

Computational Physics

Multilevel Monte Carlo methods for the Grad-Shafranov free boundary problem[☆]Howard C. Elman^a, Jiaxing Liang^b, Tonatiuh Sánchez-Vizuet^{c,*}^a Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, United States of America^b Applied Mathematics & Statistics and Scientific Computing Program, University of Maryland, College Park, United States of America^c Department of Mathematics, The University of Arizona, United States of America

ARTICLE INFO

Keywords:

Multilevel Monte Carlo Finite-Element
 Uncertainty quantification
 Free boundary Grad-Shafranov problem
 Adaptive finite element discretization

ABSTRACT

The equilibrium configuration of a plasma in an axially symmetric reactor is described mathematically by a free boundary problem associated with the celebrated Grad-Shafranov equation. The presence of uncertainty in the model parameters introduces the need to quantify the variability in the predictions. This is often done by computing a large number of model solutions on a computational grid for an ensemble of parameter values and then obtaining estimates for the statistical properties of solutions. In this study, we explore the savings that can be obtained using multilevel Monte Carlo methods, which reduce costs by performing the bulk of the computations on a sequence of spatial grids that are coarser than the one that would typically be used for a simple Monte Carlo simulation. We examine this approach using both a set of uniformly refined grids and a set of adaptively refined grids guided by a discrete error estimator. Numerical experiments show that multilevel methods dramatically reduce the cost of simulation, with cost reductions typically on the order of 60 or more and possibly as large as 200. Adaptive gridding results in more accurate computation of geometric quantities such as x-points associated with the model.

1. Introduction

Monte Carlo (MC) techniques are one of the most common strategies for dealing with the quantitative assessment of the accuracy of numerical simulations of physical models with uncertainties. The idea behind these methods is to obtain a large number of samples (typically by numerically solving the associated model) for random realizations of the uncertain parameters, and use these data to gather statistical information about the quantity of interest. However, when the model involves the solution of partial differential equations, the computational effort related to the collection of the data points required can easily become unmanageable. To overcome this difficulty methods like polynomial chaos expansions [46], stochastic Galerkin [18], and stochastic collocation [3] have been used to handle uncertainties associated with a small number of parameters. These techniques, however, often require the development of specialized numerical solvers or rely on the smooth dependence of the model with respect to the parameter values. The multilevel Monte Carlo (MLMC) method was developed [8,19,41,44] as an efficient alternative that does not require additional smoothness

assumptions and can take advantage of an existing numerical solver. Given a target numerical grid (i.e. a grid whose resolution is considered sufficiently fine) MLMC improves the efficiency of the sampling step by offsetting the bulk of the numerical computations to a sequence of coarser grids where the numerical solution is cheaper.

In the particular context of free-boundary Grad-Shafranov computations subject to parameter uncertainty, the authors have shown that the computational cost can be reduced manyfold by employing a strategy based on stochastic collocation [12]; however, due to the latent possibility of plasma-wall contacts, the smoothness of the mapping between coil currents and equilibria cannot be guaranteed. In this paper, our goal is to overcome this difficulty by approximating the expectation of the equilibrium configuration using a multilevel Monte Carlo Finite-Element (MLMC-FE) approach. We will consider two MLMC-FE algorithms: a classical strategy based on uniformly refined meshes and a variation based on meshes refined adaptively using an *a posteriori* error estimator. As we shall see, both of these approaches greatly reduce computational costs, with the adaptive strategy being somewhat more

[☆] The review of this paper was arranged by Prof. David W. Walker.^{*} Corresponding author.E-mail addresses: helman@umd.edu (H.C. Elman), jliang18@umd.edu (J. Liang), tonatiuh@math.arizona.edu (T. Sánchez-Vizuet).




Fig. 1. The plasma confinement region Ω_p is enclosed by the violet line. The rectangles represent the external coils C_k ; the gray curved line represents the exterior wall of the vacuum chamber. See online version for color images.

effective in computing the approximation of the expectation of geometric properties of the equilibrium configuration.

An outline of the paper is as follows. In Section 2, we briefly recall the Grad-Shafranov free boundary problem. Section 3 is devoted to the introduction of the Monte Carlo and multilevel Monte Carlo Finite-Element methods. The section concludes with the introduction of an algorithm to compute the optimal number of samples required at each discretization level. In Section 4 we present numerical experiments comparing the effectiveness of these strategies with a Monte Carlo strategy based on a single mesh. Concluding remarks are presented in the final Section 5. For completeness, the technical mathematical, and algorithmic aspects of the problem and the methods discussed are included as an appendix.

2. The Grad-Shafranov free boundary problem

2.1. The deterministic problem

In a cylindrically symmetric magnetic confinement device, with coordinates (r, z, φ) , the mathematical expression of the equilibrium condition between the hydrostatic and magnetic forces acting on the plasma results in the celebrated Grad-Shafranov equation [21,36,42]. This nonlinear elliptic equation relates the *poloidal flux function* $\psi(r, z)$ to the hydrostatic pressure $p(\psi)$ and the *toroidal field* $g(\psi)$ (both of which are assumed to be functions of ψ only), and the currents I_k going through external coils with cross-sectional area S_k . Posed in free space, the equation takes the form

$$-\nabla \cdot \left(\frac{1}{\mu r} \nabla \psi \right) = \begin{cases} r \frac{d}{d\psi} p(\psi) + \frac{1}{2\mu r} \frac{d}{d\psi} g^2(\psi) & \text{in } \Omega_p(\psi) \\ I_k / S_k & \text{in } \Omega_{C_k} \\ 0 & \text{elsewhere.} \end{cases} \quad (1a)$$

Above, μ is the magnetic permeability, Ω_{C_k} denotes the area occupied by the k -th external coil, Ω_p is the plasma confinement region which is *not known a priori* and must be determined as a problem unknown, making this a *free boundary problem*. A schematic of a cross-section, for $r > 0$, of a tokamak is depicted in Fig. 1. The confinement region Ω_p is characterized as the largest region that contains the magnetic axis (defined as the point where ψ has a global maximum) and that is bounded by a closed level set $\psi = \text{constant}$. The solution to this free boundary problem is ubiquitous in nuclear fusion and several computational codes have been developed over the years (see for instance [14,20,24,25,28,31] and references therein).

A common choice in the literature, first proposed in [37], for the free functions $p(\psi)$ and $g(\psi)$ in the right hand side of (1a) is

$$\frac{d}{d\psi} p(\psi) = j_0 \frac{\beta}{r_0} (1 - \psi_N^{\alpha_1})^{\alpha_2}, \quad \frac{1}{2} \frac{d}{d\psi} g^2(\psi) = j_0 \mu_0 r_0 (1 - \beta) (1 - \psi_N^{\alpha_1})^{\alpha_2}, \quad (1b)$$

where r_0 represents the outer radius of the vacuum chamber and $\psi_N \in [0, 1]$ is a normalization of the flux ψ such that $\psi_N = 1$ on the plasma boundary $\partial\Omega_p$ and $\psi_N = 0$ at the magnetic axis. The parameters α_1 and α_2 control the behavior of ψ around the magnetic axis, and β measures the ratio between the hydrostatic pressure and the magnetic pressure in the plasma, and j_0 is a normalization factor.

2.2. Incorporating uncertainty

In this article, we will consider that the uncertainty in the model (1a) is concentrated in the values of the currents I_k going through the external coils. As a result, the function ψ (and all quantities derived from it) are random variables. Obtaining a full description of their probability density functions might not be possible, but an approximate picture can be obtained by exploring the parameter space and computing sample approximations of its expectation and variance. We will model the array of currents as a d -dimensional random variable $\omega := (\omega_1, \dots, \omega_d)$, where d is the number of confinement coils in the reactor, and the k -th component of ω is the current going through the k -th coil. We will consider that ω is uniformly distributed around a baseline vector $I = (I_1, \dots, I_d)$ corresponding to the desired current values in a deterministic model. We will often refer to I as either the *reference* or *unperturbed* currents. Letting $\tau > 0$ denote the size of the possible perturbation in the current values (relative to the components of I), the vector ω is then uniformly distributed over the d -dimensional parameter space

$$W := \prod_{k=1}^d [I_k - \tau |I_k|, I_k + \tau |I_k|]. \quad (2)$$

Since coils are independent of each other, the stochastic random currents $\{\omega_k\}_{k=1}^d$ are uncorrelated and the joint density function of ω is given by $\pi(\omega) = \prod_{k=1}^d \pi_k(\omega_k) = \prod_{k=1}^d \frac{1}{2\tau |I_k|}$. The equilibrium configuration determined by the solution to (1a) is then the random variable $\psi(r, z, \omega)$; we will be primarily interested in efficiently constructing an approximation to its expected value

$$\mathbb{E}[\psi(r, z, \omega)] = \int_W \psi(r, z, \omega) \pi(\omega) d\omega, \quad (3)$$

as well as those of some derived quantities such as the plasma boundary, the location of the x-points, etc.

3. Monte Carlo and multilevel Monte Carlo Finite-Element methods

We now turn our attention to the numerical approximation of the expected value (3). Since the location and shape of the plasma boundary depend on the values of the coil currents, variations of these values could lead to contacts between the plasma and the wall or even loss of confinement. This fact translates into a possible non-smoothness of the mapping between coil currents and the solutions of (1a) which may then cause techniques such as stochastic collocation to underperform. Moreover, the computational effort associated with cubature methods scales exponentially with the dimension of the parameter space, seriously limiting their feasibility for estimating (3). This leads to the use of Monte Carlo methods, which are agnostic to both the smoothness of the mapping and the dimensionality of the problem [38], although they have a slow convergence rate (1/2) that tends to make them costly. This will be addressed through the use of a multi-level approach.

3.1. Monte Carlo Finite-Element method

We will describe the method in terms of a generic solution, u , to a PDE involving stochastic parameters and its finite element approximation u_h , where h is the mesh parameter of the discretization. Let

$\{\omega^{(i)}\}_{1 \leq i \leq N}$ be a set of N realizations of the random variable ω giving rise to a sample of N realizations $u^{(i)} = u(\omega^{(i)})$ and $u_h^{(i)} = u_h(\omega^{(i)})$ of the exact solution and its finite element discretization. We will assume that all these functions belong to a functional space Z endowed with a norm $\|\cdot\|_Z$ (see the Appendix 7.1 for details), and we will consider the standard FEM error estimate $\|u^{(i)} - u_h^{(i)}\|_Z \leq C^{(i)} h^p$, where p is the order of the FEM discretization and the constant $C^{(i)}$ depends only on the problem geometry and the particular values of $\omega^{(i)}$.

