

This is a postprint version of the following published document:

Bernal, F. (2018). An Implementation of Milstein's Method for General Bounded Diffusions. *Journal of Scientific Computing*, 79(2), 867–890.

DOI: <https://doi.org/10.1007/s10915-018-0884-6>

Copyright © 2018, Springer Science Business Media, LLC, part of Springer Nature

An implementation of Milstein's method for general bounded diffusions

Francisco Bernal*

Abstract

Despite its generality and powerful convergence properties, Milstein's method for functionals of spatially bounded stochastic differential equations is widely regarded as difficult to implement. This has likely prevented it from being utilised in applications. In this paper, we design and analyse in detail one such implementation. The presented method turns out to be on par with other, popular schemes in terms of computational cost—but with a (nearly) linear weak convergence rate under the usual smoothness requirements on coefficients and boundary. Two byproducts of theoretical interest are a new, non-standard rank-one update formula, and a connection between numerics of bounded diffusions and Eikonal equations. Three examples are worked out, confirming the accuracy and robustness of the method.

(Accepted for publication in the Journal of Scientific Computing, 2018)

Keywords: stopped diffusion, reflected diffusion, Feynman-Kac formula, stochastic numerics, rank-one updates to matrix decompositions, Eikonal equation.

AMS subject classifications: 60H35, 65C30, 65C05.

1 Introduction

Scope and motivation. We are interested in *efficient* weak schemes for Feynman-Kac functionals of stochastic differential equations (SDEs) related to *general* linear boundary value problems (BVPs). By *efficiency*, we refer to the computational cost required to bound the expected error within a given tolerance and a given confidence interval. (This depends on the weak rate of convergence w.r.t. the timestep h of the SDE scheme, and on its cost per iteration.) By *general*, we mean that no restrictions are posed on the coefficients and boundary beyond those required for solvability of the associated SDE (i.e. a minimum degree of regularity of both). In particular, we cover BVPs with mixed boundary conditions (BCs) (i.e. of Dirichlet type on some portion of the boundary and Robin

*CMAP, Ecole Polytechnique, email: francisco.bernal@polytechnique.edu

elsewhere), for which the stochastic representation is in terms of functionals of SDEs spatially enclosed by reflecting and stopping boundaries.

Such stochastic schemes are the basis for the pointwise solution, $u(x_0, t_0)$, of linear BVPs in \mathbb{R}^D via the Monte Carlo method [27]. Besides the mathematical interest of this connection by itself, the Monte Carlo approach is computationally advantageous in many applications, v.g. [2, 4, 7, 22, 24]. The stochastic numerics presented here can also be adapted to construct a variety of probability densities of a Wiener diffusion inside a domain enclosed by absorbing and/or reflecting surfaces, including the survival probability, the time spent close to the reflecting boundary, and first-passage times from the domain [20, 29, 30].

State of the art. Given a general linear BVP, it is challenging to numerically integrate its associated Feynman-Kac functionals with a satisfactory weak order of convergence w.r.t. h (δ henceforth). This is specially true when the BVP has mixed BCs—or equivalently, when the SDE is stopped on some portion of the boundary and reflected elsewhere. The method of Constantini *et al.* [8] may be considered as the standard; it is very easy to code but has a (proven) δ of just up to one half. The difficulties pertain to the determination (in a weak sense) of the first-exit point (in the presence of stopping boundaries—see [3] and references therein), and of the local time (in the presence of reflecting ones). Other schemes have a better δ , but they are less general; the following list is not exhaustive. Gobet’s half-space approximation [15] is very useful in practice. He has theoretically analysed the schemes for either killed (rather than stopped) or reflected diffusions; and proved that many Feynman-Kac functionals converge linearly with h —including those for homogeneous Dirichlet BCs, and those for Neumann (rather than Robin) BCs. Bossy *et al.* [5] put forward a straightforward method which has $\delta = 1$ for the Feynman-Kac functionals associated with homogeneous Neumann BCs. A relatively new approach (see [1] and references therein) relies on adaptive h , based on a running error estimate. It has shown very good results (including in nonconvex domains) when the diffusion is a Brownian motion, but we are not aware of extensions to more general cases. Also sticking to the Brownian motion, the method of randomisation can deal accurately with all sorts of linear BCs (see [22] and references therein).

Besides the method of Constantini *et al.*, there is another one, due to Milstein, which also is completely general *and* has a proven $O(h)$ weak convergence rate [25, 26]: we will call it MM in short. Paradoxically (since it predates all the methods quoted above), MM has gone underreported in the literature, most likely because of a reputation of being complex to implement. (In fact, we are not aware of numerical experiments with it having been published, beyond one example in [27, p. 372] and a few more in our review [3], both on purely stopped diffusions.) For instance, MM is described as follows in [8, p. 77]: “(...) two other weak discretization schemes (...) are considered in domains with smooth boundary (...) One of them achieves the rate of convergence h but, as pointed out by the author, is difficult to implement”. Moreover, quoting from [15, p. 280]: “An appropriate Markov chain approximation at random discretization

times has been studied by Milstein (...) His procedure requires at each step near the boundary to change coordinates, and, by the way, the algorithm seems to be difficult to implement”.

