured, deterministic or pseudo-random way, leading to Quasi-Monte Carlo Methods [16,17]. Finally, surrogates [18] are often used, e.g. Gaussian Process Regression [19] or Polynomial Chaos [20]. However, those methods suffer (i) when the dimension of the uncertain parameter space becomes even moderately large, (ii) when trying to predict uncertainty in the tails of a distribution.

This paper chooses an alternative approach that sets out to optimise the use of a hierarchy of finite element (FE) models for Monte Carlo simulations of composites with defects. By adopting a Multilevel Monte Carlo (MLMC) framework, it is shown that only a handful of costly fine scale computations are needed to accurately estimate statistics of structural failure loads, as opposed to thousands of fine scale samples typically needed in classical Monte Carlo analyses. The missing exploration of the variability, leading to sampling error, is taken care of by a large number of coarse simulations. Multilevel techniques [21], were first developed by Heinrich [22], and later popularised by Giles [23] in the context of option pricing in financial mathematics. Its huge potential in uncertainty quantification for engineering applications was identified by Cliffe et al. [24] where it has been motivated via a subsurface hydrology application. Since then it has been applied to a range of other applications [25–27], it has been improved [28,29] and extended to allow also for experimental data to be taken into account in a Bayesian setting [30,31].

Importantly in many engineering applications, estimating the expected load of structural failure is of limited interest, instead in design, often the probability that the failure load is less than a 'safe' value is required. Such a model has a binary output, failure (1) or no failure (0). Thus in this paper, an extension to MLMC originally proposed by Elfverson et al. [29] is also pursued and refined. For most samples, using a FE error estimator it is possible to conclude from a coarse, computationally cheap model that further model refinements will not change the binary output. It suffices that the coarse model predicts a load sufficiently far from the failure boundary. This extension of MLMC is referred to as *Multilevel Monte Carlo with Selective Refinement* (MLMC-SR). It delivers significant further computational gains over MLMC, as will be demonstrated.

The aim of this paper is to describe the multilevel Monte Carlo method in an abstract way to show that it can be applied to a broad class of problems (Section 2), before making the following novel contributions in the context of the composite applications:

1. to provide an extension of the MLMC framework that significantly accelerates the estimation of failure probabilities using MLMC-SR (Section 3);
2. to present simple-to-implement, self adaptive, practical algorithms for both MLMC (Algorithm 1) and MLMC-SR (Algorithm 2);
3. to demonstrate the MLMC framework and its extension for computing failure probabilities on two novel stochastic analyses of classical problems in aerospace composites:

   Example I (in Section 4), which explores the effects of random fibre waviness on the compressive strength of composite material, and makes new connections with Budiansky's classical (deterministic) kinking model [32].

   Example II (in Section 5), which studies the buckling performance of a skin panel with uncertain ply orientations — as far as the authors are aware, this is the first study using MLMC for a stochastic eigenvalue problem (see also [33] for our original results).

The numerical results confirm the theoretically predicted gains for both example problems with huge – up to 1000-fold – potential speed-ups versus standard Monte Carlo simulation. This level of speed-up brings stochastic simulations that would otherwise be unthinkable into the feasible range.

2. The multi-level Monte Carlo methodology

In this section the Multilevel Monte Carlo (MLMC) method, originally proposed by Giles [23], is introduced and compared to classical Monte Carlo simulation. The section is finished by providing a simple to implement, self-adaptive algorithm.

To describe the multilevel uncertainty quantification method, it is assumed that there is a finite element (FE) model of a composite structure that is subject to some uncertainty in its material properties, for example due to a defect or the misalignment of fibres. The accuracy and the computational cost of the model is directly linked to the number of degrees of freedom (M) and thus to the resolution of the FE mesh. Typically, for a particular application an engineer is interested in some scalar quantity of interest $Q := \mathcal{G}(\mathbf{u})$, where $\mathcal{G}(\cdot)$ is a (nonlinear) function of the FE solution $\mathbf{u}$. In the context of the two example problems, this is for example the expected value of the failure stress or the critical buckling load. In cases with random defects or uncertainty, engineers are therefore interested in estimating the expected value of Q, denoted $\mathbb{E}[Q]$, or perhaps the distribution of Q.

2.1. Standard Monte-Carlo

In a typical Monte Carlo (MC) analysis, a large number (N) of independent random realisations (or samples) of the underlying random parameters are created. For each sample, the FE solution is computed on a mesh with M degrees of freedom. From this solution, the quantity of interest $Q_M^{(j)}$ for the jth sample is evaluated. Then, the average of these independent samples,

$$\hat{Q}_{M,N}^{\mathrm{MC}} = \frac{1}{N}\sum_{j=1}^{N} Q_M^{(j)}, \tag{1}$$

is the standard Monte Carlo estimator for the expected value $\mathbb{E}[Q_M]$ of the random variable Q_M.

The error in the approximation $\hat{Q}_{M,N}^{\mathrm{MC}}$ can be quantified using the *root mean square error* (RMSE), given by

$$e(\hat{Q}_{M,N}^{\mathrm{MC}}) = \left(\mathbb{E}[(\hat{Q}_{M,N}^{\mathrm{MC}} - \mathbb{E}[Q])^2]\right)^{1/2}. \tag{2}$$

A particular advantage of quantifying the error in this way is the fact that the mean square error can be expanded so that

$$e(\hat{Q}_{M,N}^{\mathrm{MC}})^2 = \mathbb{E}[Q_M - Q]^2 + \frac{\mathbb{V}[Q_M]}{N}, \tag{3}$$

where $\mathbb{V}[Q_M]$ denotes the variance of the random variable Q_M.

From this expansion two sources of error in the estimator (1) can be identified. The first term in (3) is the square of the *bias error*. This arises since one is actually interested in the expected value $\mathbb{E}[Q]$ of Q, the (inaccessible) random variable corresponding to the exact solution without any FE error. Assuming convergence of the FE method for each sample, it follows that

$$|\mathbb{E}[Q_M - Q]| \approx M^{-\alpha}, \quad \text{as} \quad M \to \infty, \tag{4}$$

where $\alpha > 0$ is the order of convergence.[1] The particular value of α is problem dependent, and depends in particular on the dimension of the problem, the choice of elements, the smoothness of the model parameters and of the functional $\mathcal{G}$ that defines the quantity of interest

[1] Here and below, the notation $a_n \approx n^{-\alpha}$ for the sequence a_n means that there are constants $C_1 \geq c_1 > 0$ independent of n, such that $c_1 n^{-\alpha} \leq a_n \leq C_1 n^{-\alpha}$, i.e. the sequence decreases with a rate of $-\alpha$ as $n \to \infty$.

(see e.g. [34]). Thus, the approximation error can be reduced below any prescribed *(bias) tolerance* (ϵ_b) by making M sufficiently large.

The second term in (3) gives the *sampling error* since $\mathbb{E}[Q]$ is only approximated by averaging N samples. To ensure this term is smaller than a *(sample) tolerance* ϵ_s^2, it suffices to choose

$$N \approx V \epsilon_s^{-2}, \tag{5}$$

where $V = \mathbb{V}[Q] \approx \mathbb{V}[Q_M]$. Finally, choosing

$$\epsilon_s^2 = \theta \epsilon^2 \quad \text{and} \quad \epsilon_b^2 = (1-\theta)\epsilon^2, \quad \text{for some} \quad 0 < \theta < 1, \tag{6}$$

the total mean square error is less than ϵ^2, for any user-provided overall error tolerance $\epsilon > 0$. Thus, in order to reduce the total error in (2) it is necessary to increase both the number of degrees of freedom M and the number of samples N. This very quickly leads to an intractable problem when the cost to compute each sample to a sufficiently high accuracy is high.

The cost C_M to compute one sample $Q_M^{(j)}$ of Q_M, in terms of floating point operations (FLOPs) or CPU time, depends on the complexity of the FE solver. Typically it will grow like

$$C_M \approx M^\gamma, \tag{7}$$

for some $\gamma \geq 1$. The rate γ is again problem dependent and affected in particular by the dimension, the choice of elements, the sparsity of the stiffness matrix and the solver; it is typically somewhere between 1 and 2. Combining the bounds in (4), (5) and (7), the total cost to achieve a root mean square error $e(\widehat{Q}^{\mathrm{MC}}_{M,N}) \leq \epsilon$ with standard MC can be bounded by

$$\mathrm{Cost}(\widehat{Q}^{\mathrm{MC}}_{M,N}) \approx N M^\gamma \approx \epsilon^{-2-\gamma/\alpha}, \tag{8}$$

where the hidden constants are again independent of M, N and ϵ.

2.2. Multilevel Monte-Carlo simulation (MLMC)

Multilevel Monte Carlo simulation (MLMC) [23,24] seeks to reduce the variance of the estimator in (1) and thus to reduce computational time by recursively using a hierarchy of FE models as control variants. The standard MC estimator in the previous section was too costly because all samples were computed to the required level of accuracy to guarantee a sufficiently small discretisation (or bias) error. A hierarchy of FE models is introduced by refinement of a coarse mesh as shown in Fig. 1. Each mesh corresponds to a *level* $0 \leq \ell \leq L$ in our multilevel method with $M_0 < \cdots < M_\ell < \cdots < M_L$ degrees of freedom, respectively. Typically, M_0 is taken to be much smaller than M_L, yet how small is problem specific.

By exploiting the linearity of the expectation operator $\mathbb{E}[\cdot]$, the MLMC method avoids estimating $\mathbb{E}[Q]$ directly on the finest, most computationally expensive, level L. Instead it estimates the mean on the coarsest level, and corrects this mean by successively adding estimates of the expected values of the differences $Y_\ell = Q_{M_\ell} - Q_{M_{\ell-1}}$ between subsequent levels, using the identity

$$\mathbb{E}[Q_M] = \mathbb{E}[Q_{M_0}] + \sum_{\ell=1}^{L} \mathbb{E}[Y_\ell]\,. \tag{9}$$

The MLMC estimator for $\mathbb{E}[Q]$ is then given by

$$\widehat{Q}^{\mathrm{ML}}_M = \widehat{Q}^{\mathrm{MC}}_{M_0,N_0} + \sum_{\ell=1}^{L} \widehat{Y}^{\mathrm{MC}}_{\ell,N_\ell} \tag{10}$$

where the numbers of samples N_ℓ on each level are judiciously chosen to minimise the total cost of this estimator for a given prescribed sampling error (see below). Note that samples $Y_\ell^{(j)}$ of Y_ℓ require the FE approximations $Q^{(j)}_{M_\ell}$ and $Q^{(j)}_{M_{\ell-1}}$ on two consecutive mesh levels, i.e. two solves, but crucially both with the same random input parameters.