The Monte Carlo Finite-Element (MC-FE) estimator $A_{MC}(u_h)$ for $\mathbb{E}(u)$ is defined as the sample mean

$$A_{MC}(u_h) := \frac{1}{N} \sum_{i=1}^N u_h^{(i)}. \quad (4)$$

This estimator is easily shown to be unbiased and to satisfy $\mathbb{E}(A_{MC}) = \mathbb{E}(u_h)$. A quantity that serves as a foundation for examining the spatial and statistical accuracy of the MC-FE estimator is the *mean squared error* (MSE) defined as

$$\mathcal{E}_{A_{MC}}^2 := \mathbb{E} \left[\|\mathbb{E}(u) - A_{MC}(u_h)\|_Z^2 \right].$$

It can be shown (see, for instance [4, Theorem 4.3]) that, for linear problems, the Monte Carlo estimator accurately approximates the expected value in the sense that $\mathcal{E}_{A_{MC}}^2 \leq K(N^{-1/2} + h^p)^2$, where the constant $K > 0$ depends on the problem geometry and the expected values of the problem data. The MSE can be decomposed into terms related to the bias and variance as

$$\begin{aligned} \mathcal{E}_{A_{MC}}^2 &= \|\mathbb{E}(u) - \mathbb{E}(u_h)\|_Z^2 + \mathbb{E} \left[\|\mathbb{E}(u_h) - A_{MC}(u_h)\|_Z^2 \right] \\ &= \|\mathbb{E}(u) - \mathbb{E}(u_h)\|_Z^2 + \frac{\mathbb{V}(u_h)}{N} = \mathcal{E}_{\text{Bias}}^2 + \mathcal{E}_{\text{Stat}}^2, \end{aligned}$$

where $\mathbb{V}(u) := \mathbb{E}[\|u - \mathbb{E}(u)\|_Z^2]$ and $\mathbb{V}(A_{MC}(u)) = \mathbb{V}(u)/N$. The last two terms in the expression above implicitly define the discretization error $\mathcal{E}_{\text{Bias}}$ and the sampling (or statistical) error $\mathcal{E}_{\text{Stat}}$ respectively.

If $\mathbb{E}(u) \neq 0$, the mean squared error can be expressed as a percentage through normalization by the factor $\|\mathbb{E}(u)\|_Z^2$, leading to the *normalized mean squared error* $\hat{\mathcal{E}}_{A_{MC}}^2$. Since the exact random variable u is not available, we will approximate the relative mean squared error (nMSE) by

$$\hat{\mathcal{E}}_{A_{MC}}^2 \approx \frac{\|\mathbb{E}(u) - \mathbb{E}(u_h)\|_Z^2}{\|\mathbb{E}(u_h)\|_Z^2} + \frac{\mathbb{V}(u_h)}{N \|\mathbb{E}(u_h)\|_Z^2} = \hat{\mathcal{E}}_{\text{Bias}}^2 + \hat{\mathcal{E}}_{\text{Stat}}^2, \quad (5)$$

where $\hat{\mathcal{E}}_{\text{Bias}}$ and $\hat{\mathcal{E}}_{\text{Stat}}$ are relative analogues to the discretization and statistical errors defined above. If the number of grid points for the FEM discretization is M then, in two dimensions, it is standard to assume that $\hat{\mathcal{E}}_{\text{Bias}} = \mathcal{O}(M^{-p/2})$. Given a target tolerance ϵ , the contribution of the statistical and discretization errors towards the total nMSE can be controlled by requiring that

$$\hat{\mathcal{E}}_{\text{Bias}}^2 = \mathcal{O}(M^{-p}) \leq (1 - \theta)\epsilon^2, \quad \hat{\mathcal{E}}_{\text{Stat}}^2 = \frac{V_h}{N} \leq \theta\epsilon^2, \quad (6)$$

where $\theta \in (0, 1)$ is known as the *splitting parameter*, and $V_h := \mathbb{V}(u_h) / \|\mathbb{E}(u_h)\|_Z^2$. This in turn allows us to estimate the sample size N and the number of grid points M required to achieve the desired tolerance as

$$M \geq ((1 - \theta)\epsilon^2)^{-\frac{1}{p}}, \quad N = \left\lceil \frac{V_h}{\theta\epsilon^2} \right\rceil. \quad (7)$$

Assuming the average cost to obtain one sample (i.e. to solve (1a) for one value of the coil-currents) is $C = \mathcal{O}(M^c)$ for some $c > 0$, the total computational cost of the MC-FE estimator can be estimated as

$$C(A_{MC}) = \mathcal{O}(NM^c) = \mathcal{O}\left(\epsilon^{-2-\frac{2c}{p}}\right). \quad (8)$$

3.2. Multilevel Monte Carlo Finite-Element method

MLMC reduces the computational cost associated with sampling—which in our case involves the numerical solution of a non-linear PDE in a target computational mesh—by approximating the expectation of the quantity of interest on the finest mesh by a sequence of control variates on a set of coarse grids that are cheaper to compute [13]. Using the linearity of expectation, the MLMC estimator expresses the quantity of interest on the finest spatial grid, $\mathbb{E}(u_h)$, by a telescoping sum involving the numerical approximations of u on coarser grids. Consequently, MLMC's workload is shifted from the fine grid to coarser grids, making it more efficient than MC [8]. Cf. e.g. [35] for alternative ways to reduce the costs of MC methods.

To construct meshes that are easy to describe for both uniform and adaptive mesh refinement, we will characterize them using the number of grid points rather than the mesh size. We will refer to $\ell = 0, \dots, L$ as the *level* of a mesh $\mathcal{T}_\ell$ containing $\{M_\ell\}$ grid points. We will then consider a sequence of meshes $\mathcal{T}_0, \dots, \mathcal{T}_L$ with increasing resolution so that $\{M_\ell\}_{0 \leq \ell \leq L}$ defines an increasing sequence and $\mathcal{T}_L$ is the finest mesh, and we will denote by u_ℓ the approximation of u on the mesh $\mathcal{T}_\ell$.

The expectation of the function u can be approximated by the expectation of the finest approximation u_L . Since $u_L = u_0 + (u_1 - u_0) + (u_2 - u_1) + \dots + (u_L - u_{L-1})$, this approximation is given by

$$\mathbb{E}(u) \approx \mathbb{E}(u_L) = \mathbb{E}(u_0) + \sum_{\ell=1}^L \mathbb{E}(u_\ell - u_{\ell-1}) = \sum_{\ell=0}^L \mathbb{E}(Y_\ell), \quad (9)$$

a telescoping sum, where each of the terms

$$Y_0 := u_0 \quad \text{and} \quad Y_\ell := u_\ell - u_{\ell-1} \quad (\text{for } \ell \geq 1) \quad (10)$$

can be regarded as a correction of the coarsest approximation u_0 . If each term $\mathbb{E}(Y_\ell)$ is estimated by gathering N_ℓ samples at level ℓ and computing the sample expectations

$$\begin{aligned} \mathbb{E}(Y_0) \approx \hat{Y}_0 &:= \frac{1}{N_0} \sum_{i=1}^{N_0} u_0^{(i)}, & \mathbb{E}(Y_\ell) \approx \hat{Y}_\ell &:= \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} (u_\ell^{(i)} - u_{\ell-1}^{(i)}) \\ &(\text{for } \ell \geq 1), \end{aligned}$$

then the MLMC-FE estimator at level L will be unbiased and can be written as

$$A_{\text{MLMC}}(u_L) := \sum_{\ell=0}^L \hat{Y}_\ell = \frac{1}{N_0} \sum_{i=1}^{N_0} u_0^{(i)} + \sum_{\ell=1}^L \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} (u_\ell^{(i)} - u_{\ell-1}^{(i)}). \quad (11)$$

Recalling that $\mathbb{E}(\hat{Y}_\ell) = \mathbb{E}(Y_\ell)$ and $\mathbb{V}(\hat{Y}_\ell) = \mathbb{V}(Y_\ell)/N_\ell$ we conclude that, for the MLMC-FE estimator, it follows that $\mathbb{E}(A_{\text{MLMC}}) = \sum_{\ell=0}^L \mathbb{E}(\hat{Y}_\ell) = \mathbb{E}(u_L)$ and $\mathbb{V}(A_{\text{MLMC}}) = \sum_{\ell=0}^L \mathbb{V}(\hat{Y}_\ell) = \sum_{\ell=0}^L \mathbb{V}(Y_\ell)/N_\ell$.

As for MC-FE, the mean squared error $\mathcal{E}_{A_{\text{MLMC}}}^2$ can be split into contributions from bias and variance as

$$\begin{aligned} \mathcal{E}_{A_{\text{MLMC}}}^2 &= \mathbb{E} \left[\|\mathbb{E}(u) - A_{\text{MLMC}}(u_L)\|_Z^2 \right] \\ &= \|\mathbb{E}(u) - \mathbb{E}(u_L)\|_Z^2 + \sum_{\ell=0}^L \frac{\mathbb{V}(Y_\ell)}{N_\ell} = \mathcal{E}_{\text{Bias}}^2 + \mathcal{E}_{\text{Stat}}^2. \end{aligned}$$

Similarly, using $\mathbb{E}(u) \approx \mathbb{E}(u_L)$, the normalized mean squared error $\hat{\mathcal{E}}_{A_{\text{MLMC}}}^2$ can be approximated by

$$\hat{\mathcal{E}}_{A_{\text{MLMC}}}^2 \approx \frac{\|\mathbb{E}(u) - \mathbb{E}(u_L)\|_Z^2}{\|\mathbb{E}(u_L)\|_Z^2} + \sum_{\ell=0}^L \frac{V_\ell}{N_\ell} = \hat{\mathcal{E}}_{\text{Bias}}^2 + \hat{\mathcal{E}}_{\text{Stat}}^2, \quad (12)$$

where

$$V_\ell := \mathbb{V}(Y_\ell) / \|\mathbb{E}(u_L)\|_Z^2. \quad (13)$$

As above, the parameter $\theta \in (0, 1)$ can be used to split the contributions of the two components of the error by requiring $\hat{\mathcal{E}}_{\text{Bias}}^2 \leq (1 - \theta)\epsilon^2$ and $\hat{\mathcal{E}}_{\text{Stat}}^2 \leq \theta\epsilon^2$ where ϵ is a predetermined tolerance such that $\hat{\mathcal{E}}_{A_{\text{MLMC}}}^2 \leq \epsilon^2$.

