

Space-time multilevel Monte Carlo methods and their application to cardiac electrophysiology



S. Ben Bader^{a,*}, P. Benedusi^a, A. Quaglino^{a,b}, P. Zulian^a, R. Krause^a

^a Center for Computational Medicine in Cardiology, Institute of Computational Science, Università della Svizzera italiana, Lugano, Switzerland

^b NNAISENSE SA, Lugano, Switzerland

ARTICLE INFO

Article history:

Available online 5 February 2021

Keywords:

Uncertainty Quantification
Multilevel methods
Space-time finite elements (3D+1)
Cardiac electrophysiology
Monodomain equation

ABSTRACT

We present a novel approach aimed at high-performance uncertainty quantification for time-dependent problems governed by partial differential equations. In particular, we consider input uncertainties described by a Karhunen-Loève expansion and compute statistics of high-dimensional quantities-of-interest, such as the cardiac activation potential. Our methodology relies on a close integration of multilevel Monte Carlo methods, parallel iterative solvers, and a space-time discretization. This combination allows for space-time adaptivity, time-changing domains, and to take advantage of past samples to initialize the space-time solution. The resulting sequence of problems is distributed using a multilevel parallelization strategy, allocating batches of samples having different sizes to a different number of processors. We assess the performance of the proposed framework by showing in detail its application to the solution of nonlinear equations arising from cardiac electrophysiology. Specifically, we study the effect of spatially-correlated perturbations of the heart fibers' conductivities on the mean and variance of the resulting activation map. As shown by the experiments, the theoretical rates of convergence of multilevel Monte Carlo are achieved. Moreover, the total computational work for a prescribed accuracy is reduced by an order of magnitude with respect to standard Monte Carlo methods.

© 2021 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Many phenomena in science, engineering, and medicine can be modeled as initial boundary value problems for an unknown function. Their numerical solution is a computationally demanding task. In addition to that, the input data or the parameters of these models might not be known exactly, i.e., they are subject to a possibly large uncertainty. In order to guarantee that the outcome of simulations can be relied upon, it is therefore imperative to quantify the effects of such input uncertainties on the output quantities-of-interest. This process is called forward propagation and is part of the field of Uncertainty Quantification (UQ). This can be achieved via, e.g. sampling methods by solving, numerous times, a computationally expensive problem. This setting requires that one should not only rely on huge computational resources (e.g. parallel implementation on supercomputers), but also on mathematical state-of-the-art techniques for reducing the computational load.

In the case of cardiac electrophysiology, it has been well established that the heart muscle-fibers play a major role in the electrical conductivity. From CT and MRI imaging, the fibers orientations as well as the conductivities they carry are

* Corresponding author.

E-mail address: seif.ben.bader@usi.ch (S. Ben Bader).

extremely hard to determine precisely. Moreover, these can vary from one patient to another, leading to huge uncertainties if standard orientations or conductivities are used in patient-specific simulations. In this work, we focus only on fiber conductivities and model it as a spatially-correlated random perturbation applied to the diffusion field. In this context, a single patient-tailored simulation can take several hours on a large cluster, for example when the monodomain model is employed [26]. It follows that, e.g., UQ for such models is currently unfeasible, or extremely inefficient, with plain Monte Carlo methods.

A standard approach to overcome these difficulties is to use model reduction strategies, such as reduced or surrogate models [7,19]. These approximations are built upon the observation that in many applications it is clearly possible to distinguish between *offline* and *online* phases of the workflow, where the former is typically very expensive and encompasses everything that can be precomputed in advance, e.g., training a surrogate model before any patient-specific data becomes available, while the latter is very cheap and consists only of model evaluations. Unfortunately, this approach has two significant drawbacks. On the one hand, complex error estimates must be provided to ensure that the approximation error is within acceptable bounds. On the other hand, the offline training should be performed only once, which is unrealistic in cases where, e.g., the geometry is not known a priori. Alternative to that, physics-based reduction strategies which bypass the precomputation are also possible. In the case of electrophysiology, the activation time can be computed via the eikonal equation, which provides a physiologically-meaningful solution [31], strongly linked to the bidomain equation [6].

A more recent idea is to allow reduced models to be inaccurate, in the sense of not providing an approximation within certain error bounds. Accuracy is replaced by correlation [29] or even only statistical dependence [22] with the high-fidelity model. This approach has a twofold advantage. On the one hand, it is the statistical dependence, rather than the error bounds of coarse models based on coarse meshes, that is crucial to ensure that propagating uncertainties via the low-fidelity models provide useful information on the statistics of the high-fidelity quantity-of-interest. On the other hand, low-fidelity models are not restricted to low-resolution geometries, therefore solving them can be several orders of magnitude faster than a coarse model. However, some computational effort (~ 100 runs of the high-fidelity model) is needed for estimating the statistical dependence between the models. Moreover, there is no established approach for treating vector-valued quantities-of-interest [28], which may require a significantly different number of samples of each model for each vector component. An approach based on minimizing a norm of the error was recently proposed in [34].

In the specific context of electrophysiology, the use of Bayesian multifidelity methods has been advocated in [33], where a Bayesian regression mapping the low- and high-fidelity outputs is created. The Bayesian nature of this approach error automatically augments the estimate of the solution statistics with full probability distributions and credible intervals, obtained independently of the degree of statistical dependence between models and the number of samples employed. This makes it a very suitable approach for scenarios where resources are scarce, such as clinical applications. However, this approach is limited to scalar or low-dimensional quantities-of-interest. Moreover, it does not yield asymptotic convergence to point estimates as the number of samples tends to infinity, so it is unsuitable for scenarios where accurate estimates are desired.

All of the above methodologies are characterized by a clear segregation among sampling, models, discretizations, and solution methods. Here, we take the opposite route and intertwine the different components. The central point of our approach is to consider the multilevel Monte Carlo method [13], which extends the idea of control variables [9]. The key idea is to perform most of the simulations on a sequence of low-resolution models. Similarly to the case of multifidelity, the direct use of the highest-resolution level guarantees convergence, while a significant portion of the computational load is offset to the low-resolution hierarchy. However, in this case, no offline preprocessing is needed to find the correlation across models and vector-valued outputs are naturally handled by the framework. In order to allow for arbitrary coarsening in space and time, we further rely on a space-time discretization of the monodomain equation. By doing so, it also becomes possible to use the already computed samples for the initialization of Newton's method, which solves the space-time equations, thereby significantly reducing the amount of iterations required for convergence.

The outline of the paper is as follows: Section 2 introduces the electrophysiology equations and the uncertainty related to the fiber conductivities, Section 3 details the proposed method using the heat equation as a model problem, and Section 4 presents the results from the numerical experiments.

2. Modeling the uncertainty in the fiber conductivities

2.1. The deterministic monodomain equation

The heart assumes a pumping function that is the result of a very complex contraction and relaxation cycle occurring in the cardiac cells. This process is controlled by a non-trivial pattern of electrical activation which is at the core of the heart function. Indeed, several heart diseases are closely related to disturbances of the electrical activity. As a consequence, understanding and modeling this activity is important for a better clinical diagnosis and treatment of patients.

The propagation of the electrical potential inside the cardiac muscle is initiated in the sinoatrial node, that is on top of the left and right atriums. It generates a stimulus that gives rise to a traveling wave through the heart. Its cells have the ability to respond actively to this electrical stimulus through voltage-gated ion channels.

Starting from the first models for ionic channels [20], a large family of models has been developed by now, describing the electrical activity at the subcellular level. Combined with diffusion in space, the so-called monodomain equation has been

derived [21,25] for the electrical propagation in the cardiac tissue. The monodomain equation is written as a non-stationary reaction-diffusion equation

$$\begin{aligned} \frac{\partial u(\mathbf{x}, t)}{\partial t} - \nabla \cdot (G(\mathbf{x}) \nabla u(\mathbf{x}, t)) + I_{\text{ion}}(u(\mathbf{x}, t)) &= I_{\text{app}}(\mathbf{x}, t), \quad \forall (\mathbf{x}, t) \in D \times (0, T], \\ G(\mathbf{x}) \nabla u(\mathbf{x}, t) \cdot \mathbf{n}(\mathbf{x}) &= 0, \quad \forall (\mathbf{x}, t) \in \partial D \times (0, T], \\ u(\mathbf{x}, 0) &= 0, \quad \forall \mathbf{x} \in D, \end{aligned} \quad (1)$$

where $u = u(\mathbf{x}, t)$ is the electrical potential, $D \subset \mathbb{R}^d$ a domain representing the heart, $T \in \mathbb{R}^+$ the end time, $I_{\text{app}} : D \times [0, T] \rightarrow \mathbb{R}$ an applied stimulus modeling the sinoatrial node activation, and $I_{\text{ion}} : \mathbb{R} \rightarrow \mathbb{R}$ an ion channel model. We rely here on the Fitz-Hugh Nagumo (FHN) model [10], for which we have:

$$I_{\text{ion}}(u) := \alpha(u - u_{\text{rest}})(u - u_{\text{th}})(u - u_{\text{peak}}). \quad (2)$$

The values u_{rest} , u_{th} and u_{peak} are characteristic potential values representing the action potential behavior that a heart cell goes through during its activation. These are respectively the resting potential u_{rest} (cell is unactivated), the threshold potential u_{th} (cell is triggered) and the peak value u_{peak} (cell is activated). Here, α is a non-negative scaling parameter.

Finally, G in (1) represents the conductivity tensor modeling the fibers, along which electrical potential propagation takes place, all over the heart muscle. The tensor G plays a major role in our model as it may represent a source of uncertainty with regards to the fibers orientations and conductivities it models. We here assume that G is a scalar isotropic diffusion field and focus exclusively on the fiber conductivities.

2.2. The stochastic monodomain equation

The mathematical and numerical instruments developed in decades of research in the electrophysiology field allow in principle for virtual therapy planning. The monodomain equation is by now an established model in cardiac electrophysiology, describing with high accuracy the electrical activity in the myocardium. Nonetheless, patient-specific simulations are still not employed as a routine tool in the treatment of patients.