The cost of the MLMC estimator is

$$\mathrm{Cost}(\widehat{Q}^{\mathrm{ML}}_{M_L}) = \sum_{\ell=0}^{L} N_\ell C_\ell\,, \tag{11}$$

where C_ℓ is the cost to compute a single sample of Y_ℓ on level $\ell \geq 1$ or of Q_{M_0} on level 0. By using statistically independent samples across all the levels, the mean square error of $\widehat{Q}^{\mathrm{ML}}_M$ expands to

$$e(\widehat{Q}^{\mathrm{ML}}_M)^2 = \left(\mathbb{E}[Q_M - Q]\right)^2 + \sum_{l=0}^{L} N_\ell^{-1} V_\ell\,, \tag{12}$$

where $V_0 = \mathbb{V}[Q_{M_0}]$ and $V_\ell = \mathbb{V}[Y_\ell]$, for $\ell \geq 1$. This leads to a hugely reduced variance of the estimator since both FE solutions Q_{M_ℓ} and $Q_{M_{\ell-1}}$ converge to Q and thus

$$V_\ell = \mathbb{V}[Q_{M_\ell} - Q_{M_{\ell-1}}] \to 0 \quad \text{as} \quad M_\ell \to \infty.$$

Again, it can be assumed that there exists a $\beta > 0$ such that

$$V_\ell \approx M_\ell^{-\beta}. \tag{13}$$

As for the standard MC estimator, one can ensure that the bias error is less than ϵ_b by choosing $M = M_L$ on the finest level sufficiently large to satisfy (4). To choose the numbers of samples N_ℓ on each of the levels and thus to ensure that the sampling error is less than ϵ_s, there is still some freedom that can be used to minimise the computational cost of the overall MLMC algorithm. The samples per level are chosen by a constrained optimisation problem which minimises $\mathrm{Cost}(\widehat{Q}^{\mathrm{ML}}_M)$ in (11) with respect to $N_0, \ldots, N_\ell$, subject to the constraint that the sampling error of the multilevel estimate (9) is equal to the required tolerance, i.e.,

$$\sum_{\ell=0}^{L} N_\ell^{-1} V_\ell = \epsilon_s^2.$$

This leads to

$$N_\ell \;=\; \epsilon_s^{-2} \left(\sum_{\ell=0}^{L} \sqrt{V_\ell C_\ell} \right) \sqrt{\frac{V_\ell}{C_\ell}} \tag{14}$$

Thus, the total cost of computing (9) can be shown (cf. [24]) to grow like

$$\mathrm{Cost}(\widehat{Q}^{\mathrm{ML}}_M) \;=\; \epsilon^{-2} \left(\sum_{\ell=0}^{L} \sqrt{V_\ell C_\ell} \right)^2 \;\approx\; \epsilon^{-2-\max\left(0, \frac{\gamma-\beta}{\alpha}\right)} \quad \text{as} \quad \epsilon \to 0, \tag{15}$$

where α, β and γ are as defined in (4), (7) and (13) and ϵ is again the tolerance for the total root mean square error.[2] The number of levels L has to grow with $\log(\epsilon^{-1})$.

There are three regimes for the behaviour of the computational cost of the MLMC algorithm:

1. If the variance V_ℓ decays faster than the inverse cost C_ℓ^{-1} (with respect to ℓ), i.e $\beta > \gamma$, then the work on level 0, which is proportional to ϵ^{-2}, dominates the total cost as $\epsilon \to 0$.
2. If V_ℓ decays slower than C_ℓ^{-1}, i.e $\beta < \gamma$. then the work on level L dominates the total cost as $\epsilon \to 0$. This cost is proportional to $\epsilon^{-2-\frac{\gamma-\beta}{\alpha}}$.
3. If $V_\ell C_\ell$ is bounded, i.e. $\beta = \gamma$, then the work is spread evenly over all levels and the hidden constant in (15) grows with $L^2 \approx (\log(\epsilon^{-1}))^2$.

By estimating α, β and γ from the samples, as they come in, a user can determine which regime their problem will fall in. This then indicates the expected speed-up of MLMC when compared to standard Monte Carlo. In practice, it is irrelevant though.

[2] In practice, the values for N_ℓ need to be rounded up to the nearest integer, but that has no impact on the asymptotic cost as $\epsilon \to 0$.

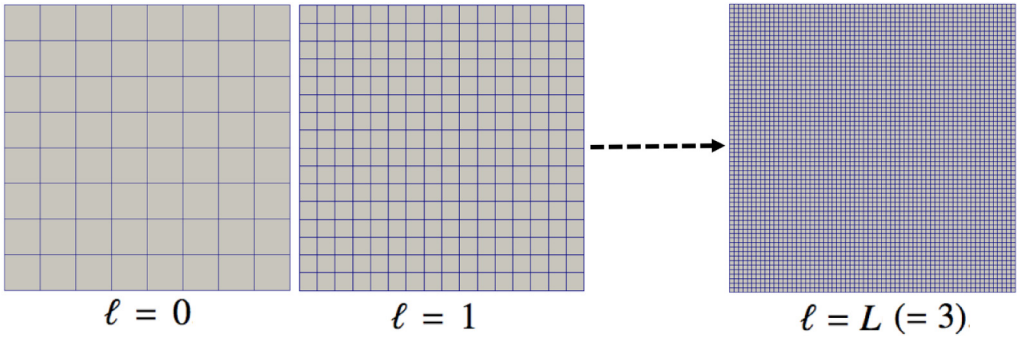

Fig. 1. Example hierarchy of two-dimensional, quadrilateral finite element meshes for the multilevel algorithm achieved through uniform refinement.

2.3. Implementation of MLMC

In this section, the practical implementation of the MLMC algorithm is discussed, as well as a way to compute the (optimal) values of L, M_ℓ and $\{N_\ell\}_{\ell=0}^{L}$ 'on the fly' from the sample averages and the sample variances of Y_ℓ. For ease of presentation, let $Y_0 = Q_{M_0}$. The presentation is also restricted to the case of uniform mesh refinement, where the mesh size is simply halved each time, i.e. $h_\ell = 2^{-\ell} h_0$, but this is not essential [35].

The aim is to estimate $\mathbb{E}[Q]$ under a prescribed tolerance ϵ on the RMSE, which is made up of two parts, the bias error and the sampling error (3). Firstly, to estimate the bias error, M_ℓ is assumed to be sufficiently large, so that the FE method is in the asymptotic regime where

$$\left|\mathbb{E}[Q_{M_\ell} - Q]\right| \approx M_\ell^{-\alpha} \tag{16}$$

and the ratio of the asymptotic constants $r = c_1/C_1$ of the asymptotic constants in $\approx$ (cf. Footnote 1) is close to 1 for $M > M_\ell$.

For uniform refinement, the number of degrees of freedom on level ℓ is given by $M_\ell \approx m^\ell M_0$. For the two-dimensional numerical examples which follow below, domains are discretised by quadrilateral elements and thus $m = 4$. It follows by the reverse triangle inequality that

$$\begin{aligned}\left|\mathbb{E}[Y_\ell]\right| &= \left|\mathbb{E}[Q_\ell - Q_{\ell-1}]\right| = \left|\mathbb{E}[Q_{\ell-1} - Q] - \mathbb{E}[Q_\ell - Q]\right| \\ &\geq \left|\left|\mathbb{E}[Q_{\ell-1} - Q]\right| - \left|\mathbb{E}[Q_\ell - Q]\right|\right|.\end{aligned} \tag{17}$$

By noting that due to (16) $\mathbb{E}[Q_{l-1} - Q] \geq r m^\alpha \mathbb{E}[Q_l - Q]$ it follows that

$$|\mathbb{E}[Y_\ell]| \geq (r m^\alpha - 1)\left|\mathbb{E}[Q_\ell - Q]\right| \tag{18}$$

In the numerical experiments, the choice $r = 1$ was made. This is equivalent to the assumption above that M_ℓ is sufficiently large and that the FE error is in the asymptotic regime. For each test case, this assumption is verified numerically in this paper. One can of course also choose more prudent values, such as $r = 0.9$ or $r = 0.7$. Rearranging (18) and replacing $|\mathbb{E}[Y_\ell]|$ by the Monte Carlo estimate $\hat{Y}^{MC}_{\ell,N_\ell}$, the bias error on level ℓ can be estimated by

$$\mathcal{E}_\ell := \left|\mathbb{E}[Q_{M_\ell} - Q]\right| \leq \frac{1}{r m^\alpha - 1}\,\hat{Y}^{\mathrm{MC}}_{\ell,N_\ell}. \tag{19}$$

The sample variance is estimated in the standard way

$$s_\ell^2 = \left(\frac{1}{N_\ell}\sum_{j=1}^{N_\ell}\left(Y_\ell^{(j)}\right)^2\right) - \left(\hat{Y}^{\mathrm{MC}}_{\ell,N_\ell}\right)^2 \approx V_\ell. \tag{20}$$

The adaptive method is summarised in Algorithm 1. Since each sample is independent, Algorithm 1 can be readily parallelised by distributing samples across many processors.

Remark 1. In this paper, a hierarchy of levels is created by uniform refinement of a coarse mesh. In this case for each random sample, via FE interpolation and projection, there is a unique connection between the models at different levels. However, to implement the MLMC approach there is no constraint that the models must be nested, or have exactly the same input variables. For example, in an engineering application the coarse model could be a beam/shell model of a structure, whilst the fine model could be a full 3D stress analysis. The only constraint is that both models must output an estimate for the same quantity of interest Q, so that the difference can be computed.

Algorithm 1: Multilevel Monte Carlo Algorithm

```
Set ϵ, θ, N*.
Set L = 0, converged = false.
while converged == false do
    Compute N_L = N* samples on level L.
    for ℓ = 0 to L do
        Estimate V_ℓ from samples on level ℓ using (20).
        Estimate optimal number of samples N_ℓ^opt on level ℓ using (14).
        if N_ℓ < N_ℓ^opt then
            Compute N_ℓ^opt − N_ℓ additional samples on level ℓ.
            Set N_ℓ = N_ℓ^opt.
        end if
    end for
    Estimate bias error E_L on level L using (19).
    if E_L < ϵ_b then
        Set converged = true
    else
        L = L + 1
    end if
end while
```

3. Multilevel Monte Carlo simulation with selective refinement for the computation of failure probabilities

For many engineering applications, estimating the expected value of a specific quantity is of limited interest. Instead, often it is of greater interest to compute the probability that the failure load λ, e.g. through buckling, is less than a 'safe' load $\lambda^\star$. Within the MLMC framework, the quantity of interest is then the binomially distributed random variable $Q = \mathbb{1}(\lambda < \lambda^\star)$, which takes the value 1 if $\lambda < \lambda^\star$ and 0 otherwise. The failure probability can then be approximated by evaluating $\mathbb{E}[Q] = \mathbb{P}(\lambda < \lambda^\star)$.