Note that each of the terms $Y_\ell^{(i)} := u_\ell^{(i)} - u_{\ell-1}^{(i)}$ appearing in (11) requires the approximation of $u^{(i)}$ on adjacent refinement levels using the same value of the parameter $\omega^{(i)}$. However, for FEM discretization, the numerical implementation of this term does not require the solution of the PDE on the two grids $\mathcal{T}_\ell$ and $\mathcal{T}_{\ell-1}$. The realization on the coarse grid can be obtained from the one on the fine grid by either Galerkin projection or interpolation. Projection is the more accurate choice, but it requires the solution of a system involving the mass matrix for each realization. Thus we prefer interpolation, for which costs are minimal and which is of second-order accuracy in our application. To avoid introducing correlation across discretization levels, none of the samples involved in the computation of Y_ℓ is reused for the finer level $\ell + 1$. That is, sampling is done such that, for $n \neq m$, the estimates Y_n and Y_m are uncorrelated. However, for any particular Y_ℓ the strong correlation between $u_\ell^{(i)}$ and $u_{\ell-1}^{(i)}$ makes the variance of the correction terms much smaller than the variance of the approximation u_L in the finest mesh, further improving the statistical approximation.

To quantify the computational effort of the MLMC-FE estimator, let $M_{\ell,i}$ be the number of grid points for the i -th sample on mesh level ℓ . We will assume that the computational cost to obtain one sample of $u_\ell^{(i)}$ is $C_{\ell,i} := C(u_\ell^{(i)}) = \mathcal{O}(M_{\ell,i}^c)$, where the exponent $c > 0$ depends on the solver, and will denote the cost of computing the correction term $Y_\ell^{(i)}$ by $C_{\ell,i} := C(Y_\ell^{(i)}) = \mathcal{O}(M_{\ell,i}^c)$ for $\ell \geq 0$ and $M_{-1,i} = 0$. For a nonlinear problem like the one at hand, the particular realization $\omega^{(i)}$ will influence the cost. We will consider the average cost to be of the form

$$C_\ell = \mathcal{O}(M_\ell^c) \quad (14)$$

and use this to estimate the total cost as $C(A_{\text{MLMC}}) = \sum_{\ell=0}^L N_\ell C_\ell = \sum_{\ell=0}^L N_\ell \cdot \mathcal{O}(M_\ell^c)$. Using the method of Lagrange multipliers, it is shown in [19] that this total cost, a function of N_ℓ , can be minimized subject to the constraint $\hat{\mathcal{E}}_{\text{Stat}}^2 \leq \theta\epsilon^2$ by the choice

$$N_\ell = \left\lceil \frac{1}{\theta\epsilon^2} \sqrt{\frac{V_\ell}{C_\ell}} \sum_{k=0}^L \sqrt{V_k C_k} \right\rceil. \quad (15)$$

With this expression for N_ℓ , the optimal total cost for the MLMC-FE estimator is

$$C(A_{\text{MLMC}}) \leq \frac{1}{\theta\epsilon^2} \left(\sum_{\ell=0}^L \sqrt{V_\ell C_\ell} \right)^2 + \sum_{\ell=0}^L C_\ell. \quad (16)$$

This analysis of computational cost is established in full generality in [19, Theorem 1].

The formula (15) suggests an iterative procedure for the approximation of $\mathbb{E}(u)$. Starting from a computational mesh $\mathcal{T}_0$, gather an initial number $\tilde{M}_0$ of samples $u_0^{(i)}$ and estimate $\hat{\mathcal{E}}_{\text{Bias}}$, $\hat{\mathcal{E}}_{\text{Stat}}$, and V_0 . If $\hat{\mathcal{E}}_{\text{Stat}}$ is larger than the prescribed tolerance, use (15) to update $\tilde{M}_0$ and gather additional samples; if $\hat{\mathcal{E}}_{\text{Bias}}$ is above the prescribed tolerance, then add an additional level of spatial refinement. The process continues adding discretization levels and collecting additional and/or samples until both $\hat{\mathcal{E}}_{\text{Bias}}$ and $\hat{\mathcal{E}}_{\text{Stat}}$ fall below the required tolerance, at which point $\mathbb{E}(u)$ is approximated using equation (9). This simple algorithm, however, presents one challenge: the term V_ℓ in equation (15) requires the computation of the term

$$\mathbb{V}(Y_\ell) = \frac{1}{N_\ell - 1} \left(\sum_{i=1}^{N_\ell} \|Y_\ell^{(i)}\|_Z^2 - \frac{1}{N_\ell} \left\| \sum_{i=1}^{N_\ell} Y_\ell^{(i)} \right\|_Z^2 \right). \quad (17)$$

However, the estimated sample sizes $\{N_\ell\}$ are available only for pre-existing discretization levels, hence whenever an additional mesh refinement is needed, N_{L+1} cannot be approximated by (15) as this formula uses V_{L+1} to compute N_{L+1} . This inconvenience can be overcome by noting that

$$\mathbb{V}(u - u_\ell) = \mathbb{E} \left[\|u - u_\ell\|_Z^2 \right] - \left\| \mathbb{E}(u - u_\ell) \right\|_Z^2 \leq \mathbb{E} \left[\|u - u_\ell\|_Z^2 \right]. \quad (18)$$

Hence, the variance $\mathbb{V}(u - u_\ell)$ can be estimated by the expectation of the squared discretization error $\|u - u_\ell\|_Z^2$. For a uniformly refined grid, we can resort to a standard *a priori* error estimate and assume that $\mathbb{E}[\|u - u_\ell\|_Z^2] = \mathcal{O}(M_\ell^{-b_1})$ for some $b_1 > 0$. It then follows that $V_\ell = \mathcal{O}(M_\ell^{-b_1})$ and following [40] we can then approximate V_{L+1} in terms of the known variance V_L by

$$V_{L+1} = (M_{L+1}/M_L)^{-b_1} V_L. \quad (19)$$

3.3. Adaptive multilevel Monte Carlo Finite-Element method

In view of the benefits of approximating the quantity of interest across a sequence of increasingly finer meshes, and with the goal of further reducing the computational cost associated with reducing the bias associated with the numerical discretization, it is natural to focus the refinement only on those parts of the mesh where the error is concentrated. Our goal is then to, starting from a computational mesh $\mathcal{T}_0$, generate a family of adaptively refined meshes $\{\mathcal{T}_\ell\}_{0 \leq \ell \leq L}$ that will produce better approximations of $\mathbb{E}(u)$ than the ones resulting from consecutive uniform refinements of the initial grid. With this goal in mind, the use of an *a posteriori* error estimator to guide the construction of the family of meshes has been proposed in the context of multilevel Monte Carlo methods [11,26,27,32].

Residual error estimator. A key ingredient in an adaptive solver is a local error estimator. In our case, for each element K on the mesh $\mathcal{T}_\ell$, we will use the simple residual-based *a posteriori* error indicator

$$\eta_{K,\ell}(\omega) := h_K^2 \left\| \nabla \cdot \left(\frac{1}{\mu r} \nabla u_\ell(\omega) \right) - f(u_\ell(\omega)) \right\|_K + h_K^{3/2} \left\| \left[\frac{1}{\mu r} \nabla u_\ell(\omega) \cdot \mathbf{n} \right] \right\|_{\partial K \setminus \partial \Omega}, \quad (20)$$

where ∂K is the boundary of the element K , h_K is the diameter of K , $\mathbf{n}$ is the outward unit normal to the element K , $[\cdot]$ denotes the jump across the edge of an interior element, and f is the source term defined piecewise on the right-hand side of (1a). Following [11,32], we will further define the mean local and mean global error estimators respectively as

$$\eta_{K,\ell} := \mathbb{E}(\eta_{K,\ell}(\omega)) \quad \text{and} \quad \eta_\ell^2 := \sum_{K \in \mathcal{T}_\ell} \eta_{K,\ell}^2. \quad (21)$$

For linear deterministic problems, estimators of this form can be shown to be such that there are constants $C_1, C_2 > 0$ such that $C_1 \eta_\ell \leq \|u - u_\ell\|_Z \leq C_2 \eta_\ell$. Therefore, the error estimator will accurately locate the regions of high error density and will decay at the same rate as the true error [34]. The global error can then be approximated by adding the local estimators over the entire triangulation. Using these error estimators, the adaptive analogue of (18) can be written as

$$\mathbb{V}(u - u_\ell) \leq \mathbb{E} \left[\|u - u_\ell\|_Z^2 \right] \approx \eta_\ell^2, \quad (22)$$

which then leads to the following adaptive analogue of the extrapolation formula (19)

$$V_{L+1} = (\eta_{L+1}/\eta_L)^2 V_L. \quad (23)$$

This estimate can then be used in combination with (15) to obtain an update for the sample size required at each adaptive level.

Adaptive solution cycle. With these definitions in place, we can then describe our strategy, which follows the “SOLVE → ESTIMATE

→ MARK → REFINES” paradigm familiar from deterministic adaptive solvers [7], as:

1. **Solve:** Starting from a fixed number of samples, the problem (1a) is solved on the initial mesh $\mathcal{T}_0$.
2. **Estimate:** The local mean error estimator is approximated from the sample gathered.
3. **Mark:** The set $\mathcal{M}_\ell$ containing the smallest possible number of elements in $\mathcal{T}_0$ satisfying

$$\sum_{K \in \mathcal{M}_\ell} \eta_{K,\ell}^2 \geq \zeta \eta_\ell^2, \quad (24)$$

for some predetermined value $\zeta \in [0, 1]$ is marked for refinement—this marking strategy is known as *Dörfler marking* in the adaptive finite element community [10].

4. **Refine:** The elements marked are refined in such a way that the resulting triangulation $\mathcal{T}_{\ell+1}$ is shape-regular and conforming. Efforts should be made to make sure that the growth of the number of elements is kept at bay. In our case, we used the implementation given in [16] of the algorithms described in [6,15,39].

The steps above are repeated until the error estimator η falls below a certain predetermined value.