Our contribution. The reason why MM is hard to implement is that the algorithms are incomplete. At a given point, the user is instructed to “take the ellipsoid tangent to the boundary”, or to perform a given rotation or a certain matrix decomposition close to reflecting boundaries; but details as to how to are glossed over both in the seminal papers as in the later book [27, chapter 6]. In this paper, we provide a practical and theoretically sound such implementation of MM—apparently the first one. (Henceforth, we will call it “our implementation” or “Algorithm 1”, to distinguish it from the original algorithm.¹)

Regarding the determination of the tangent ellipsoid (Section 4.4), we have established a novel link with an anisotropic Eikonal equation (Lemma 2) which allows one, theoretically, to do it exactly. Nonetheless, this approach is unpractical except with a constant diffusion matrix (as in Example 5.3). Therefore, a fast $O(D^2)$ half-space approximation (Lemma 3) is also provided. For the reflections (Section 4.5), we have introduced a non-standard, non-Cholesky update formula which endows the overall implementation with $O(D^2)$ complexity (Algorithm 2). This is relevant because MM is designed for Monte Carlo simulations, which are often meant for high-dimensional problems ($D \gg 1$). In addition to those three main contributions, we discuss every other aspect and provide fast recipes for them. In sum, this paper lifts what were the main obstacles to the routinary utilization of Milstein’s method for general bounded diffusions.

A preliminary, simplified version of the implementation presented here, valid only for smooth, purely stopping boundaries, was sketched in the review paper [3]. That algorithm, however, is now superseded by Algorithm 1.

Our implementation turns out to have a cost per time step comparable (less than 2 times larger) to that of the integrator by Constantini *et al.*—however, with a (nearly) linear weak convergence rate. To put this fact in perspective, the complexity of a Monte Carlo simulation (proportional to the CPU time) with error tolerance $0 < \epsilon \ll 1$ drops dramatically (by a factor $1/\epsilon$) from $\delta = 1/2$ to $\delta = 1$ [12]. Substantial further gains would be possible by combining Algorithm 1 with Multilevel [12], extrapolation [23], or both [18]. Critically, all of the previous strategies rely on the *a priori* knowledge of δ —thus highlighting the suitability of MM thanks to its sound theoretical foundation.

While MM has an asymptotically proven $\delta = 1$ under the assumptions of adequate smoothness of the coefficients and of the boundary, our implementation might not exactly reproduce it in two cases:

1. Unless Ω has a trivial shape (or is a combination thereof), the distance map will itself be a numerical approximation—thus, possibly spoiling

¹The Matlab code and data files used for this paper are available at the journal repository, and upon request from the author.

$\delta = 1$. (Note, however, that the same holds for *any* numerical scheme for bounded diffusions, not just the one presented in this paper.)

2. When the absorbing boundaries are curved and the fast recipe in Lemma 3 for constructing the smallest tangent ellipsoid is employed (as will be typically the case). The reason is that Lemma 3 relies on locally approximating the boundary by its tangent plane, but we have not rigorously proved convergence of $\delta \rightarrow 1$ as $h \rightarrow 0$.

Consequently, we claim that our algorithm has a “(nearly) linear weak convergence rate”. On the other hand, we emphasize that all of the numerical experiments involving smooth boundaries which we have carried out univocally suggest $\delta = 1$, in practice.

We close the Introduction by briefly commenting on preprocessing.

The main effort which must be independently undertaken before the Monte Carlo simulation is the generation of a signed distance map for Ω . Again, we stress that this is required for any numerical scheme for bounded SDEs—although it often goes unmentioned. In the simplest cases, like a ball or a parallelepiped, an exact distance formula is available. Otherwise, we propose solving an Eikonal equation via the Fast Marching Method (check Section 4.3).

Finally, if the exact tangent ellipsoid is needed, the anisotropic Eikonal equation in Lemma 3 must be numerically solved in advance. Example 5.3 has been crafted to illustrate a situation where this would be highly advantageous, for that problem is very hard to solve with other numerical schemes.

Since preprocessing is either not particular to the algorithm presented here, or up to some point optional—and in both cases to be tackled with an independent method—its cost has not been explicitly included in the main discussion.

Organisation of the paper. Section 2 recalls the theoretical connection between linear second-order BVPs with mixed BCs, and bounded stopped/reflected SDEs. (The form of the Feynman-Kac formulas there is not the most usual one, but one tailored to MM.) In Section 3, the original MM is described, and the new Algorithm 1 is listed. Implementation details are discussed around several new lemmas in Section 4, which is the core of the paper. Three numerical examples are worked out in Section 5; and conclusions are drawn in Section 6. To avoid clutter, all proofs of the lemmas in Section 4 have been moved into Appendix A. Finally, Appendix B lists a few relevant Matlab code snippets.