A particular reason for this can be found in the data which is acquired in clinical practice. For instance, the fiber conductivities still represent a great challenge to be determined from an imaging point of a view. One should therefore account for possible uncertainties of the diffusion field G in (1).

Modeling the uncertainty requires to account for an additional stochastic variable in the formulation of the monodomain equation. Let us denote this stochastic variable with $\omega \in \Omega$ where Ω is the set of all possible outcomes, which in our application would represent the set of all diffusion fields modeled by the uncertainty. We are now interested in estimating the statistics of $u(\mathbf{x}, t, \omega)$, i.e. $\mathbb{E}[u(\mathbf{x}, t)] = \int_{\Omega} u(\mathbf{x}, t, \omega) d\mathbb{P}(\omega)$, as a solution to the monodomain stochastic PDE, for which the main equation reads: for almost every $\omega \in \Omega$:

$$\frac{\partial u(\cdot, \omega)}{\partial t} - \nabla \cdot (G(\cdot, \omega) \nabla u(\cdot, \omega)) + I_{\text{ion}}(u(\cdot, \omega)) = I_{\text{app}}(\cdot), \quad \text{in } D \times (0, T], \quad (3)$$

where we omitted the space and time variables (\mathbf{x}, t) for the sake of readability. For evaluating $\mathbb{E}[u(\mathbf{x}, t)]$ given (3), we rely on Monte-Carlo and Multilevel Monte-Carlo techniques as described in Section 3. These methods imply a deterministic reformulation of (3), and solving the monodomain equation for different independent and identically distributed samples. Naturally, parallelization plays a major role in the performance of both methods, given the important number of samples to consider.

3. Space-time multilevel Monte Carlo

3.1. Multilevel Monte Carlo

As introduced in the previous section, we are interested in equations that are subject to uncertain or unreliable data in the aim of evaluating the information deriving from the process of simulations. This concept had been summarized under the terminology of *uncertainty quantification* (UQ) and is concerned about extracting robust and qualitative information despite the presence of this variability. In the next subsection, we briefly detail the problem we intend to solve, and introduce the reader to the notation we rely on.

3.1.1. Problem setting and Uncertainty modeling

From a mathematical point of view, the underlying task resides in computing the solution of a PDE with coefficients depending on a random input belonging to a stochastic space. For the problem (3) to be well-defined, we need to introduce the following setting.

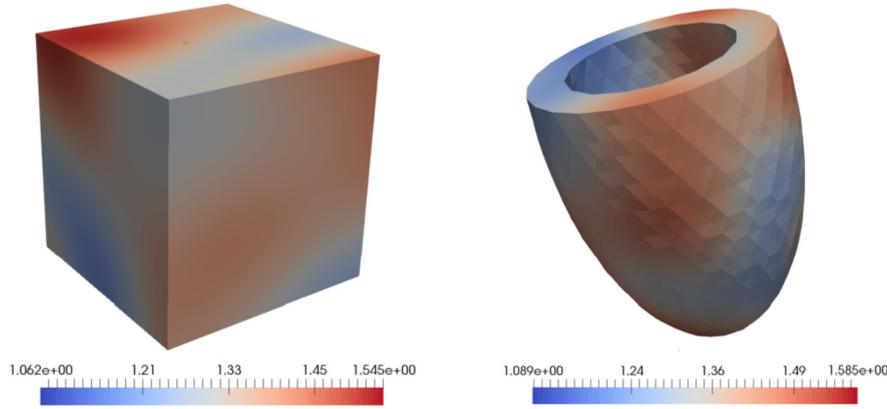


Fig. 1. Example of random fields for cube and ventricle geometries. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Let $(\Omega, \Sigma, \mathbb{P})$ be a complete and separable probability space, with $\Sigma \subset 2^\Omega$ a σ -field and \mathbb{P} a probability measure. We consider a bounded, Lipschitz domain $D \subset \mathbb{R}^d$ with $d = 1, 2, 3$. We also require the Bochner space

$$L^2(\Omega, \Sigma, \mathbb{P}; B) := \{v : \Omega \rightarrow B \mid v \text{ strongly measurable such that } \|v\|_{L^2(\Omega; B)} < \infty\},$$

where $B = L^2([0, T]; H^1(D))$ is itself a Bochner space ($[0, T]$ denotes the time domain) and

$$\|v\|_{L^2(\Omega; B)} := \left(\int_{\Omega} \|v(\cdot, \omega)\|_B^2 d\mathbb{P}(\omega) \right)^{1/2}. \tag{4}$$

The norm on the Bochner space $B = L^2([0, T]; H^1(D))$ is in turn defined as

$$\|v(\cdot, \omega)\|_B := \left(\int_{[0, T]} \int_D v(x, t, \omega)^2 dx dt \right)^{1/2}. \tag{5}$$

For the sake of brevity, we will always refer to $L^2(\Omega, \Sigma, \mathbb{P}; B)$ with $L^2(\Omega; B)$. Having everything at hand, we introduce a simplified equation model of (3), given by

find $\tilde{u} \in L^2(\Omega; L^2([0, T]; H^1(D)))$ such that for almost every $\omega \in \Omega$,

$$\partial_t \tilde{u}(X, \omega) - \operatorname{div}(G(X, \omega) \nabla \tilde{u}(X, \omega)) = f(X) \text{ in } D \times [0, T], \tag{6}$$

where $X = (\mathbf{x}, t)$ and $G(X, \omega)$ is the diffusion field subject to the uncertainty $\omega \in \Omega$, that can be interpreted as a particular diffusion field from all the possible ones we allow for the problem for. Notice that if the forcing function f was depending on $\tilde{u}(X)$ and written as $f(X) = I_{\text{app}}(X) - I_{\text{ion}}(\tilde{u}(X, \omega))$, we recover the stochastic monodomain equation (3). We will however later on for error estimates (cf. subsection 3.1.2) rely on a classical forcing function $f(X)$, independent from \tilde{u} , treating therefore the convergence rates for the quadrature methods on the stochastic heat equation. The PDE further encodes suitable boundary and initial conditions. We are interested in computing the statistics of the solution u to the stochastic PDE (6), i.e. $\mathbb{E}[\tilde{u}](X) = \int_{\Omega} \tilde{u}(X, \omega) d\mathbb{P}(\omega) \in L^2([0, T]; H^1(D))$.

More specifically, every $\omega \in \Omega$ expresses a different diffusion field $G(X, \omega)$. In our particular case, it is more convenient to think of these different diffusion fields as a superposition of a mean diffusion field $G_0(X)$ and a random perturbation all over the domain D (random field).

We furthermore assume the diffusion uncertainties to be spatially correlated, since neighboring regions are more likely to have the same conductivity in a healthy heart. We therefore talk about spatially correlated random fields [17]. Examples of random fields of cube and idealized ventricle geometries can be found in Fig. 1.

The following are obtained by expressing the stochastic diffusion coefficient $G : D \times \Omega \rightarrow \mathbb{R}$ as a truncated Karhunen-Loève (KL) expansion [33]

$$G(\mathbf{x}; \omega) = G_0(\mathbf{x}) + s \sum_{k=1}^M \sqrt{\lambda_k} \phi_k(\mathbf{x}) \psi_k(\omega), \tag{7}$$

where λ_k and $\phi_k(\mathbf{x})$ are respectively eigenvalues and eigenvectors of the covariance matrix induced by a spatial correlation kernel k . Regarding the KL truncation, the stochastic dimension M is obtained by low rank approximating the covariance matrix with pivoted Cholesky decomposition [15,17]. The $\{\psi_k(\omega)\}_{k=1}^M$ are stochastically uncorrelated random variables of

the joint density function $\rho(\mathbf{y}) = \prod_{k=1}^M \rho_k(y_k)$. We use here a uniform distribution, i.e. $\psi_i \sim \mathcal{U}(-1, 1)$. We are therefore able to identify the stochastic space Ω with its image $[-1, 1]^M$ using the mapping

$$\psi : \Omega \rightarrow [-1, 1]^M, \tag{8}$$

$$\omega \mapsto (\psi_1(\omega), \dots, \psi_M(\omega)). \tag{9}$$

Likewise by inserting a realization $\mathbf{y} = (y_1, \dots, y_M) = (\psi_1(\omega), \dots, \psi_M(\omega))$, we could reformulate the truncated KL expansion in a deterministic form for $G : D \times [-1, 1]^M \rightarrow \mathbb{R}$ as

$$G(\mathbf{x}, \mathbf{y}) = G_0(\mathbf{x}) + s \sum_{k=1}^M \sqrt{\lambda_k} \phi_k(\mathbf{x}) y_k. \tag{10}$$

Finally, $s > 0$ in (10) is chosen such that for all $x \in D$

$$s \left\| \sum_{i=1}^M \sqrt{\lambda_i} \phi_k(\mathbf{x}) y_k \right\|_{L^\infty([-1, 1]^M)} \leq \delta G_0(\mathbf{x}),$$

for $\delta = 0.5$, in order to enforce the diffusion coefficient to be uniformly elliptic and bounded, i.e. for all $\mathbf{x} \in D$ and for all $\mathbf{y} \in [-1, 1]^M$

$$0 < G_{\min}(\mathbf{x}) \leq G(\mathbf{x}, \mathbf{y}) \leq G_{\max}(\mathbf{x}) < \infty.$$

The arguments above allow also for rewriting the stochastic equation (6) in a parametric form: find $u \in L^2([-1, 1]^M; L^2([0, T]; H^1(D)))$

$$\partial_t u(X, \mathbf{y}) - \operatorname{div}(G(X, \mathbf{y}) \nabla u(X, \mathbf{y})) = f(X) \text{ in } D \times [0, T], \tag{11}$$

for all $\mathbf{y} \in [-1, 1]^M$.