A notion of mesh level. For uniformly refined grids, the notion of mesh level is natural: starting from a mesh $\mathcal{T}_\ell$, one step of uniform refinement decreases the mesh parameter h across the grid by a factor of $1/2$; the resulting mesh is said to have level $\ell + 1$ and is denoted by $\mathcal{T}_{\ell+1}$. For adaptively refined grids, where the mesh parameter is not constant through the grid, the notion of the level does not come so naturally. We will use the fact that, for a uniform refinement, the numerical error decays by a factor of $(1/2)^p$, (where p is the order of the FEM solver) with each successive level to extend the notion of mesh level to adaptively refined grids.

Consider a numerical approximation u_ℓ obtained on a mesh $\mathcal{T}_\ell$ with an associated error estimation given by η_ℓ . We will say that a mesh has level $\ell + 1$ and will denote it by $\mathcal{T}_{\ell+1}$ if it was obtained from $\mathcal{T}_\ell$ by cycling over the adaptive loop using the value $(1/2)^p \eta_\ell$ as the stopping tolerance. In other words, we will say that an adaptively refined mesh has level $\ell + 1$ if it produces a numerical solution with an error $(1/2)^p$ times smaller than one with level ℓ , just like in the uniform case. We will refer to $q := (1/2)^p$ as the decay factor. In terms of discretization accuracy, after ℓ steps of adaptive refinement, an adaptively refined mesh with level ℓ will have an associated error estimation $\eta_\ell = q^\ell \eta_0$, where η_0 corresponds to the error estimation at the initial mesh. In our numerical experiments, since the convergence rate of the piecewise linear solver is 2 (when measured in the L^2 norm), we shall use a decay factor $q = 1/4$ to define our adaptively refined meshes.

Deterministic adaptive grids. Ideally, in the stochastic setting, all the error estimations collected from the totality of samples would be used to drive the adaptive refinement forward and build an optimal set of meshes at every level. However, due to the iterative nature of the algorithm arising from (15), the optimal mesh at every level would have to be corrected with every new batch of samples and the solutions corresponding to all previous realizations $\omega^{(i)}$ would have to be recomputed. The computational cost of re-sampling in this manner quickly becomes impractical.

Instead, following [26,27,33], we will construct a sequence of deterministic adaptive grids with partial knowledge about $\mathbb{E}(u)$ as follows. Starting from a sample $\{\omega^{(i)}\}_{1 \leq i \leq N}$ (where N is small and arbitrarily chosen) and a mesh $\mathcal{T}_0$, the PDE is solved and the local error is estimated for every solution $u_0^{(i)}$, resulting in N local error estimators $\{\eta_{K,0}(\omega^{(i)})\}_{1 \leq i \leq N}$. The mean local and global error estimators $\eta_{K,0}$ and η_0 are then approximated by the sample means of the individual estimators. Using this approximation of $\eta_{K,0}$, the mesh $\mathcal{T}_0$ is refined.

Algorithm 1: Generate adaptive deterministic mesh set.

Input: Initial mesh $\mathcal{T}_0$, maximum mesh level L , element marking percentage $\zeta \in (0, 1)$, sample size N , error decay factor $q \in (0, 1)$.

Output: Adaptive mesh set $\{\mathcal{T}_\ell\}_{\ell=0}^L$.

```

1 for  $\ell = 1, \dots, L$  do
2   Refine = TRUE.
3   Set  $j = 1$  and  $\mathcal{T}^{(j)} = \mathcal{T}_{\ell-1}$ .
4   while Refine do
5     for  $i = 1, \dots, N$  do
6       (i) Draw a random sample  $\omega^{(i)}$  from  $W$ .
7       (ii) Obtain  $u^{(i)}$  by solving the free boundary problem on  $\mathcal{T}^{(j)}$  using  $\omega^{(i)}$ .
8       (iii) Calculate and accumulate the local and global error estimators  $\eta_{K,j}(\omega^{(i)})$  and  $\eta_j(\omega^{(i)})$ .
9     Compute the mean estimators  $\eta_{K,j}$  and  $\eta_j$  from the accumulated samples.
10    Mark & refine the mesh  $\mathcal{T}^{(j)}$  according to  $\eta_{K,j}$  to obtain  $\mathcal{T}^{(j+1)}$ .
11    if  $\eta_j \leq q\eta_1$  then
12      Store  $\mathcal{T}^{(j)}$  as  $\mathcal{T}_\ell$ .
13      Refine = FALSE.
14     $j = j + 1$ 

```

This process is continued until the approximated mean error estimator satisfies $\eta \leq q\eta_0$; the resulting mesh is stored and labeled as $\mathcal{T}_1$ (mesh level 1). The previous steps are repeated until a target number of meshes $\{\mathcal{T}_\ell\}_{0 \leq \ell \leq L}$ have been generated. The process is described in Algorithm 1. Since the family of meshes produced is constructed using random samples of ω , they approximately reduce the error for the approximate expectation $\mathbb{E}(u_h)$ by a factor of q with every increasing level. This family of meshes is then kept fixed during the MLMC run.

4. Numerical experiments

We now demonstrate and evaluate the performance of the methods presented for the Grad-Shafranov free boundary problem. We examine the efficiency of the three simulation approaches, comparing the CPU times (or computational cost) and the accuracy of some geometric descriptors that are generated from the approximation of $\mathbb{E}(\psi)$ obtained with each of the techniques. Following [14], we consider an ITER geometry with 12 coils and a “baseline” vector of target current intensities I given by

$$\begin{aligned}
I_1 &= -1.4 \times 10^6 \text{ A}, & I_2 &= -9.5 \times 10^6 \text{ A}, \\
I_5 &= -9 \times 10^6 \text{ A}, & I_6 &= 3.564 \times 10^6 \text{ A}, \\
I_9 &= -6.426 \times 10^6 \text{ A}, & I_{10} &= -4.82 \times 10^6 \text{ A}, \\
I_3 &= -2.0388 \times 10^7 \text{ A}, & I_4 &= -2.0388 \times 10^7 \text{ A}, \\
I_7 &= 5.469 \times 10^6 \text{ A}, & I_8 &= -2.266 \times 10^6 \text{ A}, \\
I_{11} &= -7.504 \times 10^6 \text{ A}, & I_{12} &= 1.724 \times 10^7 \text{ A}.
\end{aligned} \quad (25)$$

We will refer to these values as the *reference currents*. The profiles for p and g on the right hand side of (1a) follow the form given in (1b) with the specific values $r_0 = 6.2\text{m}$, $\beta = 0.5978$, $\alpha_1 = 2$, and $\alpha_2 = 1.395$, and $j_0 = 1.3655 \times 10^6 \text{ A/m}^2$. The reactor and coil array geometries follow the ones described in [2]. In our experiments, we will take the vector of current intensities to be subject to uncertainty modeled as a uniformly distributed perturbation of magnitude $\tau = 2\%$ centered around the reference values above.

4.1. Experiment description

In this section, we present numerical results comparing the three approaches – MC-FE, uniform MLMC-FE, and adaptive MLMC-FE. For the solution of (1), we used the finite element-based solver FEEQS.m [23] developed by Holger Heumann and collaborators as a lightweight




Fig. 2. Mean CPU time to compute 100 realizations of u_ℓ , as a function of the number of grid points M_ℓ , plotted on a logarithmic scale. The fitted curve indicates that the computational cost C_ℓ behaves approximately like $M_\ell^{1.09}$.

Matlab implementation of the code CEDRES++ [14,24]. The code implements a piecewise linear finite element discretization of a weak formulation of (1) and employs a globalized variation of Newton's method to resolve the nonlinearity. (The stopping threshold for the relative residual was set to 5×10^{-11} .) For the solution of the perturbed problems, the initial iterate of Newton's method was taken to be the solution corresponding to the reference currents I . All tests used the splitting parameter $\theta = 0.5$ in (6). The user-specified tolerances for the normalized mean squared error range from $\epsilon = 2 \times 10^{-4}$ to 8×10^{-3} . Experiments were conducted using MATLAB R2022a on a System 76 Thelio Major with 256GB RAM and a 64-Core @4.6 GHz AMD Threadripper 3 Processor.

To produce an estimate of the number of samples required on each discretization level, equation (15) requires the knowledge of two problem-dependent parameters: the power c appearing in the estimate for the computational cost (14), and the normalized variance of the correction terms V_ℓ as defined in (13). The normalization factor $\|E(u_L)\|_Z$ in (13) was estimated on the finest uniform mesh level ($\ell = 5$) to be approximately 8.57×10^{-1} . To estimate the value of c , 100 random currents are sampled for different mesh sizes M_ℓ , the processing times required to obtain the solutions are averaged for each mesh size, and c is estimated through a regression. Fig. 2 shows the behavior of the average cost as a function of the mesh size M_ℓ ; from the data displayed, the power law is estimated to be $c \approx 1.09$. Note that this cost estimate is based on Matlab timings and not on the complexity analysis of standard linear solution algorithms. The same samples are also used to estimate the sample means, $E(Y_\ell)$ or $E(u_h)$, and variances, $V(Y_\ell)$ or $V(u_h)$ dynamically using Welford's algorithm [45]. As the new samples are gathered, the mean m_w and proxy for the variance, s_w , are updated using the following formulas for the i -th sample, with $m_w^{(0)} = 0$, $s_w^{(0)} = 0$:

$$m_w^{(i)} = m_w^{(i-1)} + \frac{u^{(i)} - m_w^{(i-1)}}{i}, \quad s_w^{(i)} = s_w^{(i-1)} + \left\langle u^{(i)} - m_w^{(i-1)}, u^{(i)} - m_w^{(i)} \right\rangle.$$

Using sample size i , the variance is then given by $V^{(i)} = s_w^{(i)} / (i - 1)$.

To perform MLMC-FE simulations, the user typically defines and generates a sequence of spatial grids, where, given a tolerance ϵ , the fineness of the grid is determined by the requirement that the discretization error ($\hat{\epsilon}_{\text{Bias}}^2$ in (6)) or an estimate of it be small enough. In this study, we generated two types of grids, a set of *geometry-conforming* uniformly refined grids, and a set of adaptively refined grids constructed using the strategy presented in Section 3.3.¹ For the uniformly refined grid, we generated a total of six levels of grids. We note that, due to the

Table 1

The number of grid points M_ℓ for both geometry-conforming uniform and adaptive ($q = 1/4$) meshes as the mesh levels increase from 0 to 5.