For aerospace applications these probabilities are necessarily small, and obtaining good estimates for such rare events is difficult since, by definition, a large number of samples are required to observe even a single case of failure. One of the main issues is that a simple binomial distribution ($Q \in \{0, 1\}$) loses important information on how close a given sample is to failing. In particular, as a function of the failure load λ, Q is a step function. One proposed method for improving convergence is to 'smooth' the quantity of interest close to the critical value

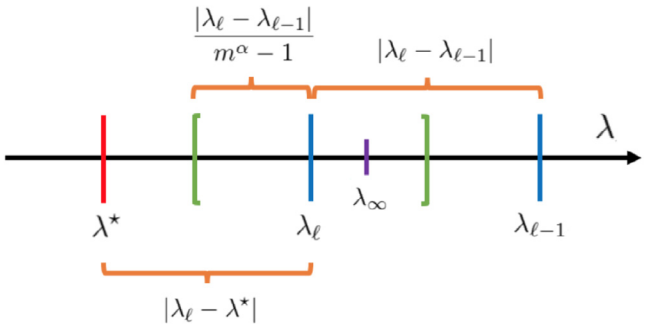

Fig. 2. Graphical representation of MLMC with selective refinement.

Algorithm 2: Selective refinement for one sample in a failure probability calculation.

1: Given critical load $\lambda^\star$, constant r, level ℓ, and sample number j.
2: **for** levels $i = 0, \dots, \ell$ **do**
3: Compute $\lambda_i^{(j)}$ on level i.
4: **if** $i > 1$ **then**
5: **if** $|\lambda_i^{(j)} - \lambda^\star| > |\lambda_i^{(j)} - \lambda_{i-1}^{(j)}|/(rm^\alpha - 1)$ **then**
6: Set $\lambda_\ell^{(j)} = \lambda_i^{(j)}$.
7: Exit **for loop**
8: **end if**
9: **end if**
10: **end for**
11: Evaluate failure probability functional $Q_\ell^{(j)} = \mathbb{1}(\lambda_\ell^{(j)} - \lambda^\star)$

$\lambda^\star$ and to allow it to take intermediate values between 0 and 1 [23]. Here, however, a different approach is proposed, which combines the error estimator in Eq. (19) and the multilevel framework presented in Section 2.2, motivated from an approach proposed in [29].

Following a similar calculation to (19), but now sample-wise, it follows that the FE error for the jth sample on level ℓ can be estimated by

$$|\lambda_\ell^{(j)} - \lambda^{(j)}| \leq \frac{|\lambda_\ell^{(j)} - \lambda_{\ell-1}^{(j)}|}{rm^\alpha - 1} \tag{21}$$

with $r \approx 1$ for M_ℓ sufficiently large. Then, using again the reverse triangle inequality

$$|\lambda^{(j)} - \lambda^\star| \geq \left| |\lambda_\ell^{(j)} - \lambda^\star| - |\lambda^{(j)} - \lambda_\ell^{(j)}| \right| \tag{22}$$

and it suffices to check whether

$$|\lambda_\ell^{(j)} - \lambda^\star| > \frac{|\lambda_\ell^{(j)} - \lambda_{\ell-1}^{(j)}|}{rm^\alpha - 1} \tag{23}$$

to ensure that $Q_{\ell'}^{(j)} = Q_\ell^{(j)}$, for all $\ell' \geq \ell$, due to Eq. (22). Thus, there exists a sample-dependent level $\ell_j \geq 0$, such that $Q_\ell^{(j)} = Q_{\ell_j}^{(j)}$, for all $\ell \geq \ell_j$. This implies that especially on the finer levels $\ell \gg 0$ most samples will not have to be computed with full accuracy, if they are sufficiently far from $\lambda^\star$ (as illustrated in Fig. 2).

This selective refinement technique for MLMC is summarised in Algorithm 2. Each sample is first computed on level 0 and then the FE mesh is selectively refined until ℓ_j is reached. Subsequent refinements will not change the value of $Q_{\ell_j}^{(j)}$ thus saving valuable computational time. The key point is that this modification simply reduces the average cost per sample on the finer levels, whilst the original multilevel algorithm (as described in Algorithm 1) remains unchanged.

Elfverson et al. [29] formalised the gains of MLMC with selective refinement (MLMC-SR) over MLMC and standard MC for computing failure probabilities. They showed that, provided $|\mathbb{E}[\lambda_\ell - \lambda]| \eqsim M_\ell^{-\alpha}$ and $C_\ell \eqsim M_\ell^\gamma$, where C_ℓ denotes the cost to compute one sample $\lambda_\ell^{(j)}$ of λ_ℓ, the expected cost to compute one sample of the failure probability functional Q_ℓ on level ℓ using the selective refinement method (as presented in Algorithm 2) behaves like

$$\mathbb{E}\left[\text{Cost}(Q_\ell)\right] \eqsim M_\ell^{\gamma-\alpha}. \tag{24}$$

The order of growth, with respect to degrees of freedom M_ℓ, is shown to be significantly reduced from the corresponding cost per sample for the basic MLMC simulation. This is due to the fact that only a fraction of samples are solved on their highest refinement level, with work instead concentrated on the lower (computationally cheaper) levels.

However, there are some significant practical challenges for MLMC that are inherited by MLMC-SR: If p^+ is the probability of observing a failure on level ℓ and not on level $\ell - 1$ (i.e. $Y_\ell = 1$) and p^- is the probability of failure on level $\ell - 1$ but not on level ℓ (i.e. $Y_\ell = -1$), then p^+ and $p^- \to 0$ as $\ell \to \infty$, making it extremely challenging to estimate the mean and variance of Y_ℓ accurately; in particular, $\mathbb{E}[Y_\ell] = p^+ - p^-$ and $\mathbb{V}[Y_\ell] = p^+ + p^- - (p^+ - p^-)^2$. In [29], the following family of estimators for the probabilities $p^\pm$ are proposed to mitigate this problem:

$$\widehat{p}_k^\pm = \frac{x^\pm + k}{N_\ell + k}, \quad \text{for} \quad k \in \mathbb{N}_0, \tag{25}$$

where $x^\pm$ denote the numbers of samples for which $Y_\ell = \pm 1$ across N_ℓ samples, respectively. These estimators are biased for $k \geq 1$. But while the relative variance of the standard, unbiased estimator $\widehat{p}_0^+$ for p^+ (i.e., with $k = 0$) explodes as $p^+ \to 0$, the relative variance of the estimators for $k \geq 1$ can be bounded such that

$$\frac{\mathbb{V}[\widehat{p}_k^+]}{\mathbb{E}[\widehat{p}_k^+]^2} \leq \frac{p^+ N_\ell}{(p^+ N_\ell + k)^2} < 1 \tag{26}$$

(cf. [29, Sect. 7]). Large values of k give a large bias in the estimator, but a smaller relative variance and thus higher accuracy. In particular, the bias of the estimator is significant if $p^+ N_\ell \ll k$ and there are too few samples to estimate p^+ accurately. However, $\widehat{p}_k^+$ still acts as an upper bound in that case. The same holds true for $\widehat{p}_k^-$. However, an additional worry is the possible cancellation in the final estimates of $\mathbb{E}[Y_\ell]$ and $\mathbb{V}[Y_\ell]$, which is critical here, since these estimates are used to bound the numerical bias and sampling error and control the stopping criteria for the MLMC algorithm.

Remark 2. The situation is less pronounced for stochastic eigenvalue problems with a nested hierarchy of grids, as considered in Section 5.2 below. Due to the min–max principle for eigenvalue problems [36], $\lambda_\ell^{(j)} \leq \lambda_{\ell-1}^{(j)}$. This guarantees that $p^- = 0$ and avoids cancellation errors.

4. Example I — compressive strength of fibre composites with random fibre misalignment

It is well established that the compressive failure of undamaged composites is primarily governed by plastic micro-buckling (or kinking) of the fibres [32,37], and this failure is initiated in regions of local fibre misalignment or waviness. The classical micro-mechanical model for the compressive strength of a composite σ given by Budiansky [32] is

$$\sigma = \frac{G}{1 + |\Phi|/\gamma_y} \tag{27}$$

where G, Φ and γ_y are the shear modulus, fibre misalignment angle and shear strain at failure, respectively. This idealised model (often referred to as *kinking theory*) assumes the misalignment or kink of known angle Φ. Observations of real fibre waviness show that the misalignment Φ is not a single value, but a complex random field, as seen in Fig. 3 (left). In practical applications, it is then unclear what value of misalignment Φ to use in (27); possible options include the root mean square or the maximum misalignment. In fact, the compressive strength is also a random variable, with a distribution intricately coupled with the statistical distribution of Φ. Here, the uncertainty in the angle is modelled using a spatial random field, as shown in Fig. 3 (right) and

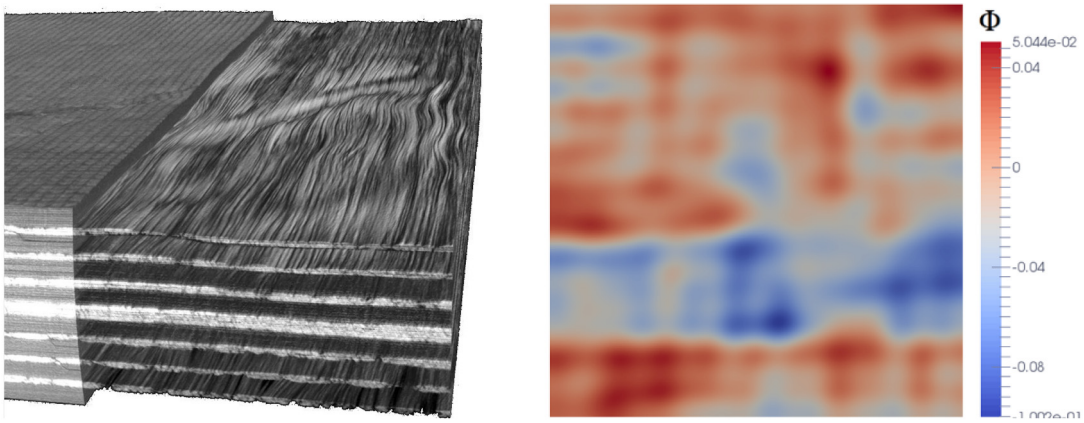

Fig. 3. Left: CT image showing random fibre waviness within a composite laminate. Right: Sample of random waviness field Φ with $N_{\mathrm{KL}} = 400$ and covariance parameters as in (39) below.

parameterised by observed statistics [3]. The computational savings of the MLMC are demonstrated by estimating the expected compressive strength of a composite $\mathbb{E}[\sigma]$ with random fibre waviness modelled in this way.