The expectation of the solution for the stochastic heat equation (11) can be rewritten as a high dimensional integral

$$\mathbb{E}[u](X) = \int_{[-1, 1]^M} u(X, \mathbf{y}) \rho(\mathbf{y}) d\mathbf{y} \approx \int_{\Omega} \tilde{u}(X, \omega) d\mathbb{P}(\omega) = \mathbb{E}[\tilde{u}](X). \tag{12}$$

3.1.2. Methods and error estimates

In order to numerically approximate the integral (12), one should rely on a quadrature formula of the type

$$\mathcal{Q}u(\cdot) = \sum_{i=1}^{N_I} w_i u(\cdot, \xi_i) \rho(\xi_i). \tag{13}$$

Approximating a stochastic integral likewise can further be understood as a sampling method, where the samples represent the quadrature points. A well known sampling method is the Monte-Carlo estimator. It is characterized by a fairly easy implementation and is straightforward to use, but it has a low convergence rate and hence requires a large number of samples, i.e. simulations. In this work, we consider a multilevel approach as an alternative to standard Monte-Carlo.

The Monte-Carlo (MC) estimator represents a direct method for calculating the solution to the parametric problem (12). Its quadrature formula can be expressed as

$$\mathcal{Q}_{MC}u(\cdot) = N_I^{-1} \sum_{i=1}^{N_I} u^i(\cdot), \tag{14}$$

i.e. the mean over solution samples u^i , $i = 1, \dots, N_I$ corresponding to N_I independent and identically distributed realization.

The definition of the MC estimator (14) would be completely auto-sufficient if a solution u^i to the sample i could be computed exactly. However, in the context the monodomain equation, an exact and analytical solution can not be provided and we rely on finite elements to numerically approximate it. Let us consider a nested sequence of shape regular and quasi-uniform tetrahedralizations $\{\mathcal{T}_l\}_{l \geq 0}$ of the spatial domain D , each of mesh size $h_l \sim 2^{-l}$. We define the continuous, piecewise linear finite element spaces

$$\mathcal{S}_l = \{v_l \in C^0(D) : v_l|_T \in \mathbb{P}_1(T), \forall T \in \mathcal{T}_l\}.$$

We now introduce the notation

$$\mathcal{Q}_{MC,l}u(\cdot) = N_I^{-1} \sum_{i=1}^{N_I} u_l^i(\cdot), \tag{15}$$

where u_i^j is the finite element approximation of the solution u^i for the sample i , lying in the space S_j . Therefore, the solution obtained through the MC method is subject not only to the stochastic error depending on the number of samples used as quadrature points, but also on the discretization error inherent to solving the PDE within the context of finite elements.

The stochastic error of the MC quadrature integration is bounded by [3]

$$\|\mathbb{E}[u] - \mathcal{Q}_{MC}u\|_{L^2([-1,1]^m; \mathcal{V})} \leq N_l^{-1/2} \|u\|_{L^2([-1,1]^m; \mathcal{V})}, \tag{16}$$

where $\mathcal{V} = L^2([0, T]; H^1(D))$ with the norm defined as in (5). On the other hand, using linear finite elements and a second order method in time (e.g. Crank-Nicolson) we get the following error estimate for a H^2 -regular solution in space [1,11,14, 35]

$$\|u - u_l\|_{\mathcal{V}} \leq C(h_l^2 + \Delta t_l^2) \|f\|_{\mathcal{V}}, \tag{17}$$

for some constant $C > 0$ that depends on the initial condition and the number of time steps considered, but not on h_l and Δt_l which are respectively the spatial and time discretization parameters for a given level l . Putting (16) and (17) together results in having the general error bound [3]

$$\|\mathbb{E}[u] - \mathcal{Q}_{MC,l}u\|_{L^2([-1,1]^m; \mathcal{V})} \leq C(N_l^{-1/2} + h_l^2 + \Delta t_l^2) \|f\|_{L^2([-1,1]^m; \mathcal{V})}. \tag{18}$$

The setting for a MC estimation is therefore closely related to that of the following error estimate, as it underlines the relationship between the mesh size and the number of samples to consider on a given discretization level. Indeed, in order to not have an error that is neither dominated by discretization in stochastics or discretization in space and time, the number of samples to consider on a level l is such that

$$N_l^{-1/2} \sim \mathcal{O}(h_l^2) \implies N_l \sim \mathcal{O}(h_l^{-4}) = \mathcal{O}(2^{4l}),$$

where we assume that $h_l \sim 2^{-l}$ and $\Delta t_l \sim \Delta t_l$. The convergence rate of the MC method can therefore be derived from (18) and is given by 2^{-2l} . The reader might at this level realize the limitation of such a method, as the number of samples to be taken gets tremendously large when considering discretization levels with small h_l and Δt_l . In addition, each one of these samples also requires a more significant computational load which places the MC method in the range of numerically costly methods for PDEs requiring a high resolution solution.

The Multilevel Monte Carlo (MLMC) estimator represents an alternative for the MC estimator. We would like to solve the equation for an initially discretized domain $D \times [0, T]$, which we will refer to as being the fine discretization level L . Let us establish a hierarchy of nested coarse discretization levels $l = 0, 1, \dots, L - 1$ obtained by uniform coarsening in space and time. We define the MLMC-estimator as follows:

$$\mathcal{Q}_{MLMC,l}u(\cdot) = \sum_{l=0}^L \mathcal{Q}_{MC}^{L-l}(u_l(\cdot) - u_{l-1}(\cdot)), \tag{19}$$

where $\{\mathcal{Q}_{MC}^l\}_{l=0}^L$ is a sequence of MC quadratures with increasing precision of order 2^{-2l} [16]. We also define $u_{-1} \equiv 0$.

To clarify the above, the underlying idea about the MLMC-estimator is to perform the dominant number of MLMC samples on a coarse level $l = 0$, and to obtain a solution for level $l = L$ by adding corrections between every different successive levels $l - 1$ and l . The MLMC and its convergence rate have been discussed in details in [5,16,36] for the case of the stochastic Poisson problem. For the stochastic equation, given (16), (17) and the above specified 2^{-2l} increasing precision for the quadrature rule, a similar result can be shown:

$$\|\mathbb{E}[u] - \mathcal{Q}_{MLMC,l}u\|_{L^2([-1,1]^m; \mathcal{V})} \leq C2^{-2L} \|f\|_{L^2([-1,1]^m; \mathcal{V})}, \tag{20}$$

for some constant $C > 0$. The MLMC has therefore the same convergence rate than MC up to a logarithmic number, but is designed to be less computationally expensive.

3.2. Space-time discretization for the linear problem

The heat equation (6) can be numerically solved by means of finite elements (FE) and finite differences (FD), which are well established for such problems. A classical approach in that sense is to perform a FE discretization in space and FD in time, i.e. using a sequential time-stepping method. The parallel scalability of this approach is eventually limited by the time integration process. In fact, when the parallelization in space saturates, sequential time integration becomes the natural bottleneck for the scalability of the solution process. This can be overcome by parallel-in-time methods. Those were developed in the last decades [27] to overcome this bottleneck and enhance parallelism in both space and time. In particular, we employ a monolithic approach in space and time, where we assemble a large space-time system that is solved in parallel [24]. For a comprehensive review of parallel-in-time methods, see [12].

We use Lagrangian FE of first order for the discretization in the space variable. Assuming that the solution $u(\mathbf{x}, t)$ is sufficiently regular on D , we derive the weak formulation from (6) (in which we neglect the stochastic variable ω) using test functions $v \in H_0^1(D)$:

$$\int_D \frac{\partial u(\mathbf{x}, t)}{\partial t} v(\mathbf{x}) d\mathbf{x} + \int_D (G(\mathbf{x}) \nabla u(\mathbf{x}, t)) \nabla v(\mathbf{x}) d\mathbf{x} = \int_D f(\mathbf{x}, t) v(\mathbf{x}) d\mathbf{x}.$$

Then, we define the semi-discrete problem:

$$M_h \frac{\partial \mathbf{u}(t)}{\partial t} + K_h \mathbf{u}(t) = \mathbf{f}(t). \tag{21}$$

Here, $M_h \in \mathbb{R}^{n \times n}$ and $K_h \in \mathbb{R}^{n \times n}$ are the standard mass and stiffness matrices obtained using n linear nodal basis functions $\{\psi_i\}_{i=1}^n \subset \mathbb{P}_1$, i.e.

$$M_h := \left[\int_D \psi_i(\mathbf{x}) \psi_j(\mathbf{x}) d\mathbf{x} \right]_{i,j=1}^n, \quad K_h := \left[\int_D (G(\mathbf{x}) \nabla \psi_i(\mathbf{x})) \nabla \psi_j(\mathbf{x}) d\mathbf{x} \right]_{i,j=1}^n, \tag{22}$$

$$\mathbf{f}(t) := \left[\int_D f(\mathbf{x}, t) \psi_i(\mathbf{x}) d\mathbf{x} \right]_{i=1}^n, \tag{23}$$

arising from the approximation:

$$u(\mathbf{x}, t) \simeq \sum_{i=1}^n u_i(t) \psi_i(\mathbf{x}), \quad \text{with } \mathbf{u}(t) = [u_1(t), u_2(t), \dots, u_n(t)]^T.$$

Consider a uniform partition of the time axis $[0, T]$ in m intervals, such that $\Delta t = T/m$ and $t_k = k\Delta t$, with $k = 0, \dots, m$. We apply the second order Crank–Nicolson method for the time discretization of (21) and obtain, for $k = 0, \dots, m - 1$

$$\left(M_h + \frac{\Delta t}{2} K_h \right) \mathbf{u}_{k+1} + \left(-M_h + \frac{\Delta t}{2} K_h \right) \mathbf{u}_k = \mathbf{f}_k \quad \text{and} \quad \mathbf{u}_k := \mathbf{u}(t_k), \tag{24}$$

$$\text{with } \mathbf{f}_k := \frac{\Delta t}{2} (\mathbf{f}(t_{k+1}) + \mathbf{f}(t_k)). \tag{25}$$