Level ℓ	0	1	2	3	4	5
Uniform	2685	8019	30449	120697	484080	1934365
Adaptive	2685	6090	25099	103968	411913	1552282

increasing accuracy of the spline approximation of the curved boundaries, these meshes are not a direct refinement of each other. Instead, each level is characterized by a mesh parameter h_ℓ being roughly half of the preceding mesh as determined by the Triangle mesh generator [43]. The adaptive refinement strategy began with the coarsest mesh from the uniform family and applied the weighted L_2 a posteriori error estimator specified in (20) and $q = 1/4$ to reflect a similar error decay as for uniform refinement, also using Triangle to generate the desired adaptively refined meshes. The number of grid points for each of these methods on different grid levels is shown in Table 1. Note that the grid sizes for the adaptive meshes are not dramatically smaller than for uniform meshes, which suggests that the solution, as a function of the spatial variables, is fairly smooth.

4.2. Computational cost

Fig. 3 shows a variety of computational results, including the error estimator, V_ℓ , and CPU time for the two versions of MLMC (and times for full MC). To investigate the convergence behavior of the discretization error, we calculate the a posteriori error estimator for both uniform and adaptive meshes in the same experiment to obtain an estimate of c for C_ℓ before conducting the simulations. The results are displayed in the top left plot in Fig. 3, with a dashed line showing a least square fit, indicating that the discretization error of both methods exhibits an asymptotic rate of $\mathcal{O}(M^{-1})$ (or $p \approx 2$). The similar convergence rate further indicates that the solution to the problem is smooth and the error is equidistributed, rendering the adaptive strategy comparable to uniform mesh refinement. Note that the estimated error is used for variance extrapolation in (23) during the MLMC simulations.

The top right plot of Fig. 3 shows the behavior of V_ℓ of (13) for both uniform and adaptive MLMC-FE methods with $\epsilon = 2 \times 10^{-4}$, using six levels of meshes. It can be seen that both methods demonstrate a decreasing trend in the values of V_ℓ as the mesh resolution increases, with a power law decay characterized by $b \approx 2$ in the least square fit. But there is a regime for a small number of grid points where the asymptotic behavior of the adaptive method is not evident, in contrast to the behavior of the uniform method. As the meshes get finer, the plots of V_ℓ for the two methods are close to being parallel.

The computational effort for uniform MLMC-FE and MC-FE scales as $\mathcal{O}(\epsilon^{-2})$ and $\mathcal{O}(\epsilon^{-3})$ respectively, as indicated by the slopes of the least square fitting lines for the red and yellow curves. This observation is consistent with the theoretical cost predictions in Theorem 1 (with $b > c$) and (8). Theorem 1 also indicates that the majority of computational work is performed on coarse grids. Table 2 shows the sample sizes obtained from (7) for MC-FE and (15) for MLMC-FE, further demonstrating a decrease in $N_\ell C_\ell$ as ℓ increases for the multilevel methods. We also found that the computational cost associated with the smallest tolerance $\epsilon = 2 \times 10^{-4}$ is so large that we were unable to generate MC-FE results on a fine mesh ($\ell = 5$) with a large sample size. In contrast, both versions of MLMC-FE successfully generated results with this tolerance.²

¹ The domain components contain curved boundaries, which we handled by treating them as polygonal structures. The mesh generation entails identifying the curved boundaries using piecewise splines and interpolating along these splines with grids of varying fineness.

² Although we could not directly generate results for MC-FE for $\epsilon = 2 \times 10^{-4}$, we could estimate the costs. In particular, we found that the variance V_h is close to constant across mesh levels. Consequently, we used (7) to estimate the number of required samples as 8000 in Table 2, approximately four times the number required for $\epsilon = 4 \times 10^{-4}$. This number was multiplied by the mean CPU




Fig. 3. Top left: weighted L_2 error (with weight μx) of estimator $C\eta_\ell$ vs. number of grid points M_ℓ plot. Top right: normalized variance V_ℓ vs. number of grid points M_ℓ plot. Bottom left: CPU time in seconds vs. tolerance ϵ . Bottom right: Monte Carlo convergence rate estimate with tolerance ϵ vs. sample size N . This plot is generated from Table 2.

Table 2

The optimal sample size estimation for MC-FE (left), uniform MLMC-FE (middle), and adaptive MLMC-FE (right). The simulations were conducted for a variety of choices of ϵ . The computational cost associated with a tolerance of $\epsilon = 2 \times 10^{-4}$ for Monte Carlo was prohibitive; the entry in the table for this tolerance (with an asterisk) is an estimate.

ϵ	Level ℓ						ϵ	Level ℓ						ϵ	Level ℓ					
	0	1	2	3	4	5		0	1	2	3	4	5		0	1	2	3	4	5
8×10^{-3}			5				8×10^{-3}	10	2	2				8×10^{-3}	10	2	2			
6×10^{-3}			7				6×10^{-3}	12	3	2				6×10^{-3}	19	3	3			
4×10^{-3}				22			4×10^{-3}	32	5	2	2			4×10^{-3}	38	5	4			
2×10^{-3}				83			2×10^{-3}	152	26	4	2			2×10^{-3}	121	18	6	2		
10^{-3}					322		10^{-3}	691	109	18	4	2		10^{-3}	697	82	49	8		
8×10^{-4}					527		8×10^{-4}	841	129	23	3	2		8×10^{-4}	1446	118	91	27	6	
6×10^{-4}					869		6×10^{-4}	1610	231	40	8	2		6×10^{-4}	2070	218	133	21	3	
4×10^{-4}					1980		4×10^{-4}	3791	589	104	15	3		4×10^{-4}	5075	484	315	61	14	
2×10^{-4}						8000*	2×10^{-4}	15859	2344	375	62	13	2	2×10^{-4}	25871	1961	1668	430	85	14

As seen in the bottom left plot of Fig. 3, the uniform and adaptive versions of MLMC-FE have similar computational costs of $\mathcal{O}(\epsilon^{-2})$, as evidenced by the similar decay rate of the error estimator and (22). According to (16), the slightly smaller magnitude of the error estimator for the adaptive MLMC-FE suggests a smaller (or comparable) computational cost in the asymptotic regime. However, when $\epsilon = 4 \times 10^{-4}$, the adaptive MLMC-FE method requires approximately twice as much CPU time (1.79×10^3 seconds) compared to the uniform MLMC-FE approach (9.29×10^2 seconds) due to a notable increase in V_ℓ around

$M_\ell = 10^4$. This also causes the speedups achieved using adaptive refinement to be somewhat smaller than for uniform refinement. Thus, the traditional advantage of adaptive mesh refinement is not clearly present. We also attribute this to the apparent smoothness of the solution. We will demonstrate some advantages of the adaptive strategy in Section 4.3.

The bottom right plot of Fig. 3 shows that the nMSE tolerance ϵ of MC-FE approach declines at $\mathcal{O}(N^{-0.51})$, which is consistent with the well-known square root convergence rate. This rate holds since $\mathbb{V}(u_h)$ remains nearly constant among all levels.

4.3. Properties of geometric parameters

Next, we will explore the plasma boundaries and geometric descriptors of the expected poloidal flux ψ resulting from the three methods. To ensure a fair comparison, we will use the results obtained from the MC-FE on the finest uniform mesh as a reference benchmark.

time observed for the computations for Fig. 2 (120.3 seconds, the largest entry appearing in the figure) to give the estimated total CPU time in Table 3. Table 3 gives quantitative data on the costs in CPU time for the three methods, as well as the speedups achieved by the two multilevel methods. It can be seen that for small tolerance ϵ , both these methods reduce the CPU times dramatically, with many examples of speedups greater than a factor of 60 and a best-case speedup of approximately 200.




Fig. 4. The overlaid plasma boundaries of 50 random realizations are displayed in the top row as violet curves. The solid violet line is the plasma boundary of the expected poloidal flux generated with tolerance $\epsilon = 4 \times 10^{-4}$. The inner and outer walls of the reactor are displayed in solid black and dark red respectively. The bottom row shows the regions close to the x-points in more detail. The dark green dots are the x-points of the expected solution. Each column from left to right corresponds to: simulation with the Monte Carlo approach, MLMC simulation on geometry-conforming uniform meshes, and adaptive MLMC simulation. All simulations were performed using the discretization level $\ell = 5$. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Table 3

The CPU time in seconds for MC-FE (left), uniform MLMC-FE (middle), and adaptive MLMC-FE (right), together with speedups for the multilevel methods, for a variety of choices of ϵ . The computational cost associated with a tolerance of $\epsilon = 2 \times 10^{-4}$ for Monte Carlo was prohibitive; the entry in the table for this tolerance (with an asterisk) is an estimate.

ϵ	MC-FE	Uniform MLMC-FE		Adaptive MLMC-FE	
	Time	Time	Speedup	Time	Speedup
8×10^{-3}	5.67e+00	4.52e+00	1.3	3.50e+00	1.6
6×10^{-3}	7.69e+00	5.25e+00	1.5	5.51e+00	1.4
4×10^{-3}	1.30e+02	2.32e+01	5.6	9.01e+00	14.4
2×10^{-3}	4.83e+02	4.62e+01	10.5	3.16e+01	15.3
10^{-3}	9.22e+03	2.47e+02	37.3	1.84e+02	50.1
8×10^{-4}	1.50e+04	2.48e+02	60.5	5.73e+02	26.2
6×10^{-4}	2.48e+04	4.13e+02	60.0	6.30e+02	39.4
4×10^{-4}	5.68e+04	9.29e+02	61.1	1.79e+03	31.7
2×10^{-4}	9.62e+05*	4.21e+03	228.5	1.21e+04	79.6

Plasma boundary. To ascertain the expected location of the plasma boundary, we first determine the expected solution to the free boundary problem, $\mathbb{E}[\psi]$, and determine the boundary of this deterministic function. This boundary is depicted, in dark violet, in Fig. 4 along with the plasma boundaries obtained from 50 random currents, which are shown in light violet curves. In Fig. 5 we present plots depicting the x-points and plasma boundaries of the expected solution ψ computed using only samples and corrections from increasingly finer grids for both uniform and adaptive MLMC-FE approaches. The data was obtained with tolerance $\epsilon = 4 \times 10^{-4}$. As can be seen when moving from left to right in Fig. 5, the result obtained using the information from the coarsest level (leftmost column) is progressively corrected with information from increasingly finer grids, leading to the desired result depicted in the rightmost column.