2 Feynman-Kac formulas in Milstein’s form

Let $D \geq 2$, $\Omega \subset \mathbb{R}^D$ be a bounded domain, and $\bar{\Omega} = \bar{\Omega} \cup \partial\Omega$, where the open connected set $\bar{\Omega}$ is the interior of the domain and $\partial\Omega$ its boundary. Consider the linear parabolic BVP of second order with mixed BCs:

$$\begin{cases} \frac{\partial u}{\partial t} = \mathcal{L}(\mathbf{x}, t)u + c(\mathbf{x}, t)u + f(\mathbf{x}, t) & \text{if } 0 < t \leq T, \mathbf{x} \in \bar{\Omega}, \\ u = p(\mathbf{x}) & \text{if } t = 0, \mathbf{x} \in \Omega, \\ u = g(\mathbf{x}, t) & \text{if } 0 < t \leq T, \mathbf{x} \in \partial\Omega_A, \\ \frac{\partial u}{\partial \mathbf{N}} = \varphi(\mathbf{x}, t)u + \psi(\mathbf{x}, t) & \text{if } 0 < t \leq T, \mathbf{x} \in \partial\Omega_R, \end{cases} \quad (1)$$

where $T > 0$, $\varphi(\mathbf{x}, t) \leq 0$, and the differential generator is given by

$$\mathcal{L}(\mathbf{x}, t)u = \frac{1}{2} \sum_{i,j=1}^D a_{ij}(\mathbf{x}, t) \frac{\partial^2 u}{\partial x_i \partial x_j} + \sum_{k=1}^D b_k(\mathbf{x}, t) \frac{\partial u}{\partial x_k}. \quad (2)$$

The matrix $A(\mathbf{x}, t) := [a_{ij}]_{i,j=1}^D$ is positive definite, and $\mathbf{b}(\mathbf{x}, t) := (b_1, \dots, b_D)^T$ is called the drift. All of the coefficient functions in (1), namely $a_{ij}, b_i, c, f, p, g, \varphi$ and ψ are assumed continuous, and complying with the compatibility conditions at time $t = 0$ (see [17] or [8, equations (2.15)-(2.17)]). The boundary is decomposed as $\partial\Omega = \partial\Omega_A \cup \partial\Omega_R \cup \partial\Omega_S$, such that $\partial\Omega_A \cap \partial\Omega_R = \partial\Omega_A \cap \partial\Omega_S = \partial\Omega_R \cap \partial\Omega_S = \emptyset$. The outward² unit normal vector \mathbf{N} is assumed to be well defined on the boundary save perhaps on a set $\partial\Omega_S$; $\partial\Omega_A$ stands for the portion of the boundary (if any) where Dirichlet BCs are imposed; and on $\partial\Omega_R$, BCs involving the normal derivative, (i.e. $\mathbf{N}^T \nabla u$) hold, where $\nabla u = (\partial u / \partial x_1, \dots, \partial u / \partial x_D)^T$. (BCs involving oblique derivatives will not be considered in this paper.) Such BCs are of Neumann type iff $\varphi = 0$, or of Robin type otherwise.

Sufficient conditions for existence of a unique classical solution to (1). When $\partial\Omega_A = \emptyset$ (respectively $\partial\Omega_R = \emptyset$) we say the BCs are purely reflecting (resp. purely stopping), while when both $\partial\Omega_A \neq \emptyset \neq \partial\Omega_R$, we say the BCs are mixed. If the BCs are purely reflecting (resp. purely stopping), and $\partial\Omega_S = \emptyset$, theorem 2.6 (resp. theorem 2.7) in [8] (see also [17]) ensure the existence of a unique classical solution, with regularity dependent on that of the BVP coefficients and of $\partial\Omega$. (By a classical solution, we mean that $u(\mathbf{x}, t)$ lives in the Hölder space $C^{1,2}([0, T] \times \bar{\Omega})$.) In the mixed BCs case, this connection is less general and more dependent on the smoothness of the boundary [21, 28].

Stochastic representation of the pointwise solution to (1). Under slightly stronger conditions, the stochastic representation of the PDE with mixed BCs (1) expresses its pointwise solution $u(\mathbf{x}_0, t)$ as the expected value of a functional of an SDE starting at \mathbf{x}_0 at time zero and being normally reflected on $\partial\Omega_R$ and stopped on $\partial\Omega_A$.

Let $\sigma(\mathbf{x}, t)$ (called the diffusion matrix) be defined by $\sigma(\mathbf{x}, t)\sigma^T(\mathbf{x}, t) = A(\mathbf{x}, t)$. (This is always possible since A is positive definite. Hence, $\det \sigma \neq 0$.) The following result is an extension of the well-known Feynman-Kac formulas for purely reflected ($\partial\Omega_A = \emptyset$) and purely stopped ($\partial\Omega_R = \emptyset$) diffusions, adapted from [8, theorem 2.5] and [27, chapter 6].

²In the SDE literature, the normal is usually taken inwards. Here we follow the PDE convention.