4.1. A two-dimensional random Cosserat continuum model

In this example problem, a square domain Ω in the (x, y) plane made up of uni-directional composite pre-preg is considered. Individual fibres are misaligned by an angle $\Phi(\mathbf{x})$ to the x-axis. This misalignment is modelled as a random field on Ω. The mean and the covariance structure of Φ will be inferred from measurements of fibre misalignment of carbon fibre pre-pregs available in the literature [3]. The random field Φ is characterised by a two-point exponential covariance function

$$c(\mathbf{x}, \mathbf{y}) = s_{\Phi}^2 \exp\left(-\frac{|x_1 - y_1|}{\omega_1} - \frac{|x_2 - y_2|}{\omega_2} \right). \tag{28}$$

The parameters s_{Φ}^2 and ω_i denote the variance and correlation length (in each direction) of the misalignment field. To generate a single realisation of this random field, the misalignment field is modelled using a Karhunen–Loève (KL) expansion, an expansion in terms of a countable set of basis functions $\phi_n(x)$, parameterised by a sequence of independent standard Gaussian random variables $\{\xi_n\}_{n\in\mathbb{N}}$, given by

$$\Phi(\mathbf{x}) = \sum_{n=1}^{\infty} \sqrt{\mu_n}\phi_n(\mathbf{x})\xi_n. \tag{29}$$

Here, $\{\mu_n\}_{n\in\mathbb{N}}$ and $\{\phi_n\}_{n\in\mathbb{N}}$ are the eigenvalues and associated (normalised) eigenfunctions of the covariance function (28). Note that the eigenvalues $\{\mu_n\}_{n\in\mathbb{N}}$ are positive and strictly decreasing which provides a natural ordering of the importance of the contribution of each term to $\Phi(\mathbf{x})$. In a computational setting it is therefore natural to truncate the KL expansion (29) after N_{KL} terms, giving a parameterisation of the random field by the set of variables $\boldsymbol{\xi} = [\xi_1, \xi_2, \ldots, \xi_{N_{\mathrm{KL}}}]$. Fig. 3 (right) shows a realisation of the random field generate using the approach described with $N_{\mathrm{KL}} = 400$. For further details on random fields and their implementation within structural applications, the reader is referred to the classical text by Spanos and Ghanem (2003) [20] or related article [38]. Note that it is possible to implement more complex covariance functions, and implement them on more complex geometries, see for example [39].

A key consideration when modelling the mechanics of such a composite, is that the shear stiffness parallel to the fibres is an order of magnitude less than the orthogonal shear stiffness; and hence, in general, the stress state is non-symmetric, i.e. $\sigma_{12} \neq \sigma_{21}$. As a result, a finite size element of composite carries a coupled stress (a moment per unit area), and the fibres bend to achieve moment equilibrium. A classical approach to capturing these internal bending effects is to model the composite as a Cosserat Continuum [37,40]. Here, under plane-strain assumptions, each material point has the conventional displacement degrees of freedom u_1 and u_2 (u and v in global coordinates), as well as an independent (Cosserat) rotational degree of freedom θ_3. Under the assumption of small deformations and rotations, this gives the small Cosserat strain and curvature measures

$$\varepsilon_{ij} = \frac{du_i}{dx_j} + e_{ij3}\theta_3 \quad \text{and} \quad \kappa_{ij} = \frac{d\theta_3}{dx_j}, \tag{30}$$

where e_{ijk} denotes the permutation tensor. The permutation tensor is defined as $e_{123} = e_{312} = e_{231} = 1$, $e_{213} = e_{132} = e_{321} = -1$ and $e_{ijk} = 0$ if any indices are repeated, e.g. $e_{112} = 0$. These strain and curvature measures are work conjugates to the Cosserat stresses σ_{ij} and coupled-stress m_{ij}, respectively. The linear Cosserat constitutive relationships (derived in [40, Sect. 2.3]) can be expressed in matrix form as

$$\underline{\sigma} = \underline{\underline{C}}\,\underline{\varepsilon} \quad \text{and} \quad \underline{m} = \underline{\underline{D}}\,\underline{\kappa}.$$

The fourth-order tensors $\underline{\underline{C}}$ and $\underline{\underline{D}}$ – identified through the double underline but here represented as matrices – can be rotated by the misalignment angle $\Phi(\mathbf{x})$ to the global coordinate system via the transformation matrices $T^{\varepsilon}_{\phi(\mathbf{x})}$ and $T^{\kappa}_{\phi(\mathbf{x})}$ so that the global tensors become

$$\underline{\underline{\tilde{C}}}_{\phi(\mathbf{x})} = (T^{\varepsilon}_{\phi(\mathbf{x})})^{-1}\,\underline{\underline{C}}\,T^{\varepsilon}_{\phi(\mathbf{x})} \quad \text{and} \quad \underline{\underline{\tilde{D}}}_{\phi(\mathbf{x})} = (T^{\kappa}_{\phi(\mathbf{x})})^{-1}\,\underline{\underline{D}}\,T^{\kappa}_{\phi(\mathbf{x})}. \tag{31}$$

The force and moment equilibrium equations for a small element of composite, in the absence of body forces and coupling are given by

$$\frac{d\sigma_{ij}}{dx_j} = 0 \quad \text{and} \quad \frac{dm_{ij}}{dx_j} + e_{ijk}\sigma_{jk} = 0. \tag{32}$$

In our model, these equilibrium equations are subject to the Dirichlet boundary conditions

$$u(\mathbf{x}) = 0 \text{ at } x_1 = 0, \qquad u(\mathbf{x}) = \Delta < 0 \text{ at } x_1 = L, \tag{33}$$

$$\text{and} \quad v(\mathbf{x}) = 0 \text{ at } x_2 = 0 \text{ and } x_2 = L. \tag{34}$$

To solve (32) using the FE method, the differential equations are recast as a variational problem: Find a solution $(\mathbf{u}, \theta_3) \in V^2 \times W$, such that for all test functions $(\mathbf{v}, \zeta_3) \in V^2 \times W$ the variational equation

$$\int_{\Omega} \underline{\underline{\tilde{C}}}_{\phi(\mathbf{x})}\,\underline{\varepsilon}(\mathbf{u}, \theta_3) : \underline{\varepsilon}(\mathbf{v}, \zeta_3) + \underline{\underline{\tilde{D}}}_{\phi(\mathbf{x})}\,\underline{\kappa}(\theta_3) : \underline{\kappa}(\zeta_3)\, d\mathbf{x} = \int_{\Gamma} \underline{t}\,\mathbf{v} + \mu\zeta_3\, d\mathbf{x}, \tag{35}$$

holds. Here, $\underline{t}$ denote the stress traction, and μ the coupled stress traction on the boundary of the domain Γ. The spaces V and W are appropriate function spaces on which the components of $\mathbf{u}$ and the Cosserat rotation θ_3 are defined. Here, an appropriate choice is the

Sobolev Space H^1; that is, the space of all square integrable functions with square integrable first derivatives satisfying the boundary conditions.

To approximate (35), the domain Ω is uniformly discretised into a set of 4-node quadrilateral elements

$$\mathcal{Q}_h = \{\Omega_e^{(i)}\}_{i=1}^{n_{\text{el}}},$$

where n_{el} denotes the number of elements and h is the side-length of the elements. The solution is approximated by restricting (35) to the finite dimensional subspace $V_h^2 \times W_h \subset V^2 \times W$. In these examples V_h and W_h are chosen to be the set of piecewise bi-linear functions on $\mathcal{Q}_h$, and the corresponding approximate solutions are denoted by $\mathbf{u}_h$ and θ_h. As for any standard FE analysis, substitution of the approximations $\mathbf{u}_h$ and θ_h allows (35) to be rewritten as a linear system of the form

$$\mathrm{K}\mathbf{d} = \mathbf{f} \tag{36}$$

where $\mathrm{K} \in \mathbb{R}^{M\times M}$ is the global stiffness matrix and $\mathbf{f} \in \mathbb{R}^M$ is the load vector due to the prescribed boundary conditions. The vector $\mathbf{d} \in \mathbb{R}^M$ contains the coefficients of all degrees of freedom in the expansions of $\mathbf{u}_h$ and θ_h above. If n_{node} is the total number of nodes in the grid, then $M \approx 3n_{\text{node}}$ (taking into account boundary conditions).

From the solution $\mathbf{d}$, the objective is to calculate the compressive strength σ_c of the composite using the quadratic failure criterion in [32]. To this end, the effective stress τ_e is defined in terms of the transverse stress σ_{22} and the shear-stress parallel to the fibres σ_{12}, such that

$$\tau_e = \sqrt{\sigma_{12}^2 + \left(\frac{\sigma_{22}}{R}\right)^2}. \tag{37}$$

The material parameter R is the ratio of the transverse and shear yield strength of the material. Failure is said to occur when the effective stress is equal to the shear strength τ_y of the material.