Among the three methods, MC-FE yields the smoothest plasma boundaries in the vicinity of the x-point, followed by adaptive MLMC-FE, while the MLMC-FE approach on geometry-conforming uniform meshes manifests the most pronounced irregularity in the plasma boundary. The boundary of the expected solution generated with the

Table 4

Geometric parameters of the expected poloidal flux $\mathbb{E}[\psi]$ from MC-FE, MLMC-FE with geometry-conforming uniform mesh set, and adaptive MLMC-FE. The results are generated with an nMSE 4×10^{-4} .

	MC-FE	Uniform MLMC-FE	Adaptive MLMC-FE
x point	(5.14,-3.29)	(5.14,-3.29)	(5.14,-3.28)
magnetic axis	(6.41,0.61)	(6.44,0.56)	(6.46,0.54)
strike points	(4.16,-3.71)	(4.16,-3.71)	(4.16,-3.71)
inverse aspect ratio	0.32	0.32	0.32
elongation	1.86	1.87	1.86
upper triangularity	0.43	0.43	0.43
lower triangularity	0.53	0.53	0.53

uniform grid MLMC-FE method exhibits irregularities as can be seen in Fig. 4. These large deformations can be primarily attributed to the additional challenges arising from the curved boundaries. We will address this point in more detail at the end of this section. The top row of Fig. 5 demonstrates that using a geometry-conforming mesh provides a more accurate approximation of the curved structure (in black) of the configuration than that in the bottom row. These observations underscore the challenge of striking a balance between preserving geometric fidelity when dealing with curved boundaries and the desired statistical accuracy of the solution.

Geometric descriptors. Table 4 reports some geometric parameters of the expected poloidal flux $\mathbb{E}[\psi]$ in (3). It is observed that these parameters are consistent across different simulation techniques, with agreement typically up to two or, in some cases, one significant digit. Having been derived from $\mathbb{E}[\psi]$ all these values are deterministic. There is, however, uncertainty associated with the corresponding quantities derived from ψ , as they are themselves random variables. Uncertainty on these quantities can be assessed by computing Monte Carlo estimates of their expectations and variances. The multi-level methodology can be used for this purpose without any modification by simply regarding them as quantities of interest in their own right. All the relevant descriptors are computed for each of the samples gathered during the compu-




Fig. 5. The violet curves represent the expected plasma boundaries for simulations using increasingly finer grids when $\epsilon = 4 \times 10^{-4}$ using the sample sizes specified in Table 2. Each sub-plot focuses on a region near the x-point, maintaining the same zoom-in ratio as the second row of Fig. 4. The dark green dots denote the locations of the x-point. The top row shows the results of MLMC-FE on a set of geometry-conforming uniform meshes, while the bottom row displays the results for adaptive MLMC.

Algorithm 2: Multilevel Monte Carlo Finite-Element method.

Input: Initial mesh level $L = 0$, sequence of mesh available $\{\mathcal{T}_\ell\}_{\ell \geq 0}$,
 root nMSE ϵ , $\zeta \in (0, 1)$, initial sample size $N_{\text{old}} = \{N_\ell\}_{\ell=0}^L$,
 counter $j = 1$, sample size corrections $dN_\ell^j = \{N_\ell\}_{\ell=0}^L$.

Output: $\{N_\ell\}_{\ell=0}^L$, $\{V_\ell\}_{\ell=0}^L$, A_{MLMC} .

```

1 while  $\sum_\ell (dN_\ell^j) > 0$  do
2   for  $0 \leq \ell \leq L$  do
3     for  $i = 1, \dots, dN_\ell^j$  do
4       Solve the free boundary problem on  $\mathcal{T}_\ell$  to get  $u_\ell^{(i)}$  for the
        i-th sample.
5   Approximate  $\{V_\ell\}_{\ell=0}^L$  by (17).
6   Update the sample size estimation  $\{N_\ell\}_{\ell=0}^L$  by (15).
7    $j = j + 1$ .
8    $dN_\ell^j = \{N_\ell\}_{\ell=0}^L - N_{\text{old}}$ .
9    $N_{\text{old}} = \{N_\ell\}_{\ell=0}^L$ .
10  if  $\sum_\ell (dN_\ell^j) = 0$  then
11    if The a posteriori error estimator falls below  $\sqrt{1 - \theta\epsilon} \|\mathbb{E}(u)\|_Z$ ,
12      then
13        Compute  $A_{\text{MLMC}}$  by (11) and terminate the loop.
14    else
15       $L = L + 1$ .
16      Approximate  $V_L$  by (19) and compute  $\{N_\ell\}_{\ell=0}^L$  and go to
        Step 1.

```

tational loop at the various discretization levels and their expectations and variances are estimated following the same process described on Algorithm 2. The expectations and variances obtained in this fashion are shown in Table 5.

Despite the advantages of low computational cost, the MLMC-FE-based methods may encounter difficulties in accurately determining the locations of x-points and magnetic axis. Note that the x-points, which correspond to saddle points of the piecewise linear approximation of ψ , can only be located at the nodes of the mesh. The numerical identification of their exact locations, which often relies on changes in the sign of the discrete gradient, can be challenging; see [5,12,28] for discussions of the computational difficulties.

In summary, simulations using the uniform MLMC-FE on non-nested geometry-conforming uniform meshes may encounter a substantial challenge in accurately identifying the x-point and achieving less ac-

curate quantities, especially for the plasma boundary, when compared to the results obtained from MC-FE. On the other hand, the adaptive MLMC-FE approach on a nested adaptively refined mesh set produces results that closely align with the MC-FE at a much lower computational cost.

Meshing curved domains and their effect on MLMC estimations. In the deterministic setting this *geometric error* has the undesired consequence of hindering the decay of the discretization error since, as the mesh is refined, the discretized computational domain does not converge to the semicircle bounded by Γ . In the stochastic setting the geometric error manifests itself in rendering the Monte Carlo estimator biased and inconsistent. The inconsistency stems from the fact that, as both the sample size and the mesh level increase, the Monte Carlo estimator does not converge to the expectation of the random variable u satisfying the free boundary problem. Instead, the estimator converges to the expectation of the random variable that satisfies a perturbation of (28) where the curve Γ is not a semicircle, but the initial polygonal approximation. If the initial mesh is fine enough, this geometric bias will likely be too small to affect the estimation.

On the other hand, if an exact descriptor of the curved boundaries is available, the aforementioned difficulty can be overcome by re-sampling the curved boundaries when building the sequence of finer grids, thus allowing for a resolution of the curved structures consistent with the respective mesh parameter. If an exact descriptor is not available it is possible to approximate it with, for instance, a cubic spline representation that interpolates the original polygonal representation. This spline surrogate is then used to re-sample the boundary as the mesh is refined. This strategy, which gives rise to what we refer to as *geometry conforming meshes*, was implemented for the numerical experiments with *uniformly refined meshes* and can be seen in use in the top row of Fig. 5, where the curved boundary is represented more accurately as the mesh is refined.

Nevertheless, even if the approximation to the problem geometry is now consistent with the discretization error, this approach creates additional challenges. Since the approximations to the curved boundaries are not fixed across mesh levels, the sequence of meshes is no longer nested—not even in the case of uniform refinements. Moreover, due to the fact that the sequence of discrete domains no longer coincides across levels, the domains of definition of the respective discrete solutions will not overlap, and an extrapolation step may be needed to compute the multilevel Monte Carlo estimator on a common computational domain. This strategy, used in our numerical experiments, introduces an addi-

Table 5

Sample mean of the geometric parameters extracted from each realization during three simulations MC-FE, uniform MLMC-FE, and adaptive MLMC-FE. The results are generated with an mNSE 4×10^{-4} .

Geometric parameters	MC		MLMC			
	-		Uniform		Adaptive	
	Mean	Variance	Mean	Variance	Mean	Variance
x point	(5.14,-3.29)	(2.37e-04,1.44e-03)	(5.14,-3.28)	(2.44e-04,6.33e-04)	(5.14,-3.29)	(4.03e-04,2.91e-03)
magnetic axis	(6.41,0.61)	(1.01e-03,6.11e-04)	(6.41,0.60)	(2.65e-02,6.28e-03)	(6.42,0.60)	(3.15e-02,5.81e-03)
strike	(4.17,-3.71)	(4.46e-03,2.23e-03)	(4.16,-3.71)	(1.23e-04,4.30e-05)	(4.17,-3.71)	(4.18e-04,2.00e-04)
points	(5.56,-4.22)	(2.88e-08,3.28e-03)	(5.56,-4.22)	(4.05e-09,5.54e-05)	(5.56,-4.21)	(2.02e-08,2.98e-04)
inverse aspect ratio	0.32	4.76e-06	0.32	4.23e-06	0.32	4.31e-06
elongation	1.86	1.50e-04	1.86	1.83e-03	1.86	2.29e-03
upper triangularity	0.43	2.49e-04	0.43	1.45e-03	0.43	2.08e-03
lower triangularity	0.53	1.08e-04	0.53	1.52e-03	0.53	1.31e-03




Fig. 6. The violet curves represent the expected plasma boundaries for post-processed simulations using increasingly finer grids when $\epsilon = 4 \times 10^{-4}$ and sample size specified in Table 2. Each sub-plot focuses on a region near the x-point, maintaining the same zoom-in ratio as the second row of Fig. 4. The dark green dots denote the locations of the x-point.

tional extrapolation error. In our case this is evident, for instance, in the fact that the plasma boundary of the expected solution $\mathbb{E}[\psi_h]$ is considerably less regular in the geometry-conforming case than it is in the non-geometry-conforming one. This can be seen in Fig. 5. The extrapolation error can be taken care of through careful post-processing. One option is to project or interpolate the numerical solutions into a subdomain common to all grids so that no extrapolation is needed for evaluation. This strategy was employed to produce Fig. 6 successfully eliminating the spurious oscillations in the plasma boundary. However, doing this requires considerable computational work and reduces the time savings obtained from MLMC.

One further difficulty is that the re-sampling of the boundaries is impossible to perform in a straightforward fashion in the case of adaptively refined meshes. Thus, the geometric approximation remains fixed at the initial level of refinement. This can be seen in the bottom row of Fig. 5, where the solid black line represents the polygonal approximation to the curved boundary of the divertor. The approximation improves as the mesh is refined for the uniformly refined mesh, but stays fixed for the adaptive strategy.