If we define $A_{t,h} := M_h + \frac{\Delta t}{2} K_h$ and $B_{t,h} := -M_h + \frac{\Delta t}{2} K_h$ the system of equations (24) can be summarized in the compact form:

$$\begin{bmatrix} A_{t,h} & & & & \\ B_{t,h} & A_{t,h} & & & \\ & \ddots & \ddots & \ddots & \\ & & B_{t,h} & A_{t,h} & \end{bmatrix} \begin{pmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_m \end{pmatrix} = \begin{pmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \\ \vdots \\ \mathbf{f}_m \end{pmatrix} \iff \mathbf{C} \mathbf{u} = \mathbf{f}, \tag{26}$$

where $\mathbf{C} \in \mathbb{R}^{nm \times nm}$ is a large space-time system that can be distributed and solved in parallel and

$$\mathbf{u} := [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m]^T \quad \text{and} \quad \mathbf{f} := [\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_m]^T.$$

3.3. Discretization of the nonlinear problem

The discretization of (1) is an extension of the assembly procedure described in Section 3.2. In particular, the linear system (26) is modified to contain the discretization of the non-linear reaction term I_{ion} :

$$\mathbf{C} \mathbf{u} + \mathbf{r}(\mathbf{u}) = \mathbf{f}, \tag{27}$$

where $\mathbf{r}(\mathbf{u}) \in \mathbb{R}^{nm}$ is given by

$$\mathbf{r}(\mathbf{u}) := (\Delta t I_N \otimes M_h) \mathbf{I}_{ion}(\mathbf{u}) \quad \text{with} \quad \mathbf{I}_{ion}(\mathbf{u}) := [I_{ion}(u_1), I_{ion}(u_2), \dots, I_{ion}(u_{nm})]^T,$$

where n and m are respectively the space and time degrees of freedom.

The Jacobian $\mathbf{J}(\mathbf{u}) \in \mathbb{R}^{nm \times nm}$ of the non-linear operator in left hand side of (27) is given by

$$\mathbf{J}(\mathbf{u}) = \mathbf{C} + (\Delta t I_m \otimes M_h) \cdot \text{diag}(I'_{ion}(u_1), I'_{ion}(u_2), \dots, I'_{ion}(u_{nm})).$$

3.4. Transfer of discrete fields

The transfer of information between finite element spaces in our multilevel Monte Carlo method is performed by means of both L^2 -projections and standard finite element interpolation methods. In the case when mesh hierarchies are generated by refinement, the elements of the refined mesh form partitions of the elements of the original coarse mesh. This allows us to employ standard finite element interpolation for transferring from the coarse space to the refined space. However, if the fine mesh is not generated by refinement we can use L^2 -projections for the same task. For transferring discrete fields from the fine space to a coarser space we also employ the L^2 -projection. L^2 -projections are proven to be optimal and stable and in general superior to interpolation [18]. In particular, we employ a local approximation of the L^2 -projection, which is constructed by exploiting the properties of the dual basis [30,37]. This local approximation allows us to construct the transfer operator explicitly in such a way that it can be applied by means of a simple sparse matrix-vector multiplication. In a parallel computing environment, where meshes are arbitrarily distributed, the assembly of the mass matrices related to L^2 -projection is not trivial. The construction of the discrete L^2 -projection requires us to detect and compute intersections between the elements of the coarse mesh and the elements of the finer mesh which might be stored in different memory address spaces (e.g., on a super-computing cluster). For this purpose we use the parallel tree-search algorithms and assembly routines described in [23].

In this work we extend the above to multilevel space-time discretizations. By exploiting the tensor-product structure of (26) we can simplify the implementation and benefit from better computational performance. As before, let us assume that the spatial mesh does not change at each time-step (i.e., no adaptive mesh refinement). In this case, the tensor-product structure of the space-time grid allows us to construct space-time operators in a convenient way which requires the assembly of the spatial transfer operator to be performed only once.

For $l = 0, \dots, L - 1$ let $\mathbf{u}^l \in \mathbb{R}^{n_l m_l}$ be the coefficients of the discrete space-time solution on level l , where n_l (resp. m_l) is the number of spatial nodes (resp. time steps) on l . With $P_{l,h} \in \mathbb{R}^{n_l \times n}$ we denote the spatial discrete representations of the L^2 -projection from the fine level to a coarser level l , and with $P_{l,t} \in \mathbb{R}^{m_l \times m}$ the temporal restriction operator, obtained by standard bisection. The space-time transfer matrix $P_l \in \mathbb{R}^{n_l m_l \times nm}$ is constructed with the following tensor product

$$P_l = P_{l,t} \otimes P_{l,h},$$

and we write $\mathbf{u}^l = P_l \mathbf{u}$. The space-time interpolation (prolongation) operator from any coarse level $l - 1$ to a finer level l is constructed in the same way, using the tensor-product between transfer operators.

3.5. Parallel solver

The non linear problem (27) is solved with Newton's method. For each of the arising linear problems in the form (26), we employ a space-time parallel PGMRES, with the preconditioner given by an ILU(0) factorization within the PETSc [2] framework. If multiple processors are used, the corresponding block-Jacobi preconditioner is employed, and the PGMRES is applied, in parallel, to the local blocks. The usage of the latter PGMRES is motivated by the spectral analysis of the space-time system in (26), carried out in [4]. As default by PETSc, the GMRES solver is restarted after 30 iterations. It is important to recall that the first iterate plays an important role for the convergence of Newton's method. For this reason we use the unperturbed reference solution as initial guess for each Monte-Carlo sample. This guarantees fast convergence in our setting.

3.6. Parallelization strategy

The MLMC method developed above reduces undoubtedly the computational complexity that is carried by a naive MC approach. However, to fully exploit the multilevel algorithm, it is necessary to rely on a parallel environment that takes advantage of all different layers under which concurrent work can be executed. The MLMC requires computing samples at different discretization levels, requiring a different amount of resources. Therefore, we could think of initially distributing the required amount of work for every level as a first parallelization layer. On every one of these levels, multiple samples need to be computed. Those are completely independent and we can consider distributing work across samples as a second parallelization layer. The third layer, in the context of finite elements, would naturally be a parallelization across the tempo-spatial grid, using space-time parallel method.

We therefore rely on a so-called three level parallelization strategy largely inspired by [8]. We furthermore adapt it to standard allocation procedures on current High Performance Computing (HPC) system, in which a request for a small amount of resources is more likely to be granted in a shorter amount of time. We issue consequently different job calls involving different amount of resources for every parallel task consisting of batches of samples to solve (cf. Fig. 2). As an example, let us suppose we have an MLMC setting of $L + 1$ levels. We start with initiating p_0, \dots, p_L processes carrying the tasks for every one of these levels. The samples to be computed for every level l are equally distributed in batches over the p_l processes created for that level. The different number of samples N_l on each level and the flexibility in choosing the number of processes p_l are such that the batch size may vary across the levels, as shown in Fig. 3. Every sample is further solved in parallel with domain decomposition techniques, fulfilling the third parallelization layer. Details about the parallel solver we relied on are provided at the end of subsection 2.2 after introducing our application, i.e. the monodomain

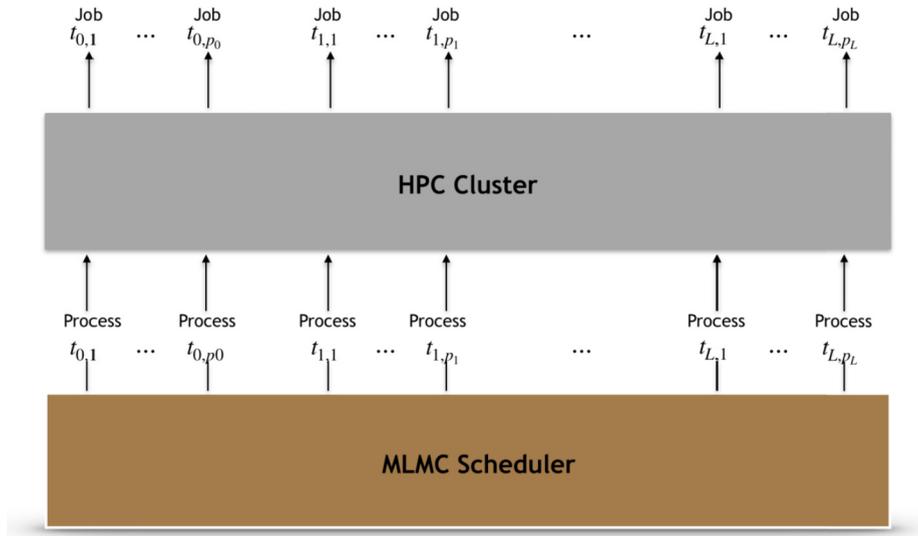


Fig. 2. MLMC tasks scheduling and job requests step. $t_{l,i}$ denotes the i -th task of level l .