Theorem 1 Assume that: i) a classical unique solution to (1) does exist; ii) there exists a constant $L > 0$ such that for $\mathbf{x}, \mathbf{y} \in \bar{\Omega}$ and $t \in [0, T]$

$$\|\sigma(\mathbf{x}, t) - \sigma(\mathbf{y}, t)\| \leq L\|\mathbf{x} - \mathbf{y}\|, \quad (3)$$

$$\|\mathbf{b}(\mathbf{x}, t) - \mathbf{b}(\mathbf{y}, t)\| \leq L\|\mathbf{x} - \mathbf{y}\|; \quad (4)$$

iii) $\partial\Omega$ is piecewise C^1 (i.e. C^1 save on maybe a set $\partial\Omega_S$); and iv) either $\bar{\Omega}$ is convex, or $u(\mathbf{x}, t)$ can be extended to a function $C^{1,2}([0, T] \times \mathbb{R}^d)$. Then, for $0 \leq t \leq T$ the pointwise solution of (1) admits the following stochastic representation:

$$u(\mathbf{x}_0, t) = \mathbb{E}[\phi_\tau] := \mathbb{E}\left[q(\mathbf{X}_\tau)Y_\tau + Z_\tau\right], \quad (5)$$

where

$$q(\mathbf{X}_\tau) = \begin{cases} g(\mathbf{X}_\tau, T - \tau), & \text{if } \tau < T, \\ p(\mathbf{X}_T), & \text{if } \tau \geq T, \end{cases} \quad (6)$$

and the processes $(\mathbf{X}_t, Y_t, Z_t, \xi_t)$ are governed by the following set of SDEs driven by a standard D -dimensional Wiener process \mathbf{W}_t :

$$\begin{cases} d\mathbf{X}_t = (\mathbf{b}(\mathbf{X}_t, T - t) - \sigma(\mathbf{X}_t, T - t)\boldsymbol{\mu})dt + \sigma(\mathbf{X}_t, T - t)d\mathbf{W}_t - \mathbf{N}(\mathbf{X}_t)d\xi_t & \mathbf{X}_0 = \mathbf{x}_0, \\ dY_t = c(\mathbf{X}_t, T - t)Y_tdt + \varphi(\mathbf{X}_t, T - t)Y_td\xi_t + Y_t\boldsymbol{\mu}^T d\mathbf{W}_t & Y_0 = 1, \\ dZ_t = f(\mathbf{X}_t, T - t)Y_tdt + \psi(\mathbf{X}_t, T - t)Y_td\xi_t + Y_t\mathbf{F}^T d\mathbf{W}_t & Z_0 = 0, \\ d\xi_t = \mathbf{1}_{\{\mathbf{X}_t \in \partial\Omega_R\}}dt & \xi_0 = 0. \end{cases} \quad (7)$$

Above, $\mathbf{1}_{\{H\}}$ is the indicator function (1 if H is true and 0 otherwise); $\tau = \inf_t\{\mathbf{X}_t \in \partial\Omega_A\}$ is the first exit time (or first passage time) from Ω ; which takes place at the first exit point $\mathbf{X}_\tau \in \partial\Omega_A$; and ξ_t is called the local time.

The functions $\boldsymbol{\mu} : \Omega \mapsto \mathbb{R}^D$ and $\mathbf{F} : \Omega \mapsto \mathbb{R}^D$ —the former a consequence of Girsanov's theorem and the latter of the expectation of Ito's integral being zero—are at this point rather arbitrary. However, properly choosing the function $\boldsymbol{\mu}$ will later be crucial for the WoE numerical scheme.

In the remainder of the paper, \mathbf{F} is set to zero. (See [2] for an application where it is not.)

Many expected values pertaining to a population of particles can be accommodated into this framework—see [8]. For instance, if $p = g = c = \varphi = \psi = 0$, $f = 1$, then $u(\mathbf{x}_0, t)$ is the mean absorption time of a particle starting at \mathbf{x}_0 and driven by the drift \mathbf{b} and diffusion matrix σ in the interval $[0, T]$. Analogously, if $p = g = c = \varphi = f = 0$, $\psi = 1$, $u(\mathbf{x}_0, t)$ is the mean local time.

Elliptic equations. Equation (1) can be formally transformed into an elliptic BVP with mixed BCs by assuming that $\partial u / \partial t = 0$, thus dropping the dependence on time from u and all the coefficients; letting $T \rightarrow \infty$; and dropping the initial condition p . If $\partial\Omega_R = \emptyset$ (purely stopped diffusions / Dirichlet BCs), the stochastic representation derived from Theorem 1 still holds as long as $c \leq 0$

and $\mathbb{E}[\tau] < \infty$ [16, section 4.4.5]. (Note that in the purely reflected case, the latter condition would be impossible.) To the best of our knowledge, there is no rigorously proved stochastic representation for elliptic BVPs with mixed BCs. Therefore, we will assume in the remainder of this paper that, if: i) the time-independent equivalent conditions of those in Theorem 1 are in place; ii) $c \leq 0$; and iii) $\mathbb{E}[\tau] < \infty$; then the time-independent equivalent representation given by Theorem 1 holds.