In the results which follow, the compressive strength is estimated by first computing $\mathbf{d}$ for a prescribed compressive end-shortening Δ. In order to remove the influence of boundary conditions, the maximum value f^* of $f = \tau_e/\tau_y$ is then found over all integration points within elements contained in a central square subregion Ω_{int} of Ω, which has area $|\Omega_{\text{int}}|$. As the problem under consideration is linear, the compressive strength σ_c is then given by

$$\sigma_c = \frac{f^*}{|\Omega_{\text{int}}|}\int_{\Omega_{\text{int}}} \sigma_x \, d\mathbf{x}. \tag{38}$$

4.2. Results

For the experiments that follow, material parameters for unidirectional pre-preg AS4/8552 are considered, with material constants taken from the Hexcel Data Sheet [41]; in particular:

$$v_f = 0.59, \quad E_f = 230 \text{ GPa}, \quad E_m = 9.25 \text{ GPa}, \quad G_f = 95.83 \text{ GPa},$$
$$G_m = 5.13 \text{ GPa}, \quad d = 7 \text{ μm} \quad \text{and} \quad \tau_c = 114 \text{ MPa}.$$

The stochastic model for random misalignment is parameterised based on data in the literature; in particular, the measurements of in-plane waviness in pre-preg given by Sutcliffe et al. [3] which agree well with other values given in [37,42,43]. However, the correlation lengths ω_1 and ω_2 in (28) are defined differently to those in [3], which will be denoted $\omega_1^\star$ and $\omega_2^\star$ in the following. There, the correlations lengths were defined as the lag at which the auto-correlation function equals 0.1, i.e. when $c(\mathbf{x},\mathbf{y})/\sigma_\phi^2 = 0.1$. Therefore, $\omega_i = -\omega_i^\star/\log(0.1)$. In summary, the following values for the parameters in the covariance function (28) are chosen:

$$\omega_1 = 229d, \quad \omega_2 = 61d \quad \text{and} \quad s_\Phi = 0.035 \text{ rad}. \tag{39}$$

Fig. 3 (right) shows a random field generated with the above parameters. Having fixed the correlation lengths of the wrinkles, the domain size is chosen to be $L = 2.5\omega_1$. Furthermore, Ω_{int} (as introduced in (38)) is chosen to be the square subdomain centred in Ω with sides of length $1.25\omega_1$. The coarsest finite element grid (level $\ell = 0$) has a mesh size of $h_0 = L/8$ (i.e. with 64 elements and $M_0 = 243$ degrees of freedom), and subsequent grids are generated by uniform refinement as shown in Fig. 1. The number of KL modes is also increased with the levels, such that $N_{\text{KL},\ell} = 50 + 50\ell$.

To compare the MLMC algorithm with standard MC, it is first useful to study how the computational cost scales with ℓ and to estimate the parameter γ in (7). Averaging the measured CPU times over 100 samples for $\ell = 0, \dots, 5$ leads to an estimate of

$$C_\ell \approx M_\ell^{1.3}, \tag{40}$$

and hence, a value of $\gamma \approx 1.3$, which is used in the calculations that follow.

For the quantity of interest, the compressive strength $Q = \sigma_c$ defined in (38), the parameters α in (4) and β in (13) can be estimated from a series of sufficiently accurate MLMC simulations. Fig. 4 shows the log–log plots of the estimated means and variances of Q_ℓ and $Y_\ell = Q_\ell - Q_{\ell-1}$, for $\ell = 0, \dots, 5$, with respect to the total number of degrees of freedom M_ℓ on each level. Looking first at the behaviour of the expectation of Q_ℓ and Y_ℓ (left), one can observe that

$$\mathbb{E}[Y_\ell] \approx M_\ell^{-0.786}$$

approximately, and hence $\alpha \approx 0.786$. The variance plot (right) implies $\beta \approx 0.740$, i.e.,

$$\mathbb{V}[Y_\ell] \approx M_\ell^{-0.740}$$

From those estimates, Eq. (15) allows to predict that the cost of the MLMC simulation is expected to grow proportionally to $\epsilon^{-2.68}$, whilst the cost of the standard MC simulation should grow like $\epsilon^{-3.64}$, as the overall tolerance for the estimator $\epsilon \to 0$.

Fig. 5 (left) compares the computational cost of the MLMC simulation versus standard MC across a range of tolerance values from $\epsilon = 42.13$ MPa to 3.06 MPa, corresponding to relative errors of 3% to 0.2%. The predicted cost estimates above are well matched by the numerical experiments: Considering the gradients of the plots in Fig. 5 (left), the costs of MLMC and standard MC can be seen to be approximately proportional to $\epsilon^{-2.64}$ and $\epsilon^{-3.22}$, respectively. More detail about the numerical results can be found in Table 1, including the optimal numbers of samples N_ℓ across the levels as given by (14), as well as the actual CPU times and the speed-up factors for each of the tolerances. In particular, it can be seen that for an absolute error of $\epsilon = 3.06$ MPa the MLMC algorithm reduces the computational cost by a factor of more than 16 over standard MC; in absolute terms, this reduces the computational time from more than 28 hours to well under 2 hours.

Whilst the primary aim of this paper has been to demonstrate the computational savings of the MLMC methodology, the results also provide an opportunity to compare the results to theoretical and experimental work in the literature. A parameter of particular interest is the standard derivation s_Φ of the misalignment field, and the effect it has on the compressive strength of a composite material. Using the multilevel methodology (with $L = 4$), Fig. 5 (right) shows the estimated means $\mathbb{E}(\sigma_c)$ (blue circles) and the estimated 10th percentiles (red stars), as well as the worst cases for 8,000 samples on level 4 (red diamonds), as functions of the standard deviation s_Φ of the misalignment field. The results are compared to the classical Budiansky 'kinking' model (27) and also the Hexcel data sheet value of $\sigma/\tau_y = 13.43$ for AS4/8552. The estimates for the 10th percentile agree very well the Budiansky model for $s_\Phi \geq 2$. Both models predict a significant decrease in compressive strength with increasing variance of fibre misalignment. Discrepancies between the two models are observed at lower variances in misalignment angles. This suggests that small misalignment angles are not dominated by a shear failure, but by failure in the σ_{22} direction. Through thickness tensile failures are not accounted for in the Budiansky model [32], and therefore this is expected.

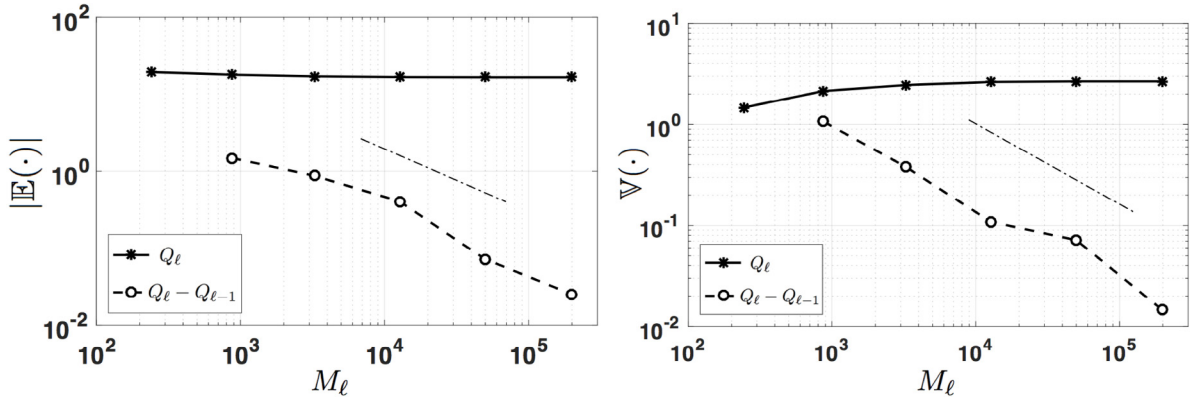

Fig. 4. Expected value (left) and variance (right) of Q_ℓ and $Y_\ell = Q_\ell - Q_{\ell-1}$ against degrees of freedom M_ℓ for Example I. The dashed lines show $\alpha \approx 0.786$ (left) and $\beta \approx 0.740$ (right).

Table 1
Optimal sample sizes and detailed cost comparison for MLMC vs. MC (Example I).

ϵ	N_ℓ						Cost (in minutes)		Saving Factor
	0	1	2	3	4	5	MLMC	MC	
3.01%	513	237	34	8	–	–	0.10	0.34	3.32
0.63%	22,014	6,191	1,449	337	123	–	6.65	42.26	6.36
0.22%	240,427	67,611	15,822	3684	1347	283	103.84	1685.50	16.23

5. Example II – buckling performance of a wing skin panel with uncertain ply orientations

In this section, a model problem to test the multilevel Monte Carlo method with selective refinement (MLMC-SR) as described in Section 3 is presented. Here, as an illustrative example for the new methodology, the structural performance of a wing skin panel subject to a typical in-service load is considered.

5.1. Model setup and mathematical description

Consider a rectangular composite plate of thickness t, length $L_x = 636$mm and width $L_y = 212$mm, with the un-deformed mid-plane of the plate occupying the domain $\Omega = [0, L_x] \times [0, L_y]$ with boundary Γ. The laminate is made up of 8 identical, orthotropic, composite plies characterised by the elastic tensor $\underline{\underline{C}}$, thickness 0.8 mm and arranged in a fully uncoupled (Winckler) stacking sequence

$$\boldsymbol{\psi} = [45°, -45°, -45°, 45°, -45°, 45°, 45°, -45°].$$

The elastic ply properties, are taken from the IM7-8552 data sheet, so that $E_{11} = 130.0$ GPa, $E_{22} = 9.25$ GPa, $G_{12} = 5.13$ GPa, $\nu = 0.36$ and $G = 5.13$ GPa.

In this example problem, the as-manufactured ply orientations are assumed to be uncertain due to angle tolerances in the laying machine. Therefore, a small, spatially homogeneous, random perturbation ϕ_i is added to each pristine ply angle ψ_i, for $i = 1, \dots, 8$. In this way, a new "defective" stacking sequence, $\boldsymbol{\psi}^d = [\psi_1^d, \dots, \psi_8^d]$, is obtained. The random angle perturbations are assumed to be normally distributed with $\phi_i \sim \mathcal{N}(0, 3^2)$. The standard deviation of the perturbations has been chosen to conform with the accuracy of automated fibre placement (AFP) machines in the industry. Typically, machines have an allowable error tolerance of 5°. Hence, in order to obtain sample perturbations ϕ_i satisfying this error tolerance with 95% confidence, the required standard deviation is $5°/1.65 \approx 3°$ (where $1.65 = z_{.05}$ is the critical z value for the one-sided 95% confidence interval of a normal distribution).