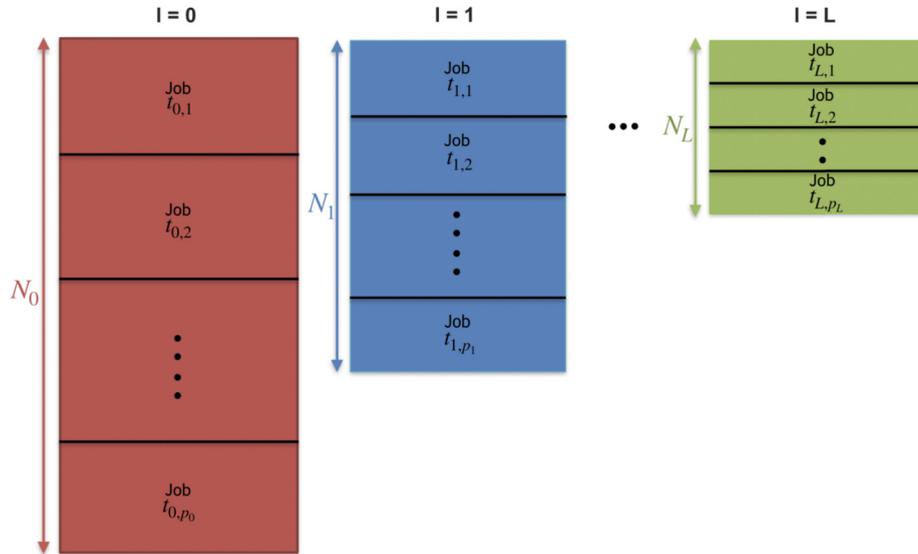


Fig. 3. The different processes executing concurrently different batches of samples.

equation. Notice that another advantage of the following cluster scheduling approach is to avoid wasting resources when the processes have different time execution.

4. Numerical experiments

The numerical experiments were realized using SLOTH [32] a UQ python library developed at the Institute of Computational Science (ICS) of Lugano. For this work, we extended it to the monodomain equation (and in general to all types of 3D+1 PDEs) by employing Utopia [38] for the finite element formulation. The experiments have been performed on the multicore partition of the supercomputer Piz Daint of the Swiss National Supercomputing Centre.

Parameters for the experiments The parameters of monodomain equation (1) used for the numerical tests are:

- $G_0(\mathbf{x}) = 3.325 \cdot 10^{-3} \text{ cm}^2 \text{ mV ms}^{-1}$ from (7) is the mean diffusion field.
- $\alpha = 1.4 \cdot 10^{-3} \text{ mV}^{-2} \text{ ms}^{-1}$, $u_{\text{rest}} = 0 \text{ mV}$, $u_{\text{th}} = 28 \text{ mV}$, and $u_{\text{peak}} = 115 \text{ mV}$ are the values for the ion channel model $I_{\text{ion}}(u)$ in (2).

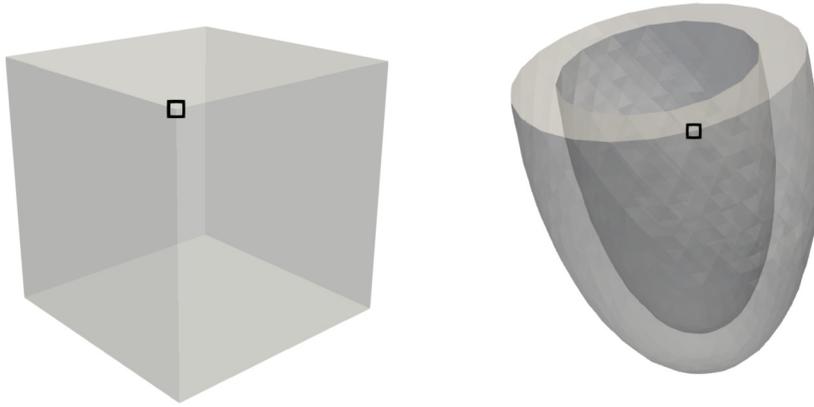


Fig. 4. Stimulus locations for the cube and the ventricle.

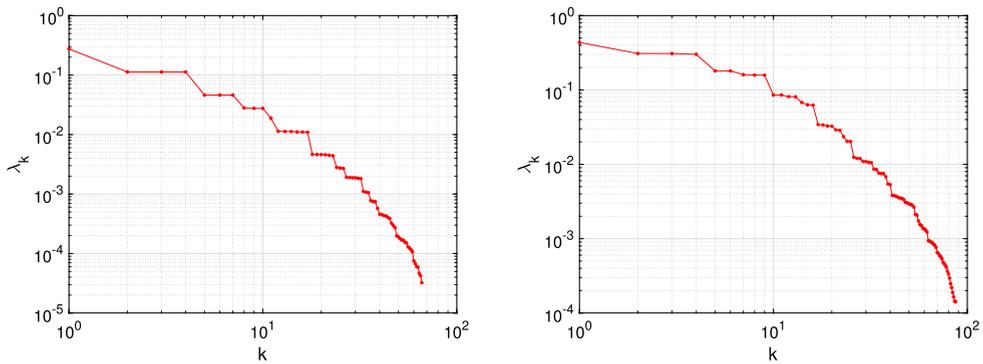


Fig. 5. KL eigenvalues decay for the cube (left) and the idealized ventricle (right).

- $I_{\text{app}}(\mathbf{x}, t) = \left(u_{\text{rest}} + u_{\text{peak}} \exp\left(-\frac{(\mathbf{x} - \mathbf{x}_0)^2}{\sigma^2}\right) \right) \chi_{[0, t_1]}(t)$ where $t_1 = \Delta t = 0.005$ ms is the function we rely on for the applied stimulus. For the cube, $\sigma = 0.5$ cm whereas for the ventricle it was set $\sigma = 1$ cm. Here, \mathbf{x}_0 is the location of stimulus and are shown for both geometries in Fig. 4.

Parameters for KL expansion We use the correlation kernel

$$k(\mathbf{x}, \mathbf{y}) = e^{-\frac{d(\mathbf{x}, \mathbf{y})^2}{\sigma_{\text{KL}}^2}},$$

for $\sigma_{\text{KL}} = 0.25$ for the cube and $\sigma_{\text{KL}} = 0.5$, while $d(\mathbf{x}, \mathbf{y})$ is the Euclidean distance between \mathbf{x} and \mathbf{y} . The *dumping factor* s in (7) for enforcing uniform ellipticity is set to $s = 0.3$ for both geometries. The low rank pivoted Cholesky decomposition on the covariance matrices induced by the correlation kernels above yielded the stochastic dimension $M = 66$ for the cube and $M = 87$ for the idealized ventricle. We furthermore report the KL eigenvalues decay in Fig. 5.

4.1. Numerical experiments in 1D+1

We intend to verify the convergence rate of both MC and MLMC as developed in Section 3. We consider an interval and use a hierarchy nested meshes, obtained by uniform coarsening in space and time. In Table 1, we report all the details regarding the considered mesh hierarchy, i.e. the number of elements, the mesh size, the number of time steps and the timestep size.

We create a reference solution, denoted as v_{ref} , by calculating the mean over 20,000 samples on the finest level $L = 5$. For the purpose of verifying the MC convergence rate, we run several MCs on each level l with $N_l \sim 2^{4l}$, and compute the difference between the expectation calculated at each level, denoted as v_l , and the reference solution. We evaluate the rooted mean square RMSE = $\mathbb{E}[\|e_l\|^2]^{1/2}$, where $e_l = |v_l - v_{\text{ref}}|$ and

$$\|e_l\|^2 = \int_{[0, T]} \|e_l(\cdot, t)\|_{L^2(D)}^2 dt = \int_{[0, T]} \int_D e_l^2(x, t) dx dt.$$

Table 1

Details about the considered Mesh hierarchy. Δt and h are the time and space discretization steps, m and n are the number of subdivisions for time and space intervals.

l	0	1	2	3	4	5
n	31	62	125	250	500	1000
h	0.032	0.016	0.008	0.004	0.002	0.001
m	4	8	16	32	64	128
Δt	0.16	0.08	0.04	0.02	0.01	0.005

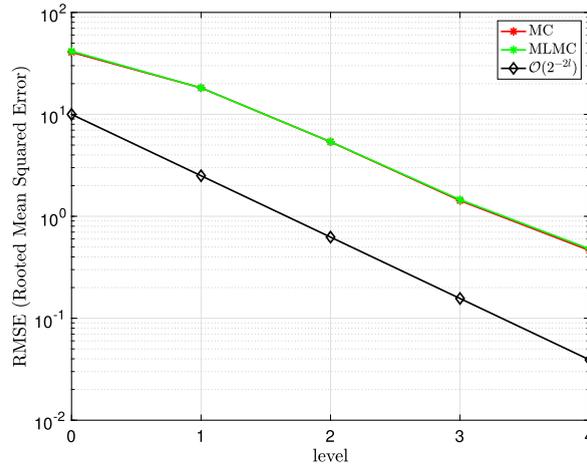


Fig. 6. Controlled convergence graphs for MC and MLMC for the 1D+1 example.

Table 2

Details about the mesh hierarchy for the cube geometry.

l	0	1	2	3
Dof's	5400	65,219	898,317	13,301,753
h	0.2	0.1	0.05	0.025
Δt	0.04	0.02	0.01	0.005

We plot the rooted mean square error versus the level, in what we will refer to as the controlled convergence graph. We use this terminology as it indicates that we are controlling the number of samples in order to have a stochastic and discretization errors of the same order. The plot is shown in Fig. 6.

In the same plot, we show the controlled convergence for the MLMC. For each level L , we perform an MLMC run involving as many levels by setting the number of samples on the coarse level to 2^{4L} . The number of samples for the following levels is then calculated by dividing this number with the ideal sampling ratio, which corresponds in our case to 16. This number can be deduced by enforcing that $h_l^2 N_l^{-1/2} = h_{l+1}^2 N_{l+1}^{-1/2}$ which represents the error contribution from levels l and $l + 1$.

4.2. Numerical experiments in 3D+1

We intend in the next subsections to estimate and verify the convergence behavior for the MC and the MLMC estimators, as done for the 3D+1 case above. Furthermore, a comparison in terms of work to accuracy for both methods will be conducted.

We rely on the cube and the ventricle geometries, introduced in Section 3 (cf. Fig. 1). As opposed to a classical MLMC approach, where one would refine an initially coarse mesh in order to reduce the variance, we here consider the realistic setting of an application for which only one mesh of the given geometry is provided. This mesh is provided at the fine discretization level, therefore the goal for the MLMC is to reduce the work load by considering coarser meshes. We rely on a nested mesh hierarchy of 3 and 4 levels respectively for the ventricle and the cube geometries. The number of space-time degrees of freedom (Dof's), the space and time discretization steps of all the considered levels for the cube and the ventricle are respectively reported in Tables 2 and 3.