3 Overview of Milstein’s method and the proposed implementation (Algorithm 1)

3.1 Notation

In order to approximate the SDE system (7) numerically, a small, constant timestep $h > 0$ is set, and each random realization of $(\mathbf{X}_t)_{0 \leq t \leq \min(\tau, T)}$ is replaced by a chain (which we may call a “trajectory”) $\hat{\mathbf{X}}_0 = \mathbf{x}_0, \hat{\mathbf{X}}_1, \dots, \hat{\mathbf{X}}_K$, where $\hat{\mathbf{X}}_k \approx \mathbf{X}_{kh}$, $k = 0, 1, \dots, K$, $\hat{\tau} := Kh \approx \tau$, and $t_k := kh$. The chains $\{\hat{Y}_k\}_{k=0}^K$, $\{\hat{Z}_k\}_{k=0}^K$ and $\{\hat{\xi}_k\}_{k=0}^K$ are defined analogously. In order to lighten the notation, we drop the hats and \mathbf{X}_k, Y_k etc. are to be understood as the discretized counterparts $\hat{\mathbf{X}}_k, \hat{Y}_k$, etc. unless stated otherwise. In general, functions with subindex k are evaluated at $(\mathbf{X}_k, T - t_k)$, like v.g. $c_k = c(\mathbf{X}_k, T - t_k)$.

Vectors are by default column vectors, and $\|\cdot\|$ is always the Euclidean norm. For $\mathbf{x} \in \mathbb{R}^D$ and $\partial\Omega$ smooth, let us define:

- $\mathbf{x}^{\partial\Omega} = \arg \min_{\mathbf{y} \in \partial\Omega} \|\mathbf{x} - \mathbf{y}\|$ (i.e. the closest point on the boundary).
- $|d(\mathbf{x})| = \|\mathbf{x} - \mathbf{x}^{\partial\Omega}\|$ (i.e. the Euclidean distance to the boundary).
- $\mathbf{N}(\mathbf{x}) = \mathbf{N}(\mathbf{x}^{\partial\Omega})$ (i.e. the normal vector pointing outwards).
- $\Pi(\mathbf{x}) = \{\mathbf{y} \mid \mathbf{N}^T(\mathbf{x})(\mathbf{y} - \mathbf{x}^{\partial\Omega}) = 0\}$ (i.e. the hyperplane tangent to $\partial\Omega$ at $\mathbf{x}^{\partial\Omega}$).

We use the shorthand notation $|d_k| := |d(\mathbf{X}_k)|$, $\mathbf{N}_k := \mathbf{N}(\mathbf{X}_k)$, and $\Pi_k := \Pi(\mathbf{X}_k)$.

$B(\mathbf{x}, R)$ stands for a D -dimensional ball of radius R centred at \mathbf{x} , and $\partial B(\mathbf{x}, R)$ for its surface. Let M be a $D \times D$ symmetric positive definite matrix with (real) spectrum given by $M\mathbf{m}_i = \mu_i\mathbf{m}_i$, $\mu_1 \geq \mu_2 \geq \dots \geq \mu_D > 0$. For $\rho > 0$, define the D -dimensional ellipsoid centred at \mathbf{x} and oriented according to M as

$$E_M^\rho(\mathbf{x}) = \{\mathbf{y} \mid (\mathbf{y} - \mathbf{x})^T M^{-1} (\mathbf{y} - \mathbf{x}) = \rho^2\}. \quad (8)$$

The major semiaxis of $E_M^\rho(\mathbf{x})$ has length $\rho\mu_1^{1/2}$ and points along \mathbf{m}_1 ³, and so successively until the minor semiaxis, which runs along \mathbf{m}_D with length $\rho\mu_D^{1/2}$. When $M = A(\mathbf{X}_k, T - t_k)$, we may simply write $E_k^\rho := E_{A(\mathbf{X}_k, T - t_k)}^\rho(\mathbf{X}_k)$.

³The eigenvalues and eigenvectors of M^{-1} are the inverse and the same as those of M , respectively.

In particular, there is one value of ρ , which we call $\bar{\rho}$, such that $E_M^{\bar{\rho}}(\mathbf{x})$ is tangent to $\partial\Omega$. We denote it as $\Sigma_M(\mathbf{x}) := E_M^{\bar{\rho}}(\mathbf{x})$ —respectively, $\Sigma_k := E_k^{\bar{\rho}}$. The surfaces of those ellipsoids are denoted as $\partial E_M^{\bar{\rho}}(\mathbf{x})$, $\partial E_k^{\bar{\rho}}$, $\partial\Sigma_M(\mathbf{x})$, and $\partial\Sigma_k$.

The eigenvalues of $A(\mathbf{x}, T-t)$ are $\lambda_1(\mathbf{x}, T-t) \geq \lambda_2(\mathbf{x}, T-t) \geq \dots \geq \lambda_D(\mathbf{x}, T-t) > 0$ (or simply $\lambda_1 \geq \dots \geq \lambda_D$ if the context permits.)