The deformation of the plate is described by the vertical displacement $w(x, y)$ and rotations of the mid-plane $\boldsymbol{\theta}(x, y) = [\theta_1, \theta_2]^T$. The plate is subjected to uniform, uniaxial compressive stress, whilst being

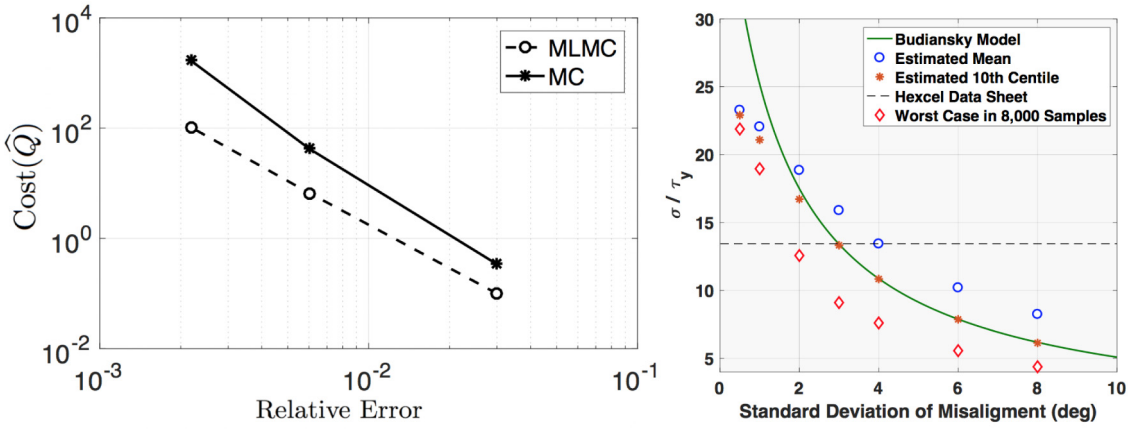

Fig. 5. (Left) Computational cost (CPU time in minutes) against relative error for standard MC (cost $\approx \epsilon^{-3.22}$) and MLMC simulations (cost $\approx \epsilon^{-2.64}$) for Example I. (Right) Normalised compressive strength σ_c/τ_y against standard deviation s_Φ of the misalignment field.

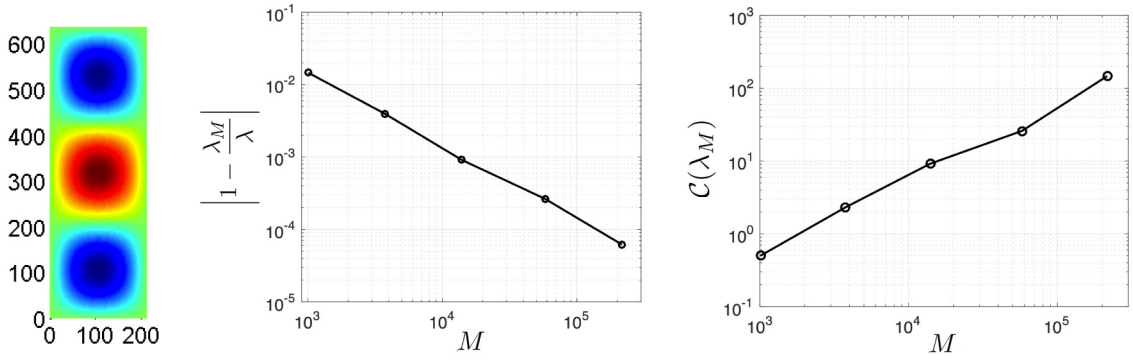

Fig. 6. (Left) Plot of the critical buckling mode corresponding to the critical buckling load of 278.59kN for the pristine panel in Example II. (Middle) Log–Log plot of the relative FE error in the buckling load $|1-\lambda_M/\lambda|$ against M showing a convergence with order $\alpha \approx 1$. (Right) Log–Log plot of Cost (CPU-time in seconds) against M showing a growth with rate $\gamma \approx 1.17$.

simply-supported around all boundaries. The critical buckling load for the plate is calculated using Reissner–Mindlin (RM) plate theory, since the advantages over Kirchhoff Plate theory are well documented [44]. This reduces the problem to a 2D problem in Ω, by applying classical laminate theory (CLT) [45], which gives the laminate stiffness tensors

$$\underline{\underline{A}} = \sum_{k=1}^{K} \underline{\underline{C}}^{(k)}(z_k - z_{k-1}), \quad \underline{\underline{B}} = \frac{1}{2}\sum_{k=1}^{K} \underline{\underline{C}}^{(k)}(z_k^2 - z_{k-1}^2) \quad \text{and} \quad \underline{\underline{D}} = \frac{1}{3}\sum_{k=1}^{K} \underline{\underline{C}}^{(k)}(z_k^3 - z_{k-1}^3), \tag{41}$$

where z_k is the distance from the top edge of the kth ply to the neutral axis of the plate and where $\underline{\underline{C}}^{(k)}$ is the elastic tensor of the kth ply in global coordinates. These homogenised tensors connect in-plane strains $\underline{\varepsilon}$ and out-of-plane curvatures $\underline{\kappa}(\boldsymbol{\theta}) = \frac{1}{2}\left(\nabla\boldsymbol{\theta} + \nabla\boldsymbol{\theta}^T\right)$, with in-plane stress and plate bending moments. Under the additional assumption that the in-plane and out-of-plane behaviour are decoupled, it follows that the in-plane stress and the moment are then given by

$$\underline{\sigma} = t^{-1}\,\underline{\underline{A}}\,\underline{\varepsilon} \quad \text{and} \quad \underline{\mu} = \underline{\underline{D}}^{*}\,\underline{\kappa}, \tag{42}$$

respectively. Here, $\underline{\underline{D}}^{*} = \underline{\underline{D}} - \underline{\underline{B}}^T\,\underline{\underline{A}}^{-1}\,\underline{\underline{B}}$, which conservatively knocks down the bending resistance of the panel to account for coupling effects.

In the absence of body forces, a moment equilibrium for the RM plate gives the linear eigenvalue problem

$$\nabla\cdot(\underline{\underline{D}}^{*}\,\underline{\kappa}(\boldsymbol{\theta})) - kG_{13}(\nabla w - \boldsymbol{\theta}) = \lambda\nabla\cdot(\underline{\sigma}\nabla w) \quad \text{such that} \quad w = 0 \ \text{ and } \ \underline{\mu}\,\mathbf{n} = 0 \ \text{ on } \ \Gamma, \tag{43}$$

where $G(= G_{13} = G_{23} = G_{12})$ is the through thickness shear stiffness and $k = 5/6$ is the shear correction (both constants). This equation is solved using a FE method, and therefore the weak form of the eigenvalue problem is used, such that the problem becomes: Find the smallest (positive real) eigenvalue λ and associated (buckling) eigenmode $0 \neq (\boldsymbol{\theta}, w) \in V^2 \times V$ such that

$$\int_\Omega \underline{\underline{D}}^{*}\,\underline{\kappa}(\boldsymbol{\theta}) : \underline{\kappa}(\boldsymbol{\zeta})\,d\mathbf{x} + kG_{13}\int_\Omega (\nabla w - \boldsymbol{\theta})\cdot(\nabla z - \boldsymbol{\zeta})\,d\mathbf{x} = \lambda\int_\Omega \sigma\nabla w\cdot\nabla z\,d\mathbf{x} \quad \forall(\boldsymbol{\zeta}, z) \in V^2 \times V. \tag{44}$$

The solutions of (44) are again approximated using piecewise bilinear FEs on a quadrilateral mesh $\mathcal{Q}_h$, such that $\boldsymbol{\theta}$ and w are interpolated with the same shape functions $\{\phi_i(\mathbf{x})\}_{i=1}^{n_{\text{node}}}$. The matrix form of the FE approximation of (44) is

$$\mathrm{K}^{\mathrm{B}}\boldsymbol{d}^{\mathrm{B}} = \lambda\,\mathrm{K}^{\mathrm{G}}\boldsymbol{d}^{\mathrm{B}}, \tag{45}$$

where $\mathrm{K}^{\mathrm{B}} \in \mathbb{R}^{M\times M}$ is the global stiffness matrix (related to the LHS of (44)) whilst $\mathrm{K}^{\mathrm{G}} \in \mathbb{R}^{M\times M}$ is the geometric stiffness matrix (related to the RHS of (44)). Further details of the exact FE formulation for a similar eigenvalue problem are provided in [46].

5.2. Comparison between MC, MLMC and MLMC-SR

Before comparing MC, MLMC and MLMC-SR for the second example problem, first the convergence rates for the FE approximation of the critical buckling load λ, as well as the associated computational cost (Cost) are investigated, again using uniform mesh refinement. For the pristine stacking sequence, the buckling load converges to a value of 278.59kN. The corresponding buckling mode is shown in Fig. 6 (left). Fig. 6 (middle) shows the convergence of the relative error in λ, which converges at a rate of $\alpha \approx 1$ with respect to the number of degrees of freedom M, i.e.,

$$\left|1 - \frac{\lambda_M}{\lambda}\right| \approx M^{-1}$$

and agrees with the theoretically predicted convergence rate for buckling modes for this element. The scaling of the cost (CPU-time in seconds) with M is shown in Fig. 6 (right). The gradient of the line shows that

$$C(\lambda_M) \approx M^{1.17},$$

i.e., a value of $\gamma \approx 1.17$. All computations were carried out using MATLAB's inbuilt sparse eigensolver `eigs(. . .)`, which is a wrapper to the widely used package ARPACK [47], which implements an implicitly restarted Arnoldi method and exploits UMFPACK [48] as the default direct solver. For this implementation the computational cost is made up of matrix assembly for (45) and of the calculation of the smallest eigenvalue of (45). For the size of 2D problems considered here ($M \leq 10^6$), the CPU-time is dominated by the matrix assembly, which scales linearly with M. For larger problem sizes ($M > 10^6$), the eigenvalue solve dominates the CPU-time and the value of γ can be expected to increase.

In this example, the aim is to estimate the probability $\mathbb{P}(\lambda < \lambda^\star = 272.47$ kN), which corresponds to estimating the mean value of the quantity of interest

$$Q(\lambda) = \begin{cases} 1 & \text{if} \quad \lambda < \lambda^\star, \\ 0 & \text{if} \quad \lambda \geq \lambda^\star. \end{cases} \tag{46}$$

The coarsest mesh ($\ell = 0$) consists of 32×32 elements, created by five uniform refinements of a single rectangular element. This level of refinement is necessary to ensure that some failures are observed on the coarsest level and the bias error estimate is reliable, since buckling load only reduces with mesh refinement due to the min–max principle [36].