Table 3
Details about the mesh hierarchy for the ventricle geometry.

l	0	1	2
Dof's	154,546	2,120,420	31,184,747
h	0.1	0.05	0.025
Δt	0.02	0.01	0.005

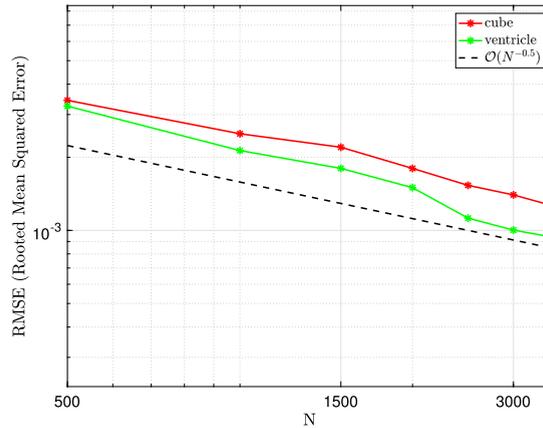


Fig. 7. Convergence rate for the MC estimator for the cube and the ventricle.

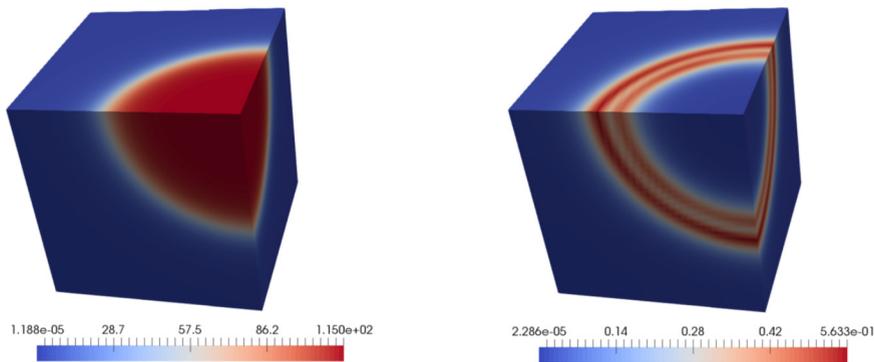


Fig. 8. Mean (left) and variance (right) at the final time state for the cube.

4.2.1. Monte-Carlo in 3D+1

We conduct a Monte-Carlo study of $N = 10000$ samples on the fine discretization levels in order to obtain a reference solution. In Figs. 8 and 9, we show the mean and the variance at the final time state for the cube and the ventricle, respectively.

In Section 3, we present the theoretical semi-linear convergence of the MC method (cf. Equation (16)). We verify this in Fig. 7 in which we plot the rooted mean square error (RMSE) against the number of samples (N) used for the estimation process. It can be observed how the RMSE indeed behaves as $RMSE = \mathcal{O}(N^{-0.5})$. Notice that these plots rely on averaged multiple runs.

4.2.2. Multilevel Monte-Carlo in 3D+1

The setting for the MLMC is very similar to that of the 1D+1 case. We underline the importance of the ideal sampling ratio as detailed at the end of Section 3. We provide in Figs. 10 and 11 a visualization of the MLMC estimator (mean at coarse level and corrections between successive levels) for the cube and the ventricle respectively.

Fig. 12 presents the controlled convergence graph of MC and MLMC for both the cube and the ventricle. It demonstrates that the general convergence rate obtained for these two methods is as expected. MC and MLMC have therefore a similar controlled convergence rate.

Work behavior for a fixed MLMC setting: As mentioned in the previous paragraph, the plots for controlled error in Fig. 12 (or controlled convergence) implicitly indicate that MLMC achieves the same precision with less effort. In order to better

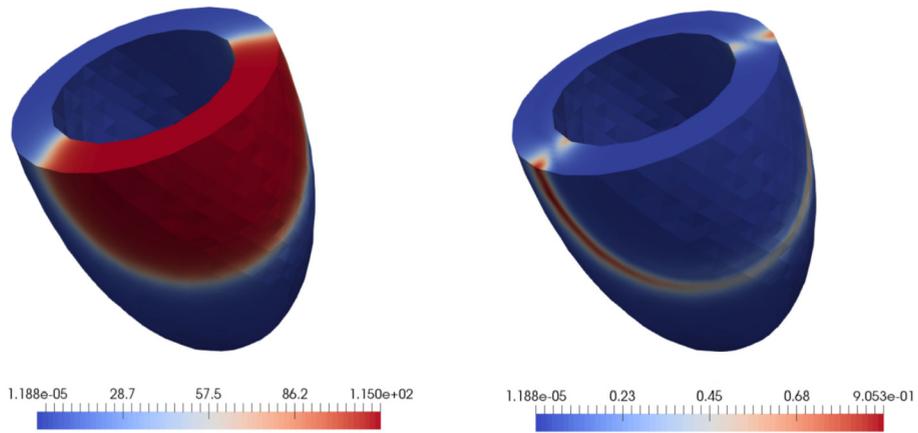
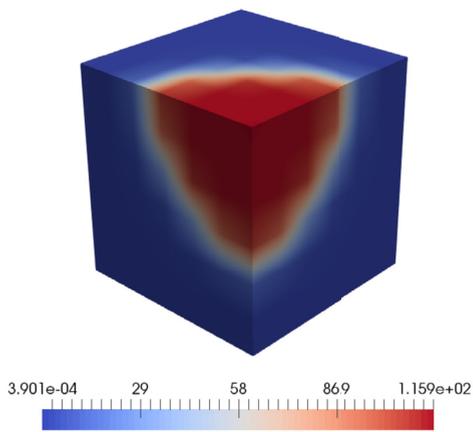
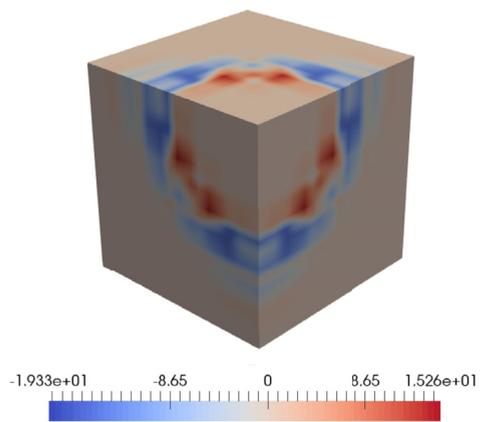


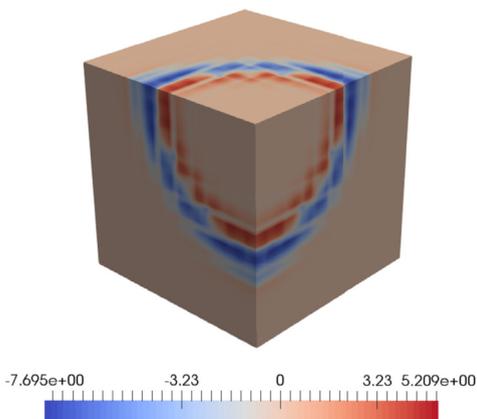
Fig. 9. Mean (left) and variance (right) at the final time state for the ventricle.



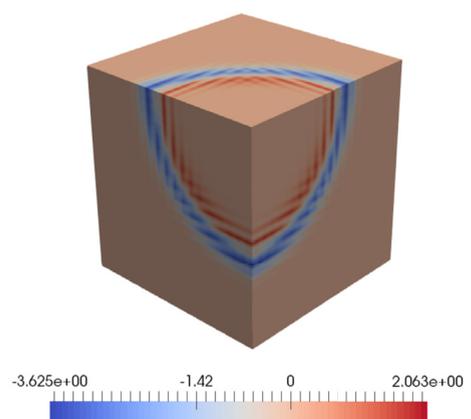
(a) Mean at coarse level.



(b) Correction between levels 0 and 1.



(c) Correction between levels 1 and 2.



(d) Correction between levels 2 and 3.

Fig. 10. Mean at coarse level and corrections between successive levels for the cube.

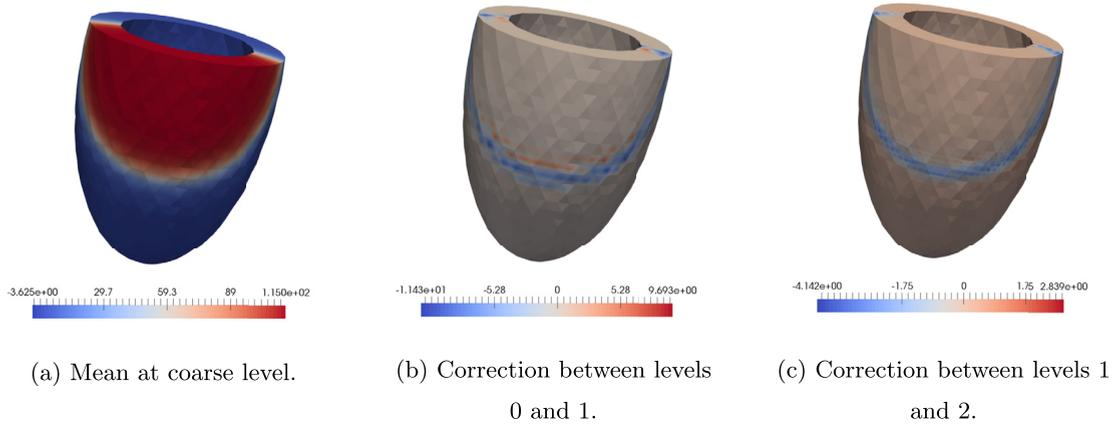


Fig. 11. Mean at coarse level and corrections between successive levels for the ventricle.