5. Concluding remarks

The objective of this study is to evaluate the performance of MC-FE and several variants of MLMC-FE for the Grad-Shafranov free boundary problem under high-dimensional uncertainties in the currents. The plasma equilibrium problem presents challenges and complexities associated with the physical properties of the system, the values of the parameters of the experiment, and even the approximation of geometric features of the model. Examples include the degradation of the stability of the equilibrium associated with an increase of plasma elongation or beta; the presence of finite current near the plasma boundary or, as in this study, the complex boundary structures, which must be treated carefully in the context of multilevel methods. These physically grounded challenges typically get reflected in the computations, for instance, as an increase in the difficulty of a linearized algorithm to converge to a desired tolerance or as the need for higher precision in the

numerical computation required to resolve the desired features. In other words, complicated physical scenarios typically result in an increase in computational costs which might easily become overwhelming when combined with uncertainty quantification efforts. In cases like these, an MLMC strategy can help to mitigate the costs required to gather statistical information at the expense of potentially overlooking subtle behavior that may not be resolved in the samples gathered in coarse levels of refinement.

It is important to remark that the multilevel Monte Carlo techniques presented rely only on the independence between different realizations *across discretization levels* and not on the particular probability distribution obeyed by the stochastic parameters. In this study, the current intensities were modeled as independent uniformly distributed random variables, but the method would have worked without any modification even if the current intensities had obeyed correlated distributions *as long as realizations across discretization levels are independent of each other*. Moreover, further stochastic components—such as the parameters appearing in the profile models in equation (1b)—can be included and handled simultaneously.

In this study, we handled the geometric difficulties in the uniform MLMC-FE case using a sequence of uniform meshes that conform to the curved boundaries in the geometry. We also found that a feature of adaptive MLMC-FE is that adaptive gridding leads to a somewhat better representation of geometric quantities; in contrast, errors introduced by computations on uniform coarse meshes lead to some distortions of some features. The traditional advantage of adaptive refinement reduced computational cost, was less clearly present. But the overall advantage of the MLMC approach in reducing costs is dramatic.

6. Acknowledgments

All the free boundary computations were done using the code FEEQS.M [23,24]. The authors are deeply grateful to Holger Heumann, the INRIA CASTOR team, and all the development team of the CEDRES++ free boundary solver for sharing access to the code and for helping us get up to speed with its usage.

Howard Elman has been partially supported by the U.S. Department of Energy under grant DE-SC0018149 and by the U.S. National Science Foundation under grant NSF-DMS-1819115. Jiaying Liang has been partially funded by the U.S. Department of Energy under grant DE-SC0018149 and by the U.S. National Science Foundation under grant NSF-DMS-1819115. Tonatiuh Sánchez-Vizuet has been partially funded by the U.S. National Science Foundation through the grant NSF-DMS-2137305.

7. Appendix

We now describe some technical mathematical and implementation details related to the work presented. This additional material is provided for the sake of completeness but is not needed to follow the discussion from the previous sections.

7.1. Weak formulation of the Grad-Shafranov equation

Deterministic problem. We will start by introducing the weak formulation for the Grad-Shafranov problem in the case where all the parameters involved are deterministic. Consider a semi-circle centered at the origin, boundary Γ and with radius ρ such that it fully contains all the relevant reactor components depicted in Fig. 1. If Ω denotes the region surrounded by Γ then, by construction, for any in Ω^c the right-hand side of (1a) will vanish identically. We will then, following [22], consider the space of real-valued functions

$$Z := \left\{ u : \Omega \rightarrow \mathbb{R} \mid \int_{\Omega} u^2 x dx dy < \infty, \right. \\ \left. \int_{\Omega} \frac{|\nabla u|^2}{x} dx dy < \infty; \text{ and } u(0, y) = 0 \right\} \cap C^0(\overline{\Omega}). \quad (26)$$

This space arises naturally when testing equation (1a) using the weighted L_2 inner product defined by

$$\langle u, v \rangle := \int_{\Omega} uvx dx dy,$$

which leads to the finite energy requirement appearing in the second inequality in the definition (26). The third requirement ($u(0, y) = 0$) is a result of the anti-symmetry of the problem with respect to reflections across the axis of symmetry of the reactor and has the effect of ensuring that the quantity

$$\|u\|_Z := \left(\int_{\Omega} \frac{|\nabla u|^2}{x} dx dy \right)^{1/2} \quad (27)$$

does indeed define a norm in the space Z . We will refer to this norm as the *energy norm*. The space Z defined in (26) endowed with the energy norm (27) is the natural function space to look for variational solutions to the deterministic linearized Grad-Shafranov equation (1). The variational formulation of the problem, as derived in [1,17,29], is that of finding $\psi \in Z$ such that for every test $\varphi \in Z$ it holds that:

$$\int_{\Omega_p} \frac{1}{\mu x} \nabla \psi \cdot \nabla \varphi dx dy - \int_{\Omega_p} \left(rp'(\psi) + \frac{1}{\mu_0 x} f f'(\psi) \right) \varphi dx dy \\ + \int_{\Gamma} \psi N \varphi dc \\ + \int_{\Gamma} \int_{\Gamma} (\psi(x_1) - \psi(x_2)) M(x_1, x_2) (\varphi(x_1) - \varphi(x_2)) dc(x_1) dc(x_2)$$

$$= \sum_{k=1}^{M_c} \frac{I_k}{S_k} \int_{\Omega_{C_k}} \varphi dx dy. \quad (28)$$

Above, the magnetic permeability μ is either a function of ψ inside a region occupied by a ferromagnetic material, $\mu = \mu(|\nabla \psi|^2/x^2)$, or a constant $\mu = \mu_0$ elsewhere, $\mathbf{x}_i = (x_i, y_i)$ denote position vectors, Ω_{C_k} denotes the region occupied by the k -th coil, and the total number of coils is denoted by M_c . The following quantities, appearing on the left-hand side of (28), are related to the Green's function associated with the operator on the left-hand side of (1a)

$$N := \frac{1}{x} \left(\frac{1}{\delta_+} + \frac{1}{\delta_-} - \frac{1}{\rho} \right), \quad \delta_{\pm} := \sqrt{x^2 + (\rho \pm y)^2},$$

$$\kappa(\mathbf{x}_1, \mathbf{x}_2) := \sqrt{\frac{4x_1x_2}{(x_1+x_2)^2 + (y_1-y_2)^2}},$$

$$M(\mathbf{x}_1, \mathbf{x}_2) := \frac{\kappa(\mathbf{x}_1, \mathbf{x}_2)}{2\pi(x_1x_2)^{3/2}} \\ \times \left(\frac{2 - \kappa^2(\mathbf{x}_1, \mathbf{x}_2)}{2 - 2\kappa^2(\mathbf{x}_1, \mathbf{x}_2)} E(\kappa(\mathbf{x}_1, \mathbf{x}_2)) - K(\kappa(\mathbf{x}_1, \mathbf{x}_2)) \right),$$

where $E(\kappa(\mathbf{x}_1, \mathbf{x}_2))$ and $K(\kappa(\mathbf{x}_1, \mathbf{x}_2))$ are complete elliptic integrals of the first and second kind respectively [30].

Accounting for stochasticity. We now consider the stochasticity in the currents and allow the vector of currents to be a d -dimensional random variable ω uniformly distributed over the parameter space W defined in (2). It is clear that in this case for any particular realization of the currents ω we will obtain a *different* equilibrium configuration $\psi(\omega)$ that belongs to the Banach space Z defined in (26). Moreover, since for every $\omega \in W$ the resulting equilibrium configuration has finite energy, it then holds for the expected value of the equilibrium that

$$\mathbb{E}(\|\psi\|_Z^2) < \infty.$$

In mathematical terms, we say that the stochasticity of the currents transforms the solution ψ to (1) into a Banach space-valued random variable with finite expected energy. If ω belongs to a complete and separable probability space $(W, \Sigma, \mathbb{P})$, the class of such random variables forms what is known as a *Bochner space* [9]. In our particular case, solutions to (1) are mappings from the parameter space W to the Banach space Z that, as functions of ω , belong to the Bochner space

$$L^2(W, \Sigma, \mathbb{P}; Z)$$

$$:= \{u : W \rightarrow Z \mid u \text{ strongly measurable, } \|u\|_{L^2(W, Z)} < \infty\},$$

where the norm $\|\cdot\|_{L^2(W, Z)}$ is precisely defined in terms of the expected energy

$$\|u\|_{L^2(W, Z)} := \left(\int_W \|u(\cdot, \omega)\|_Z^2 d\mathbb{P}(\omega) \right)^{1/2} = (\mathbb{E}(\|u(\cdot, \omega)\|_Z^2))^{1/2}.$$

7.2. Mathematical results on multilevel Monte Carlo

The proceeding analysis on the computational cost to compute the MLMC-FE estimator in terms of the desired relative accuracy ϵ is based on the following Theorem (established and proven in [8]), which quantifies the distribution of the computational effort across discretization levels in terms of the relation between the decay rate of the variance of the correction terms, V_{ℓ} , and the increase of the computational cost, C_{ℓ} , as the mesh is refined. In particular, it states that if the variance of the correction terms decays faster than the increase of computational cost, the dominant computational expense takes place on the coarsest grid—see also [19].

Theorem 1. Suppose there exist positive constants a, b, c such that $a \geq \frac{1}{2} \min(b, c)$,

- (i) $\|E(u - u_\ell)\|_Z = \mathcal{O}(M_\ell^{-a})$,
- (ii) $V_\ell = \mathcal{O}(M_\ell^{-b})$,
- (iii) $C_\ell = \mathcal{O}(M_\ell^c)$.

Then for any positive $\epsilon < e^{-1}$ small enough, there exists level L and sample size N_ℓ for which the multilevel estimator A_{MLMC} has an nMSE with

$$\frac{\mathbb{E}[\|E(u) - A_{MLMC}(u_L)\|_Z^2]}{\mathbb{E}[\|E(u)\|_Z^2]} < \epsilon^2,$$

and the total computation cost C with bound

$$C(A_{MLMC}) = \begin{cases} \mathcal{O}(\epsilon^{-2}), & b > c, \\ \mathcal{O}(\epsilon^{-2}(\log \epsilon)^2), & b = c, \\ \mathcal{O}(\epsilon^{-2 - \frac{(c-b)}{a}}), & 0 < b < c. \end{cases}$$

CRedit authorship contribution statement

Howard C. Elman: Formal analysis. **Jiaxing Liang:** Formal analysis. **Tonatiuh Sánchez-Vizuet:** Formal analysis.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Tonatiuh Sanchez-Vizuet reports financial support was provided by the U. S. National Science Foundation through the grant NSF-DMS-2137305. Howard C. Elman reports financial support was provided by the U. S. National Science Foundation through the grant NSF-DMS-1819115 and by the U.S. Department of Energy through the grant DE-SC0018149. Jiaxing Liang reports financial support was provided by the U. S. National Science Foundation through the grant NSF-DMS-1819115 and by the U.S. Department of Energy through the grant DE-SC0018149.