Fig. 7 (top) shows again the behaviour of the expected value and of the variance of Q_ℓ and of Y_ℓ, with respect to the degrees of freedom M_ℓ. From this, approximate rates $\alpha \approx \beta \approx 1$ can be estimate, such that

$$\mathbb{E}[Y_\ell] \approx M_\ell^{-1} \quad \text{and} \quad \mathbb{V}[Y_\ell] \approx M_\ell^{-1}. \tag{47}$$

The bottom two plots in Fig. 7 are related to computational cost. The lower-left plot compares the expected cost per sample for basic MLMC

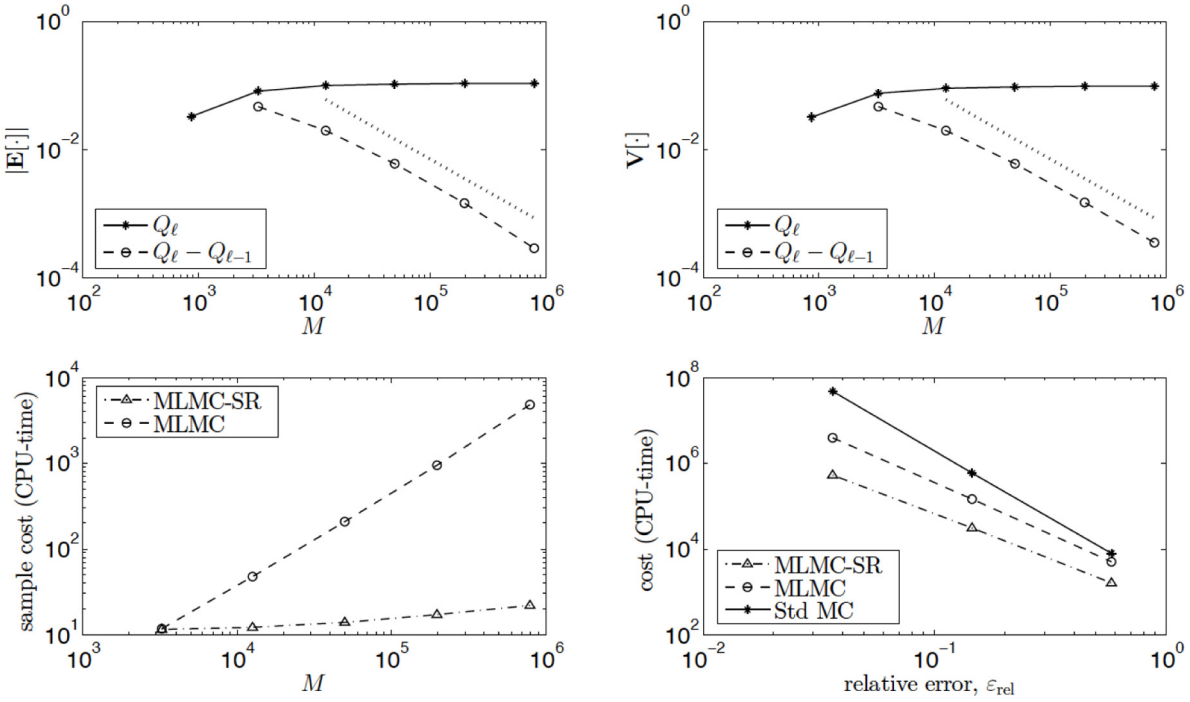

Fig. 7. (Top-left and top-right) Expected values and variances of Q_ℓ and Y_ℓ for Example II, plotted against degrees of freedom M, respectively. The dotted lines have gradient -1. (Bottom-Left) Comparison of expected cost per sample (CPU-time in sec) for MLMC and MLMC-SR; the observed gradients of the curves are 1.17 and 0.12, respectively. (Bottom-Right) Comparison of cost (CPU-time in sec) for MLMC-SR, MLMC and standard MC, against relative error tolerance ϵ_{rel}; the observed gradients of the curves are -2.03, -2.28 and -3.14, respectively.

and for MLMC with selective refinement (MLMC-SR). While the rate of growth for MLMC is $\gamma \approx 1.17$, as estimated above, using (24) the growth of the cost per sample for MLMC-SR is expected to grow proportionally to $M_\ell^{\gamma-\alpha} = M_\ell^{0.17}$, which is a huge gain. The numerical results confirm this clearly. The numerically observed growth is proportional to $M_\ell^{0.12}$.

The lower-right plot shows the computational cost of the MLMC-SR simulation versus that of basic MLMC and standard (single-level) MC, for a range of relative tolerances $\epsilon_{\text{rel}} \geq 3.6\%$. In this regime, clearly $\alpha < \gamma < 2\alpha$, for which the cost of the MLMC-SR simulation is predicted to grow proportionally to $\epsilon_{\text{rel}}^{-2}$ in [29]. This agrees exactly with the observed rate. The costs of the MLMC and standard MC simulations, on the other hand, are predicted to grow like $\epsilon_{\text{rel}}^{-2.13}$ and $\epsilon_{\text{rel}}^{-3.13}$, respectively, and again there is good agreement with the numerically observed rates. Table 2 lists the optimal numbers of samples N_ℓ across the levels as given by (14) for each of the methods to achieve the required error tolerances. The total computational costs of the simulations are also included. For the smallest error tolerance considered, 3.6%, the MLMC-SR simulation reduces the computational cost by a factor of 7.43 compared to the basic MLMC simulation, and by a factor of 90.32 compared to standard MC. The error is again split equally between bias and sampling error ($\theta = 1/2$).

It should be noted in Table 2 that MLMC-SR leaves more of the samples on the finer levels than MLMC, but the samples are of course much cheaper. In Algorithm 2, failure for each sample is computed adaptively, always starting on the coarsest level and then refining the mesh level by level. For the majority of samples, failure can already be conclusively detected or rejected on a coarse mesh, thus significantly reducing the average cost per sample.

5.3. Estimation of rare events

Finally, this section considers a much smaller failure probability of about 1/150. This will push the computational demand well out of the reach of standard Monte Carlo, and demonstrate the potential computational benefits of adopting the multilevel strategies for the estimation of rare events, especially by exploiting the selective refinement strategy. The same setup as above is used in this section, but with the much rarer failure load of $\lambda^\star = 268$ kN. In that case

$$\mathbb{P}(\lambda \leq \lambda^\star = 268 \text{ kN}) = 0.00645 \approx 1/150.$$

However, in this case the standard multilevel approach gives rise to somewhat of a paradox. Its success is based on the telescoping sum (9), which requires estimating the differences $\mathbb{E}[Y_\ell]$. However, in the stochastic eigenvalue problem $Y_\ell = 1$ only if failure occurs on level ℓ but not on level $\ell - 1$. Of course, this conditional probability is a much rarer event than just failure occurring on level ℓ. Thus, for the higher levels, many more simulations are necessary to see just one case where the two adjacent levels differ and thus $Y_\ell = 1$. However, with the use of selective refinement most samples are pre-screened by coarser/cheaper model solves; so only very rarely, when there is a discrepancy at higher levels, expensive solves are required. Since failure is rare, most samples of λ are sufficiently far away from $\lambda^\star$ so that (23) is satisfied for a very coarse model. For example, in the calculations below the average computational cost of 0.239 s to compute a sample on level 2 is little different to the 0.244 s to compute a sample on level 5.

This highlights that significant computational gains for rare events can be achieved by using MLMC-SR with just two levels, i.e., using the estimator

$$\mathbb{E}[Q] \approx \hat{Q}_0 + \hat{Y}_{L,0}, \quad \text{where} \quad Y_{L,0} = Q_L - Q_0.$$

Note however, that all levels of refinement from 0 to L are still used to calculate Q_L via the selective refinement procedure. Because of the plateau in cost for MLMC-SR for rare events, it is less efficient to use all levels in the estimator, yet the two level method still provides significant variance reduction in the results to follow. However, most of the computational gains come from the selective refinement approach.

Because of the scale of the calculations, these simulations were distributed over 1,024 cores of `Isca`, University of Exeter's supercomputer with about 400 nodes, each with 8-core Intel Xeon E5-2650v2 Ivybridge processors running at 2.6 GHz. The computational savings

Table 2
Comparison of optimal number of samples N_ℓ and computational cost for MLMC-SR, MLMC, and standard MC simulations for Example II.

ϵ_{rel}	Method	N_ℓ						Cost (h)	Saving
		0	1	2	3	4	5		
3.6%	MLMC-SR	26,883	16,489	10,348	5345	2407	1,029	147	–
	MLMC	73,226	44,365	14,382	3754	879	189	1,097	7.43
	MC	–	–	–	–	–	11,897	13,341	90.32
15%	MLMC-SR	1,610	988	620	320	145	–	8.50	–
	MLMC	3,523	2,135	692	181	43	–	41	4.81
	MC	–	–	–	–	744	–	162	19.24

Table 3
Number of solves on each level and total computational cost to estimate a rare event in Example II, demonstrating the huge relative savings of MLMC-SR over standard MC for a range of tolerances. The computational cost is CPU time on 1,024 cores of the University of Exeter's supercomputer Isca. (⋆ indicates that this calculation was only estimated due to the excessive computational cost.).

ϵ_{rel}	Method	Term	Solves on level					Cost	Saving Factor
			0	1	2	3	4		
4.3%	MLMC-SR	$\hat{Q}_0$	3.65e5	-	–	–	-	35.7 s	69
		$\hat{Y}_{2,0}$	2.54e5	2.53e5	348	–	–		
	MC	$\hat{Q}_2$	–	-	3.18e5	–	-	41 min	–
1.4%	MLMC-SR	$\hat{Q}_0$	3.32e6	–	-	–	-	5.46 min	124
		$\hat{Y}_{3,0}$	2.41e6	2.40e6	4,268	965	–		
	MC	$\hat{Q}_3$	–	-	–	2.93e6	-	11.24 h	–
0.2%	MLMC-SR	$\hat{Q}_0$	1.63e8	–	-	–	-	4.4 h	1173
		$\hat{Y}_{4,0}$	1.19e8	1.18e8	2.19e5	5.61e4	8,348		
	MC⋆	$\hat{Q}_4$	–	-	–	–	1.44e8	218 days⋆	–

of MLMC-SR over standard MC (without selective refinement) over a range of tolerances are summarised in Table 3. In each case, bias and sampling error are balanced (i.e $\theta = 1/2$). The results show huge computational savings across the range of tolerances. Especially for the finest tolerance of 0.2%, the estimated savings are such that MLMC-SR would be 1173 times faster than standard MC. More importantly, the scale of computations is such that standard MC would take 218 days of computation, even using 1,024 cores on a supercomputer, while MLMC-SR reduces this to a very reasonable 4.4 h. From an engineering perspective, this scale of savings opens the opportunity to new studies of rare events.

6. Conclusions

In this paper, multilevel Monte Carlo (MLMC) simulation has been successfully demonstrated on two typical aerospace example problems. From the numerical results, the advantages of MLMC over standard MC are apparent, with huge computational savings being observed.