The notation $\omega \sim W$ means that ω is one realization drawn from a distribution W . In particular, \mathcal{S}_D is the uniform distribution of points over $\partial B(\mathbf{0}, 1)$ (recall it is D -dimensional); \mathcal{B} is the distribution taking ± 1 with equal probability; and $\mathcal{N}(m, s^2)$ is the Gaussian with mean m and variance s^2 .

3.2 Description of Milstein's method

Remark. For the sake of clarity, let us emphasize that what we refer to as MM is the combination of two different algorithms: one for purely stopped diffusions, called “Walk on Ellipsoids” (WoE) by Milstein, [26][27, section 6.3.2]; plus the reflection off Ω_R handled according to Milstein's one-step scheme [25][27, section 6.6] for purely reflected diffusions. Under sufficient smoothness of the boundary and the coefficients (leading to a $C^4([0, T] \times \bar{\Omega})$ solution of (1)), both algorithms enjoy proven $\mathcal{O}(h)$ weak convergence separately. The combination of both schemes (called MM by us) is theoretically analysed (and nowhere tested) in [27, theorem 6.2]; under the same constraints as above, it enjoys linear weak convergence rate as well. Hence, it furnishes a method for stopped/reflected SDEs and hence suitable for Monte Carlo solutions of BVPs with mixed BCs.

Let us now explain MM, starting with its first ingredient, WoE. In contrast with Euler-like integrators, WoE takes bounded steps in order to avoid overshooting the boundary. When \mathbf{X}_k is not interacting with it (in a sense that will be clarified in a moment), the drift is first removed from (7) by setting $\boldsymbol{\mu} = \sigma^{-1}\mathbf{b}$. Then, WoE takes $\mathbf{X}_{k+1} = \mathbf{X}_k + r\sigma\omega$, where $r = \sqrt{Dh}$ and $\omega \sim \mathcal{S}_D$, whence

$$\|\omega\|^2 = 1 = \frac{1}{r^2}(\mathbf{X}_{k+1} - \mathbf{X}_k)^T A_k^{-1}(\mathbf{X}_{k+1} - \mathbf{X}_k). \quad (9)$$

By (9), the possible values of \mathbf{X}_{k+1} are distributed over ∂E_k^r . As long as $t_k < T$, WoE has the sequence $\{\mathbf{X}_k\}_{k=1,2,\dots}$ hop over $\partial E_{0'}^r, \partial E_{1'}^r, \dots$ until coming close enough to the boundary (at $k = i$, say) that \mathbf{X}_i interacts with it. Then, one of the following happens: i) a reflection on $\partial\Omega_R$ takes place, yielding \mathbf{X}_{i+1} ; ii) a special ellipsoid must be taken for \mathbf{X}_{i+1} ; or iii) $\partial\Omega_A$ is deemed hit and the trajectory stopped.

If \mathbf{X}_i is so close to the absorbing boundary that E_i^r may intersect it (but still farther than $r^2 = Dh$), WoE dictates that the tangent ellipsoid Σ_i be taken instead for \mathbf{X}_{i+1} —in order to ensure that the trajectory cannot trespass the boundary. This may happen only when $|d_i| < r\sqrt{\lambda_1(\mathbf{X}_i, T-t_i)}$. Eventually, either the trajectory is stopped at the projection on $\partial\Omega_A$ (this happens when the distance to it is less than r^2), or the initial condition is hit (i.e. $t_k \geq T$).

If \mathbf{X}_i is closer than r to $\partial\Omega_R$, the one-step scheme in [27, section 6.6] is used to handle the normal reflection on the boundary. This involves a change of coordinates such that the new first component is pointing along $-\mathbf{N}_i$.

In a nutshell, Milstein’s MM is structured in the following way:

1. If $t_k \geq T$, the trajectory has ‘reached’ the initial condition without being stopped by $\partial\Omega_A$. Then, read the initial condition and finish.
2. If not, and the trajectory is closer than r^2 to $\partial\Omega_A$, deem it stopped. Read the Dirichlet BC and finish.
3. If neither of the above, and the trajectory is closer than r to $\partial\Omega_R$, perform the reflection, and continue.
4. Otherwise, hop onto the surface of E_k^r for \mathbf{X}_{k+1} —if that ellipsoid is fully contained in Ω . If not, hop onto the surface of $\Sigma_{k,r}$ and continue.

As it stands, MM cannot be used, for several points must be clarified first:

- the (fast) determination of the boundary data: $|d(\mathbf{x})|$, $\mathbf{N}(\mathbf{x})$, $\mathbf{x}^{\partial\Omega}$, and $\Pi(\mathbf{x})$.
- the (fast) determination of $\lambda_1(\mathbf{x}, T - t)$ for general diffusions,
- the (fast) determination of $\Sigma_{A(\mathbf{x}, T-t)}(\mathbf{x})$ for general domains, and
- the rotation needed for handling the reflections (in a fast way).

By “clarified”, we mean that specific recipes must be put in place for each of the above points. Very importantly, they must be cost-efficient, since they will be performed at potentially many time steps with each of the $N \gg 1$ realizations involved in the Monte Carlo estimate of the expectation (5).