Data availability

No data was used for the research described in the article.

References

- [1] R. Albanese, J. Blum, O.D. Barbieri, On the solution of the magnetic flux equation in an infinite domain, in: S. Methfessel, G. Thomas (Eds.), 8th European Conference on Computational Physics: Computing in Plasma Physics, in: Europhysics Conference Abstracts, vol. 10D, European Physical Society, 1986, pp. 41–44.
- [2] V. Amoskov, A. Belov, V. Belyakov, Y. Gribov, A. Kavin, V. Kukhtin, E. Lamzin, K. Lobanov, N. Maximenkova, A. Mineev, S. Sytchevsky, Stray magnetic field at plasma initiation produced by ferromagnetic elements of the ITER tokamak complex, Plasma Devices Oper. 17 (4) (Dec. 2009) 238–249.
- [3] I. Babuška, F. Nobile, R. Tempone, A stochastic collocation method for elliptic partial differential equations with random input data, SIAM J. Numer. Anal. 45 (3) (2007) 1005–1034.
- [4] A. Barth, C. Schwab, N. Zollinger, Multi-level Monte Carlo finite element method for elliptic PDEs with stochastic coefficients, Numer. Math. 119 (1) (2011) 123–161.
- [5] C. Bierig, A. Chernov, Estimation of arbitrary order central statistical moments by the multilevel Monte Carlo method, Stoch. Partial Differ. Equ., Anal. Computat. 4 (1) (2016) 3–40.
- [6] C. Carstensen, An adaptive mesh-refining algorithm allowing for an h 1 stable l 2 projection onto courant finite element spaces, Constr. Approx. 20 (4) (Oct. 2003) 549–564.
- [7] C. Carstensen, M. Feischl, M. Page, D. Praetorius, Axioms of adaptivity, Comput. Math. Appl. 67 (6) (Apr. 2014) 1195–1253.
- [8] 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.
- [9] G. Da Prato, J. Zabczyk, Encyclopedia of Mathematics and Its Applications: Stochastic Equations in Infinite Dimensions, Series Number 152, 2 edition, Cambridge University Press, Cambridge, England, Apr. 2014.
- [10] W. Dörfler, A convergent adaptive algorithm for Poisson's equation, SIAM J. Numer. Anal. 33 (3) (1996) 1106–1124.
- [11] M. Eigel, C. Merdon, J. Neumann, An adaptive multilevel Monte Carlo method with stochastic bounds for quantities of interest with uncertain data, SIAM/ASA J. Uncertain. Quantificat. 4 (1) (2016) 1219–1245.
- [12] H.C. Elman, J. Liang, T. Sánchez-Vizuet, Surrogate approximation of the Grad-Shafranov free boundary problem via stochastic collocation on sparse grids, J. Comput. Phys. 448 (2022) 110699.
- [13] H.R. Fairbanks, A. Doostan, C. Ketelsen, G. Iaccarino, A low-rank control variate for multilevel Monte Carlo simulation of high-dimensional uncertain systems, J. Comput. Phys. 341 (2017) 121–139.
- [14] B. Faugeras, H. Heumann, FEM-BEM coupling methods for Tokamak plasma axisymmetric free-boundary equilibrium computations in unbounded domains, J. Comput. Phys. 343 (2017) 201–216.
- [15] S. Funken, D. Praetorius, P. Wissgott, Efficient implementation of adaptive P1-FEM in Matlab (extended preprint). Technical report, Institute for Analysis and Scientific Computing, Vienna University of Technology, 2008.
- [16] S. Funken, D. Praetorius, P. Wissgott, Efficient implementation of adaptive P1-FEM in Matlab, Comput. Methods Appl. Math. 11 (4) (2011) 460–490.
- [17] G. Gatica, G. Hsiao, The uncoupling of boundary integral and finite element methods for nonlinear boundary value problems, J. Math. Anal. Appl. 189 (2) (1995) 442–461.
- [18] R. Ghanem, P. Spanos, A stochastic Galerkin expansion for nonlinear random vibration analysis, Probab. Eng. Mech. 8 (3–4) (1993) 255–264.
- [19] M.B. Giles, Multilevel Monte Carlo methods, Acta Numer. 24 (2015) 259–328.
- [20] P.-A. Gourdain, J.-N. Leboeuf, R. Neches, High-resolution magnetohydrodynamic equilibrium code for unity beta plasmas, J. Comput. Phys. 216 (1) (July 2006) 275–299.
- [21] H. Grad, H. Rubin, Hydromagnetic equilibria and force-free fields, in: Proc. Second International Conference on the Peaceful Uses of Atomic Energy, Geneva, New York, vol. 31, United Nations, Oct 1958, p. 190.
- [22] V. Grandgirard, Modelisation de l'équilibre d'un plasma de tokamak, Technical report, Université de Franche-Comté, 1999.
- [23] H. Heumann, FEEQS.M, <http://www-sop.inria.fr/members/Holger.Heumann/Software.html>.
- [24] H. Heumann, J. Blum, C. Boulbe, B. Faugeras, G. Selig, J.-M. Ané, S. Brémont, V. Grandgirard, P. Hertout, E. Nardon, et al., Quasi-static free-boundary equilibrium of toroidal plasma with CEDRES++: computational methods and applications, J. Plasma Phys. 81 (3) (2015) 905810301.
- [25] H. Heumann, F. Rapetti, A finite element method with overlapping meshes for free-boundary axisymmetric plasma equilibria in realistic geometries, J. Comput. Phys. 334 (2017) 522–540.
- [26] H. Hoel, E. von Schwerin, A. Szepessy, R. Tempone, Adaptive multilevel Monte Carlo simulation, in: Numerical Analysis of Multiscale Computations, in: Lect. Notes Comput. Sci. Eng., vol. 82, Springer, Heidelberg, 2012, pp. 217–234.
- [27] H. Hoel, E. von Schwerin, A. Szepessy, R. Tempone, Implementation and analysis of an adaptive multilevel Monte Carlo algorithm, Monte Carlo Methods Appl. 20 (1) (2014) 1–41.
- [28] F. Hofmann, FBT - a free-boundary tokamak equilibrium code for highly elongated and shaped plasmas, Comput. Phys. Commun. 48 (2) (Feb. 1988) 207–221.
- [29] G.C. Hsiao, S. Zhang, Optimal order multigrid methods for solving exterior boundary value problems, SIAM J. Numer. Anal. 31 (3) (1994) 680–694.
- [30] J.D. Jackson, Classical Electrodynamics, 2nd edition, Wiley, New York, NY, 1975.
- [31] J. Johnson, H. Dalhed, J. Greene, R. Grimm, Y. Hsieh, S. Jardin, J. Manickam, M. Okabayashi, R. Storer, A. Todd, D. Voss, K. Weimer, Numerical determination of axisymmetric toroidal magnetohydrodynamic equilibria, J. Comput. Phys. 32 (2) (Aug. 1979) 212–234.
- [32] A. Khodadadian, M. Parvizi, C. Heitzinger, An adaptive multilevel Monte Carlo algorithm for the stochastic drift-diffusion-Poisson system, Comput. Methods Appl. Mech. Eng. 368 (2020) 113163.
- [33] R. Kornhuber, E. Youett, Adaptive multilevel Monte Carlo methods for stochastic variational inequalities, SIAM J. Numer. Anal. 56 (4) (2018) 1987–2007.
- [34] S. Larsson, V. Thomée, Partial Differential Equations with Numerical Methods, Texts in Applied Mathematics, vol. 45, Springer-Verlag, Berlin, 2003.
- [35] F. Law, A. Cerfon, B. Pebersdorfer, Accelerating the estimation of collisionless energetic particle confinement statistics in stellarators using multifidelity Monte Carlo, Nucl. Fusion 62 (2022) 076019.
- [36] R. Lüst, A. Schlüter, Axialsymmetrische magnetohydrodynamische Gleichgewichtskonfigurationen, Z. Naturforsch. 12a (1957) 850–854.
- [37] J. Luxon, B. Brown, Magnetic analysis of non-circular cross-section tokamaks, Nucl. Fusion 22 (6) (jun 1982) 813–821.
- [38] N. Metropolis, S. Ulam, The Monte Carlo method, J. Am. Stat. Assoc. 44 (1949) 335–341.

- [39] P. Morin, R.H. Nochetto, K.G. Siebert, Data oscillation and convergence of adaptive FEM, *SIAM J. Numer. Anal.* 38 (2) (Jan. 2000) 466–488.
- [40] M. Motamed, D. Appelö, A multiorder discontinuous Galerkin Monte Carlo method for hyperbolic problems with stochastic parameters, *SIAM J. Numer. Anal.* 56 (1) (2018) 448–468.
- [41] F. Nobile, F. Tesei, A multi level Monte Carlo method with control variate for elliptic PDEs with log-normal coefficients, *Stoch. Partial Differ. Equ., Anal. Computat.* 3 (3) (2015) 398–444.
- [42] V.D. Shafranov, On magnetohydrodynamical equilibrium configurations, *Sov. Phys. JETP* 6 (1958) 545–554.
- [43] J.R. Shewchuk, Delaunay refinement algorithms for triangular mesh generation, *Comput. Geom.* 22 (1–3) (May 2002) 21–74.
- [44] A.L. Teckentrup, R. Scheichl, M.B. Giles, E. Ullmann, Further analysis of multi-level Monte Carlo methods for elliptic PDEs with random coefficients, *Numer. Math.* 125 (3) (Mar. 2013) 569–600.
- [45] B.P. Welford, Note on a method for calculating corrected sums of squares and products, *Technometrics* 4 (3) (1962) 419–420.
- [46] D. Xiu, *Numerical Methods for Stochastic Computations: A Spectral Method Approach*, Princeton University Press, Princeton, NJ, 2010.