MLMC simulation is not limited to easy problems, and in fact the gains are more pronounced in cases where the discretisation error is large. A particular aim was to demonstrate the versatility of the approach, showing that the method is not restricted to problems in which the quantity of interest is a smooth functional of the PDE solution, but that it can readily be applied to the estimation of failure probabilities with significant computational speed-ups.

From an engineering viewpoint, whilst the example problems are chosen to showcase the typical gains achieved with MLMC, the physical insight and the engineering implications of parameter uncertainty are of independent interest in both cases. In the buckling test problem, the numerical results show that random variations in ply angles, perhaps unsurprisingly, increase the risk of buckling failure significantly. With ply angles of the order typically observed in an Automated Fibre Placement (AFP) machine ($\pm 5°$) significant variability is observed in buckling performance. As for our numerical results on the effect of random fibre waviness on the compressive strength of composites, high fidelity stochastic simulations show a remarkable agreement with Budiansky's classical kinking model [32] if the misalignment angle is taken to be the standard deviation of the misalignment random field.

Current and future research is exploring the use of sample-dependent adaptive grids, to exploit the computational gains offered by adaptive finite elements [35]; as well as integrating the multilevel framework with experimental data in a Bayesian setting to quantify and reduce modelling uncertainties as proposed by theoretical methodology introduced in [31]. From an engineering application perspective the MLMC methods set out in this paper are being applied to nonlinear analysis where the quantities of interest include non-smooth functionals. One such application is the classical study of imperfection sensitivity of thin cylindrical shells under compression. Since the models for nonlinear analysis are orders of magnitude more expensive the development of methods which optimise sampling like MLMC are essential.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by the EPSRC, United Kingdom 'Multiscale Modelling of Aerospace Composites' project (EP/K031368/1), Dodwell is supported by a Turing AI, United Kingdom fellowship (EP/N510129/1) and Butler held a Royal Academy of Engineering-GKN Aerospace Research Chair in Composites, United Kingdom.

References

[1] US Department of Transportation, Composite Aircraft Structure, Advisory Circular, 2010, 20–107B.
[2] E. Acar, R.T. Haftka, Reliability-based aircraft structural design pays, even with limited statistical data, J. Aircr. 44 (3) (2007) 812–823.
[3] M.P.F. Sutcliff, S.L. Lemanski, A.E. Scott, Measurement of fibre waviness in industrial composite components, Compos. Sci. Technol. 72 (2012) 2016–2023.

[4] A.T. Rhead, T.J. Dodwell, R. Butler, The effect of tow gaps on compression after impact strength of robotically laminated structures, Comput. Mater. Contin. 35 (1) (2013) 1–16.
[5] T.J. Dodwell, R. Butler, G.W. Hunt, Out-of-plane ply wrinkling defects during consolidation over an external radius, Compos. Sci. Technol. 105 (2014) 151–159.
[6] T.A. Fletcher, R. Butler, T.J. Dodwell, Anti-symmetric laminates for improved consolidation and reduced warp of tapered C-sections, Adv. Manuf.: Polym. Compos. Sci. 1 (2) (2015) 105–111.
[7] J. Belnoue, O. Nixon-Pearson, A. Thompson, D. Ivanov, K. Potter, S.R. Hallett, Consolidation-driven defect generation in thick composite parts, J. Manuf. Sci. Eng. 140 (2018).
[8] A. Sandhu, A. Reinarz, T.J. Dodwell, A Bayesian framework for assessing the strength distribution of composite structures with random defects, Compos. Struct. (2018).
[9] R.E. Melchers, Structural Reliability Analysis and Prediction, second ed., John-Wiley, 1999.
[10] I. Elishakoff, S. van Manent, P.G. Vermeulent, J. Arbocz, First-order second-moment analysis of the buckling of shells with random imperfections, AIAA J. 25 (8) (1987) 1113–1117.
[11] A. Shaw, S. Sriramula, P.D. Gosling, M.K. Chryssanthopoulos, A critical reliability evaluation of fibre reinforced composite materials based on probabilistic micro and macro-mechanical analysis, Composites B 41 (2010) 446–453.
[12] C.Y. Park, N.H. Kim, R.T. Haftka, The effect of ignoring dependence between failure modes on evaluating system reliability, Struct. Multidiscip. Optim. 52 (2015) 251–268.
[13] I. Elishakoff, Notes on philosophy of the Monte Carlo method, Int. Appl. Mech. 39 (7) (2003) 753–762.
[14] B.P. Smarslok, R.T. Haftka, L. Carraro, D. Ginsbourger, Improving accuracy of failure probability estimates with separable Monte Carlo, Int. J. Reliab. Saf. 4 (2010) 393–414.
[15] A. Chaudhuri, R.T. Haftka, Separable Monte Carlo combined with importance sampling for variance reduction, Int. J. Reliab. Saf. 7 (3) (2013) 201–215.
[16] R. Caflisch, Monte Carlo and quasi-Monte Carlo methods, Acta Numer. 7 (1998) 1–49.
[17] J. Dick, F.Y. Kuo, I.H. Sloan, High-dimensional integration: The quasi-Monte Carlo way, Acta Numer. 22 (2013) 133–288.
[18] D. Allaire, K. Willcox, Surrogate modeling for uncertainty assessment with application to aviation environmental system models, AIAA J. 48 (8) (2010) 1791–1803.
[19] C.E. Rasmussen, C.K.I. Williams, Gaussian Processes for Machine Learning, MIT Press, 2006.
[20] P.D. Spanos, R.G. Ghanem, Stochastic Finite Elements: A Spectral Approach, Courier Corporation, 1991.
[21] M.B. Giles, Multilevel Monte Carlo methods, Acta Numer. 24 (2015) 259–328.
[22] S. Heinrich, Multilevel Monte Carlo methods, in: International Conference on Large-Scale Scientific Computing, Springer, Berlin, Heidelberg, 2001, pp. 58–67.
[23] M.B. Giles, Multilevel Monte Carlo path simulation, Oper. Res. 56 (3) (2008) 981–986.
[24] K.A. Cliffe, M.B. Giles, R. Scheichl, A.L. Teckentrup, Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients, Comput. Vis. Sci. 14 (1) (2011) 3–15.
[25] A. Barth, Ch. Schwab, N. Zollinger, Multi-level Monte Carlo finite element method for elliptic PDE's with stochastic coefficients, Numer. Math. 119 (2011) 123–161.
[26] F. Müller, P. Jenny, D.W. Meyer, Multilevel Monte Carlo for two phase flow and Buckley–Leverett transport in random heterogeneous porous media, J. Comput. Phys. 250 (2013) 685–702.
[27] S. Mishra, C. Schwab, J. Sukys, Multi-level Monte Carlo finite volume methods for shallow water equations with uncertain topography in multi-dimensions, SIAM J. Sci. Comput. 34 (2012) 761–784.
[28] N. Collier, A.L. Haji-Ali, F. Nobile, E. vo. Schwerin, R. Tempone, A continuation multilevel Monte Carlo algorithm, Bit Numer. Math. 55 (2015) 399–432.
[29] D. Elfverson, F. Heilman, A. Malqvist, A multilevel Monte Carlo method for computing failure probabilities, SIAM/ASA J. Uncertain. Quantif. 4 (1) (2016) 312–330.
[30] V.H. Hoang, Ch. Schwab, A.M. Stuart, Complexity analysis of accelerated MCMC methods for Bayesian inversion, Inverse Problems 29 (2013) 085010.
[31] T.J. Dodwell, C. Ketelsen, R. Scheichl, A.L. Teckentrup, A hierarchical multilevel Markov chain Monte Carlo algorithm with applications to uncertainty quantification in subsurface flow, SIAM/ASA J. Uncertain. Quantif. 3 (1) (2015) 1075–1108.
[32] B. Budiansky, Micromechanics, Comput. Struct. 16 (1983) 3–12.
[33] R. Butler, T.J. Dodwell, R.T. Haftka, Nam H. Kim, T. Kim, S. Kynaston, R. Scheichl, Uncertainty quantification of composite structures with defects using multilevel Monte Carlo simulations, in: 17th AIAA Non-Deterministic Approaches Conference, AIAA 2015-1598 Kissimmee, Florida, 2015.
[34] A.L. Teckentrup, R. Scheichl, M.B. Gile,s, E. Ullmann, Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients, Numer. Math. 125 (3) (2013) 569–600.
[35] G. Detommaso, T.J. Dodwell, R. Scheichl, Continuous level Monte Carlo and sample-adaptive model hierarchies, SIAM/ASA J. Uncertain. Quantif. 7 (1) (2019) 93–116.
[36] B.N. Parlett, The Symmetric Eigenvalue Problem, SIAM, New Jersey, 1980.
[37] D. Liu, N.A. Fleck, M.P.F. Sutcliffe, Compressive strength of fibre composites with random fibre waviness, J. Mech. Phys. Solids 52 (2004) 1481–1505.
[38] P.D. Spanos, B.A. Zeldin, Monte Carlo treatment of random fields: a broad perspective, Appl. Mech. Rev. 51 (3) (1998) 219–237.
[39] C. Scarth, S. Adhikari, P.H. Cabral, G.H.C. Silva, A. Pereira do Prado, Random field simulation over curved surfaces: Applications to computational structural mechanics, Comput. Methods Appl. Mech. Engrg. 345 (2019) 283–301.
[40] T.J. Dodwell, Internal wrinkling instabilities in layered media, Phil. Mag. 95 (2015) 3225–3243.
[41] Hexcel Composites, Hexply 8552 Epoxy Matrix Product Data, Technical report, 2008.
[42] P.M. Jelf, N.A. Fleck, Compression failure mechanisms in unidirectional composites, J. Compos. Mater. 26 (1992) 2706–2762.
[43] M.R. Wisnom, The effect of fibre waviness on the relationship between compressive strength of unidirectional composites, J. Compos. Mater. 28 (1994) 66–76.
[44] T.J.R. Hughes, The Finite Element Method: Linear Static and Dynamic Finite Element Analysis, Dover, Mineola, New York, 2000.
[45] Z. Gurdal, R.T. Haftka, P. Hajela, Design and Optimisation of Laminated Composite Materials, Wiley, 1999.
[46] T.J. Dodwell, R. Butler, A.T. Rhead, Optimum fiber steering of composite plates for buckling and manufacturability, AIAA J. 54 (3) (2016) 1146–1149.
[47] R.B. Lehoucq, D.C. Sorensen, C. Yang, ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods, SIAM Publications, Philadelphia, 1998.
[48] T. Davis, Algorithm 832:UMFPACK V4.3 - an unsymmetric-pattern multifrontal method, ACM Trans. Math. Softw. 30 (2) (2004) 196–199.