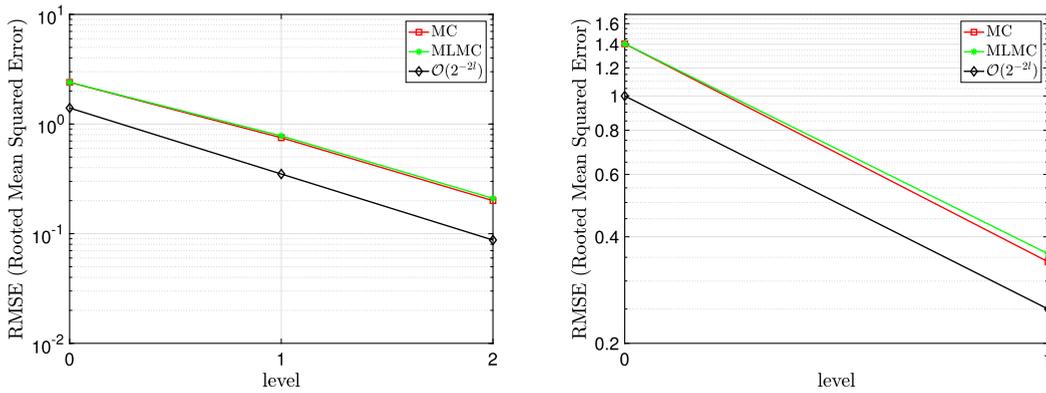


Fig. 12. Controlled convergence (MC and MLMC) graphs for the cube (left) and the ventricle (right).

underline this aspect, we attempt here to evaluate this difference of produced work for both methods, whenever the MLMC setting (fine level, number of levels, mesh hierarchy,...) had been previously set. Let us denote with W'_l the expected work for solving a sample at level $l, \forall l = 0, \dots, L$. We also introduce W_l to be defined as

$$W_l = W'_{l-1} + W'_l, \quad \forall l = 0, \dots, L$$

with $W'_{-1} \equiv 0$. We can therefore formulate the total work expression for both methods in the following way

$$W_{MC,L} = W'_L N_L, \quad W_{MLMC,L} = \sum_{l=0}^L W_l N_l, \tag{28}$$

where N_l is the number of samples computed on each level $l, \forall l = 0, \dots, L$. We furthermore can use the sampling ratio between levels for MLMC, which we will denote here with β , to rewrite the MLMC total work in function of the number of samples at the fine level exclusively

$$W_{MLMC,L} = \left(\sum_{l=0}^L \beta^{L-l} W_l \right) N_L.$$

The latter expression allows us to select different values of the number of samples N_l at the fine level, therefore evaluating the work for different samplings. In Fig. 13, we plot the RMSE versus the total amount of work for both methods. We here define the expected work W'_l for computing a sample at level l to be the total running time for solving a sample at the same level, averaged through solving a hundred of samples. Furthermore, we normalize with respect to the time for solving a sample on the fine level.

Asymptotical work behavior of MC/MLMC: Equivalently to the controlled convergence concept, we would like now to study the work in the context of the controlled error. Following the theory, in order to achieve an error of order 2^{-2l} , the cost for

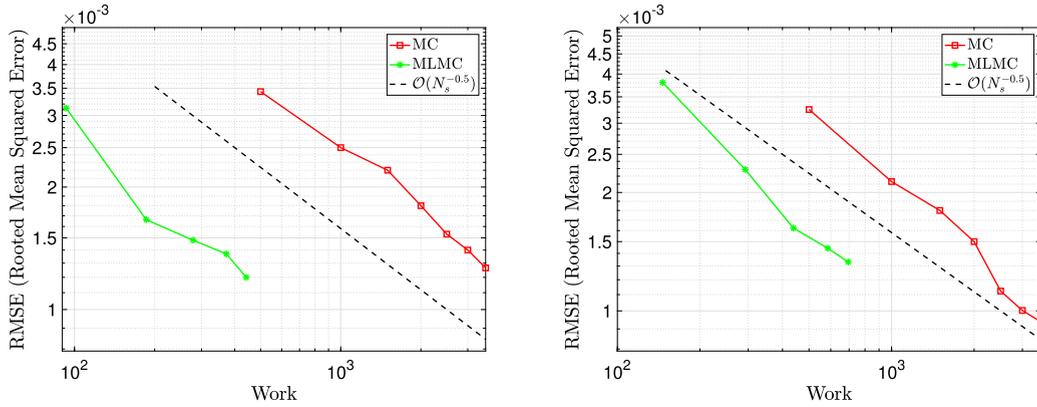


Fig. 13. Work comparison between MC and MLMC for the cube (left) and the ventricle (right). Works ($W_{MC,L}$ and $W_{MLMC,L}$ as defined in (28)) are measured on the basis of CPU time.

solving a sample with linear FE is given by

$$C_{FE,l} = 2^{\gamma dl},$$

where γ is the complexity of the solution method used and d is the dimension for the PDE problem considered (equal to 4 in our case). The number of samples to execute on a level l for getting an error of order 2^{-2l} is given by

$$N_{MC,l} = 2^{4l}.$$

From these, one can recover the total amount of work required by both methods for a fine discretization level L

$$W_{MC,L} = C_{FE,L} N_{MC,L} = 2^{\gamma dL} 2^{4L} = 2^{(\gamma d+4)L} \tag{29}$$

and

$$W_{MLMC,L} = \sum_{l=0}^L C_{FE,l} N_{MC,l} = \sum_{l=0}^L 2^{\gamma dl} 2^{4(L-l)} = 2^{4L} \sum_{l=0}^L 2^{(\gamma d-4)l} \tag{30}$$

This can be rewritten as

$$W_{MLMC,L} = \begin{cases} L2^{4L} & \text{if } \gamma d = 4, \\ \frac{2^{\gamma dL} - 2^{4L}}{2^{\gamma d-4} - 1} & \text{if } \gamma d \neq 4. \end{cases} \tag{31}$$

Therefore, in case dimension d and complexity γ are known, equations (29) and (31) give a theoretical estimate on the asymptotical behavior for the work produced by both MC and MLMC. In our case, since we are measuring the work in terms of time to solution (CPU time), we can estimate the product of those values as

$$\gamma d = \log_2(W'_{l+1}/W'_l). \tag{32}$$

The following estimates yield

- $\gamma d \approx 4$ for the cube
- $\gamma d \approx 2$ for the ventricle

The MC and MLMC produced work for the cube experiments are therefore given by

$$W_{MC,L} = 2^{8L}, \quad W_{MLMC,L} = L2^{4L}. \tag{33}$$

For the ventricle experiments, we obtain

$$W_{MC,L} = 2^{6L}, \quad W_{MLMC,L} = \frac{3(2^{4L} - 2^{2L})}{4}. \tag{34}$$

This demonstrates very clearly that the asymptotic behavior of MLMC in terms of computational work is significantly reduced in comparison with standard MC. We illustrate this in Fig. 14. We recall that the controlled RMSE e_l at level l is of

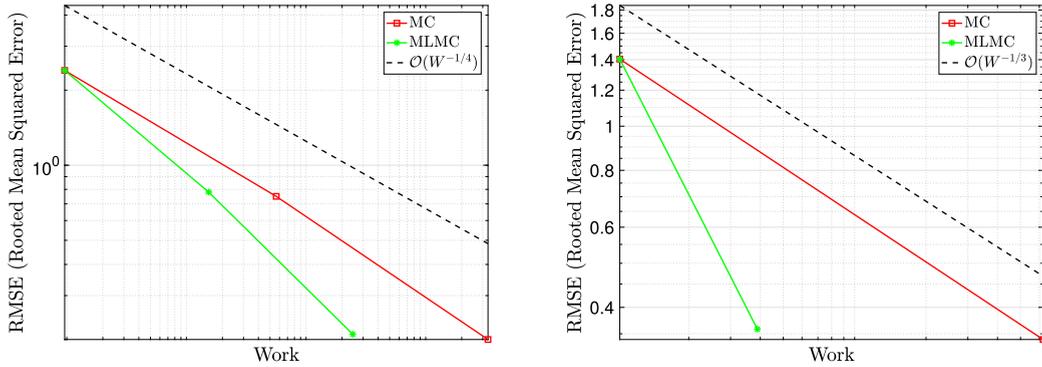


Fig. 14. MC and MLMC asymptotical work for the cube (left) and the ventricle (right). The “dimensional complexity” (γd as defined in (32)) is measured on the basis of CPU time.

order 2^{-2l} . From this, we can deduce that $e_l^{-0.5} \sim 2^l$. Putting this together with equations (33) and (34), shows that for the cube we have

$$e_l = \mathcal{O}(W_{MC}^{-1/4}),$$

whereas for the ventricle

$$e_l = \mathcal{O}(W_{MC}^{-1/3}).$$

This is illustrated in Fig. 14.

5. Conclusions

We presented a novel approach aimed at performing uncertainty quantification for time-dependent problems, in the presence of high-dimensional input uncertainties and output quantities-of-interest. We assessed the performance of the proposed framework by showing in details its application to cardiac electrophysiology. The results show that the proposed method attains the ideal convergence rates predicted by the theory. Our methodology relies on a close integration of multi-level Monte Carlo methods, parallel iterative solvers, and a space-time discretization. This allows to exploit their synergies, such as for the initialization of Newton’s method using the solution of past samples when they become available. The proposed method can also be extended to space-time adaptivity and time-varying domains, both of which are very relevant in the context of cardiac electrophysiology and will be addressed in future works.

CRedit authorship contribution statement

S. Ben Bader: Conceptualization, Investigation, Methodology, Software, Visualization, Writing – original draft, Writing – review & editing. **P. Benedusi:** Software, Writing – original draft, Writing – review & editing. **A. Quaglino:** Conceptualization, Methodology, Software, Writing – original draft, Writing – review & editing. **P. Zulian:** Software, Writing – original draft, Writing – review & editing. **R. Krause:** Conceptualization, Funding acquisition, Methodology, Resources, Supervision.

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

The authors would like to thank the Swiss National Science Foundation (SNSF) for their support through the project “Multilevel Methods and Uncertainty Quantification in Cardiac Electrophysiology” in collaboration with the University of Basel (grant agreement SNSF-205321_169599), but also through the project “HEARTFUSION: Imaging-driven Patient-specific Cardiac Simulation” in collaboration with the University of Bern (grant agreement SNSF-169239). Furthermore, the authors would like to thank the Deutsche Forschungsgemeinschaft (DFG) which has supported parts of the software used here in the SPPEXA program “EXASOLVERS – Extreme Scale Solvers for Coupled Problems” and SNSF under the lead agency grant agreement SNSF-162199. Part of the software tools utilized in this paper were developed as part of the activities of the Swiss Centre for Competence in Energy Research on the Future Swiss Electrical Infrastructure (SCCER-FURIES), which is financially supported by the Swiss Innovation Agency (Innosuisse – SCCER program). The authors gratefully acknowledge financial support by the Theo Rossi di Montelera Foundation, the Mantegazza Foundation, and FIDINAM to the Center of Computational Medicine in Cardiology.

References

- [1] Kadir A. Aziz, Peter Monk, Continuous finite elements in space and time for the heat equation, *Math. Comput.* 52 (186) (1989) 255–274.
- [2] Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman, Lisandro Dalcin, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Karl Rupp, Barry F. Smith, Stefano Zampini, Hong Zhang, Hong Zhang, PETSc users manual, Technical Report ANL-95/11 - Revision 3.7, Argonne National Laboratory, 2016.
- [3] Andrea Barth, Christoph Schwab, Nathaniel Zollinger, Multi-level Monte Carlo finite element method for elliptic PDEs with stochastic coefficients, *Numer. Math.* 119 (1) (2011) 123–161.
- [4] Pietro Benedusi, Carlo Garoni, Rolf Krause, Xiaozhou Li, Stefano Serra-Capizzano, Space-time FE-DG discretization of the anisotropic diffusion equation in any dimension: the spectral symbol, *SIAM J. Matrix Anal. Appl.* 39 (3) (2018) 1383–1420.
- [5] Julia Charrier, Robert Scheichl, Aretha L. Teckentrup, Finite element error analysis of elliptic PDEs with random coefficients and its application to multilevel Monte Carlo methods, *SIAM J. Numer. Anal.* 51 (1) (2013) 322–352.
- [6] Piero Colli Franzone, Luciano Guerri, Spreading of excitation in 3-D models of the anisotropic cardiac tissue. I. Validation of the eikonal model, *Math. Biosci.* 113 (1993) 145–209, [https://doi.org/10.1016/0025-5564\(93\)90001-q](https://doi.org/10.1016/0025-5564(93)90001-q).
- [7] Cesare Corrado, Jean-Frédéric Gerbeau, Philippe Moireau, Identification of weakly coupled multiphysics problems: application to the inverse problem of electrocardiography, *J. Comput. Phys.* 283 (2015) 271–298.
- [8] Daniel Drzisga, Björn Gmeiner, U. Rüde, Robert Scheichl, Barbara I. Wohlmuth, Scheduling massively parallel multigrid for multilevel Monte Carlo methods, *SIAM J. Sci. Comput.* 39 (5) (2017) S873–S897.
- [9] George Fishman, Monte Carlo: Concepts, Algorithms, and Applications, corrected edition, Springer, New York, ISBN 978-0-387-94527-9, February 2003.
- [10] Richard FitzHugh, Impulses and physiological states in theoretical models of nerve membrane, *Biophys. J.* 1 (6) (1961) 445–466.
- [11] Donald A. French, J. Thomas King, Analysis of a robust finite element approximation for a parabolic equation with rough boundary data, *Math. Comput.* 60 (201) (1993) 79–104.
- [12] Martin J. Gander, 50 years of time parallel time integration, in: *Multiple Shooting and Time Domain Decomposition Methods*, Springer, 2015, pp. 69–113.
- [13] Michael B. Giles, Multilevel Monte Carlo methods, *Acta Numer.* 24 (2015) 259–328, <https://doi.org/10.1017/S096249291500001X>.
- [14] Wei Gong, Michael Hinze, Error estimates for parabolic optimal control problems with control and state constraints, *Comput. Optim. Appl.* 56 (1) (2013) 131–151.
- [15] Helmut Harbrecht, Michael Peters, Reinhold Schneider, On the low-rank approximation by the pivoted Cholesky decomposition, *Appl. Numer. Math.* 62 (4) (2012) 428–440, <https://doi.org/10.1016/j.apnum.2011.10.001>.
- [16] Helmut Harbrecht, Michael Peters, Markus Siebenmorgen, On multilevel quadrature for elliptic stochastic partial differential equations, in: *Sparse Grids and Applications*, Springer, 2012, pp. 161–179.
- [17] Helmut Harbrecht, Michael Peters, Markus Siebenmorgen, Efficient approximation of random fields for numerical applications, *Numer. Linear Algebra Appl.* 22 (4) (2015) 596–617.
- [18] Christian Hesch, Peter Betsch, A comparison of computational methods for large deformation contact problems of flexible bodies, *Z. Angew. Math. Mech.* 86 (10) (2006) 818–827, <https://onlinelibrary.wiley.com/doi/abs/10.1002/zamm.200610289>.
- [19] Jan S. Hesthaven, Gianluigi Rozza, Benjamin Stamm, et al., *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, Springer, 2016.
- [20] Alan L. Hodgkin, Andrew F. Huxley, A quantitative description of membrane current and its application to conduction and excitation in nerve, *J. Physiol.* 117 (4) (1952) 500–544.
- [21] Daniel E. Hurtado, Duvan Henao, Gradient flows and variational principles for cardiac electrophysiology: toward efficient and robust numerical simulations of the electrical activity of the heart, *Comput. Methods Appl. Mech. Eng.* 273 (2014) 238–254.
- [22] Phaedon-Stelios Koutsourelakis, Accurate uncertainty quantification using inaccurate computational models, *SIAM J. Sci. Comput.* 31 (5) (2009) 3274–3300.
- [23] Rolf Krause, Patrick Zulian, A parallel approach to the variational transfer of discrete fields between arbitrarily distributed unstructured finite element meshes, *SIAM J. Sci. Comput.* 38 (3) (2016) C307–C333.
- [24] Eleanor McDonald, Andy Wathen, A simple proposal for parallel computation over time of an evolutionary process with implicit time stepping, in: *Numerical Mathematics and Advanced Applications, ENUMATH 2015*, Springer, 2016, pp. 285–293.
- [25] Walter T. Miller, David B. Geselowitz, Simulation studies of the electrocardiogram. I: the normal heart, *Circ. Res.* 43 (2) (1978) 301–315.
- [26] Steven Niederer, Lawrence Mitchell, Nicolas Smith, Gernot Plank, Simulating human cardiac electrophysiology on clinical time-scales, *Front. Physiol.* 2 (2011).
- [27] Jürg Nievergelt, Parallel methods for integrating ordinary differential equations, *Commun. ACM* 7 (12) (1964) 731–733.
- [28] Benjamin Peherstorfer, Karen Willcox, Max Gunzburger, Optimal model management for multifidelity Monte Carlo estimation, *SIAM J. Sci. Comput.* 38 (5) (Jan 2016) A3163–A3194, <https://doi.org/10.1137/15m1046472>.
- [29] Benjamin Peherstorfer, Karen Willcox, Max Gunzburger, Survey of multifidelity methods in uncertainty propagation, inference, and optimization, *SIAM Rev.* 60 (3) (2018) 550–591.
- [30] Alexander Popp, Barbara I. Wohlmuth, Michael W. Gee, Wolfgang A. Wall, Dual quadratic mortar finite element methods for 3D finite deformation contact, *SIAM J. Sci. Comput.* 34 (4) (2012) B421–B446.
- [31] Andrew J. Pullan, Karl A. Tomlinson, Peter J. Hunter, A finite element method for an eikonal equation model of myocardial excitation wavefront propagation, *SIAM J. Appl. Math.* 63 (1) (2002) 324–350, <https://doi.org/10.1137/S0036139901389513>.
- [32] Alessio Quaglino, Seif Ben Bader, Rolf Krause, SLOTH: an open-source Python library for UQ, available at <https://bitbucket.org/quagla/sloth/src/master/>, 2017, software.
- [33] Alessio Quaglino, Simone Pezzuto, Phaedon-Stelios Koutsourelakis, Angelo Auricchio, Rolf Krause, Fast uncertainty quantification of activation sequences in patient-specific cardiac electrophysiology meeting clinical time constraints, *Int. J. Numer. Methods Biomed. Eng.* 34 (7) (2018) e2985.
- [34] Alessio Quaglino, Simone Pezzuto, Rolf Krause, High-dimensional and higher-order multifidelity Monte Carlo estimators, *J. Comput. Phys.* 388 (2019) 300–315.
- [35] Alfio Quarteroni, Riccardo Sacco, Fausto Saleri, *Numerical Mathematics*, vol. 37, Springer Science & Business Media, 2010.
- [36] Aretha L. Teckentrup, Robert Scheichl, Michael B. Giles, Elisabeth Ullmann, Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients, *Numer. Math.* 125 (3) (2013) 569–600.
- [37] Barbara I. Wohlmuth, A mortar finite element method using dual spaces for the Lagrange multiplier, *SIAM J. Numer. Anal.* 38 (3) (2000) 989–1012.
- [38] Patrick Zulian, Alena Kopaničáková, Maria Chiara Giuseppina Nestola, Andreas Fink, Nur Fadel, Victor Magri, Teseo Schneider, Eric Botter, Jan Mankau, Rolf Krause Utopia, A C++ embedded domain specific language for scientific computing, software, <https://bitbucket.org/zulian/utopia>, 2016.