We prove that all of the points above can be determined at $\mathcal{O}(D^2)$ cost without further assumptions than those for MM. Therefore, Algorithm 1—which is our implementation of MM and the main result of this paper—also has an overall complexity $\mathcal{O}(D^2)$ per time step, which is the minimum attainable order: that of computing the matrix-vector product $\sigma d\mathbf{W}_t^4$ in (7). For that reason, we claim that our algorithm is “fast”.

4 Detailed implementation

The proofs of all lemmas in this Section are compiled in Appendix A.

⁴The product of a lower triangular matrix by a vector involves 1 (first row) + 2 (second row) + ... + D (bottom row) = $D(D + 1)/2$ multiplications and as many additions; i.e. an $\mathcal{O}(D^2)$ cost.

Algorithm 1 A practical implementation of MM for general bounded diffusions

- 1: **Data:** $1 \gg h > 0$, $r = \sqrt{Dh}$, a signed distance map $d(\mathbf{x})$ such that $d(\mathbf{x} \in \overline{\Omega}) < 0$
- 2: Let $\mathbf{X}_0 = \mathbf{x}_0 \in \overline{\Omega}$, $Y_0 = 1$, $Z_0 = 0$, $\xi_0 = 0$, $t_0 = 0$, $d_0 < 0$, $\mathbf{X}_0^{\partial\Omega}$, \mathbf{N}_0 , and $k = 0$
- 3: **while** neither the initial condition nor $\partial\Omega_A$ have been hit **do**
- 4: **if** $t_k \geq T$ (initial condition hit) **then**
- 5: Let $\mathbf{X}_T = \mathbf{X}_k$, $Y_T = Y_k$, $Z_T = Z_k$, $\xi_T = \xi_k$, evaluate $p(\mathbf{X}_T)$ in (6) and **finish**
- 6: **else if** $|d_k| \leq r^2$ and $\mathbf{X}_k^{\partial\Omega} \in \partial\Omega_A$ (absorption) **then**
- 7: Let $\tau = t_k$, $\mathbf{X}_\tau = \mathbf{X}_k^{\partial\Omega}$, $Y_\tau = Y_k$, $Z_\tau = Z_k$, $\xi_\tau = \xi_k$, evaluate $g(\mathbf{X}_\tau, T - \tau)$ in (6) and **finish**
- 8: **else if** $|d_k| \leq r$ and $\mathbf{X}_k^{\partial\Omega} \in \partial\Omega_R$ **then**
- 9: Evaluate c_k, f_k at $(\mathbf{X}_k, T - t_k)$; let $\bar{\varphi} = \varphi(\mathbf{X}_k^{\partial\Omega}, T - t_k)$, $\bar{\psi} = \psi(\mathbf{X}_k^{\partial\Omega}, T - t_k)$
- 10: Compute Givens entries $\{\cos \theta_2, \dots, \sin \theta_D\}$ (Section 4.5 and App. B)
- 11: Compute $\hat{\mathbf{b}} = Q\mathbf{b}(\mathbf{X}_k, T - t_k)$, $(\hat{\varphi}'_1, \dots, \hat{\varphi}'_D)^T = Q\nabla\varphi(\mathbf{X}_k^{\partial\Omega}, T - t_k)$ and $(\hat{\psi}'_1, \dots, \hat{\psi}'_D)^T = Q\nabla\psi(\mathbf{X}_k^{\partial\Omega}, T - t_k)$ by (24)
- 12: Draw $\vec{v} = (v_1, \dots, v_{D-1})$ with $v_i \sim \mathcal{B}$
- 13: Compute $\Lambda_*\vec{v}$ according to Algorithm 2, using $\sigma(\mathbf{X}_k, T - t_k)$
- 14: Compute $(\hat{A}_{11}, \dots, \hat{A}_{D1})^T$, the first column of $QA(\mathbf{X}_k, T - t_k)Q^T$, by (39)
- 15: Let $\chi = (\chi_1, \chi_2, \dots, \chi_D)^T = \begin{cases} \chi_1 = \sqrt{\hat{A}_{11}r^2 + d_k^2} - |d_k| - \hat{b}_1r^2, \\ \chi_i = -\bar{\varphi}\hat{A}_{1i}r^2 + (\Lambda_*\vec{v})_{i-1}r, \text{ for } i = 2, \dots, D \end{cases}$
- 16: Perform the reflection according to [27, theorem 6.1]:

$$\begin{cases} \mathbf{X}_{k+1} = \mathbf{X}_k + Q^T(\chi + \hat{\mathbf{b}}r^2) & \text{(using (24), see also Appendix B)} \\ Y_{k+1} = Y_k + \left(c_k r^2 + \bar{\varphi}(1 + \bar{\varphi}|d_k|)\chi_1 - \left(\sum_{j=2}^D \hat{A}_{1j}\hat{\varphi}'_j \right) r^2 + \bar{\varphi}^2 \chi_1^2 \right) Y_k, \\ Z_{k+1} = Z_k + \left(f_k r^2 + \bar{\psi}(1 + \bar{\varphi}|d_k|)\chi_1 - \left(\sum_{j=2}^D \hat{A}_{1j}\hat{\psi}'_j \right) r^2 + \bar{\varphi}\bar{\psi}\chi_1^2 \right) Y_k, \\ \xi_{k+1} = \xi_k + r^2, \quad t_{k+1} = t_k + r^2. \end{cases} \quad (10)$$
- 17: **else**
- 18: **if** $|d_k| \leq r\sqrt{\lambda_1(\mathbf{X}_k, T - t_k)}$ **then**
- 19: Let $r_{k+1} = |d_k| / \|\sigma_k^T \mathbf{N}_k\|$ (tangent ellipsoid to hyperplane, i.e. $\Sigma_k \approx \tilde{\Sigma}_k$)
- 20: **else**
- 21: Let $r_{k+1} = r$ (regular ellipsoid E_k^r)
- 22: **end if**
- 23: Let $\omega_{k+1} \sim \mathcal{S}_D$, $\mu = \sigma^{-1}\mathbf{b}$ and take one step inside Ω [27, algorithm 3.7]:

$$\begin{cases} \mathbf{X}_{k+1} = \mathbf{X}_k + \sigma_k \omega_{k+1} r_{k+1}, \\ Y_{k+1} = Y_k + Y_k c_k r_{k+1}^2 / D + Y_k \mu^T \omega_{k+1} r_{k+1}, \\ Z_{k+1} = Z_k + Y_k f_k r_{k+1}^2 / D, \\ \xi_{k+1} = \xi_k, \quad t_{k+1} = t_k + r_{k+1}^2 / D. \end{cases} \quad (11)$$
- 24: **end if**
- 25: Update d_{k+1} and $(\mathbf{N}_{k+1}, \mathbf{X}_{k+1}^{\partial\Omega})$ according to (14), and let $k = k + 1$
- 26: Safeguard: **If** $d_{k+1} > 0$ (overshoot) **then** let $\mathbf{X}_{k+1} = \mathbf{X}_{k+1}^{\partial\Omega}$ and $d_{k+1} = 0$, **end**
- 27: **end while**

4.1 Gershgorin estimation of the largest eigenvalue

In order to avoid overshooting the boundary with a hop from \mathbf{X}_k , an upper bound of $\lambda_1(\mathbf{X}_k, T - t)$ is needed. In the event that there is no closed formula for it, accurately determining $\lambda_1(\mathbf{X}_k, T - t_k)$ (for every timestep of every trajectory in the Monte Carlo simulation) may add up to a prohibitive computational overhead. For that reason, Milstein's WoE takes a uniform upper bound $\tilde{\lambda}_1 \geq \max_{\mathbf{x} \in \Omega} \max_{0 \leq t < T} \lambda_1(\mathbf{x}, T - t)$. However, $\tilde{\lambda}_1$ itself may not be straightforward to estimate, either; or unnecessarily large, thus requiring too small an h in order to reach the required accuracy. In Algorithm 1, we propose substituting it by the non-uniform, sharper upper bound given below, with no detriment to δ .

Lemma 1 $\lambda_1(\mathbf{x}, T - t)$ can be bounded above at $O(D^2)$ cost by

$$\lambda_1(\mathbf{x}, T - t) \leq \max_{j=1, \dots, D} \sum_{i=1}^D |a_{ij}(\mathbf{x}, T - t)|. \quad (12)$$

The cost of this approach is thus $O(D)$ cheaper than extracting the spectrum.

4.2 Drift removal

In order to remove the drift from (7), $\boldsymbol{\mu} = \sigma^{-1} \mathbf{b}$ must be calculated, which may be costly, specially in high dimension—unless σ is lower triangular, where $\boldsymbol{\mu}$ can be computed easily by forward substitution at cost $O(D^2)$. In the solution of BVPs, σ can always be chosen lower triangular as the Cholesky matrix of $A(\mathbf{x}, t)$, because A is positive definite.

On the other hand, imagine that the data is a non-lower-triangular diffusion matrix $\hat{\sigma}$, and the point of the Monte Carlo calculation is to compute population densities such as those mentioned in Section 2 (of the mean first exit time, etc.). Then, $\hat{\sigma}$ should be replaced in (7) by the (lower triangular) Cholesky matrix of $\hat{\sigma} \hat{\sigma}^T$. In that scenario, it may be critical to obtain the lower triangular σ analytically before start, for computing the Cholesky factorization numerically at every time step involves an $O(D^3)$ cost per step [13]. In the remainder of the paper, we shall assume without loss of generality that σ is lower triangular.

4.3 Construction of the distance map

Determining the triple of boundary data $\{d(\mathbf{x}), \mathbf{x}^{\partial\Omega}, \mathbf{N}(\mathbf{x})\}$ for a point \mathbf{x} may be time-consuming when Ω has a nontrivial shape. In general, the distance function (or distance map) inside Ω obeys the Eikonal equation in \mathbb{R}^D

$$\|\nabla d(\mathbf{x})\| = 1, \quad u(\mathbf{x} \in \partial\Omega) = 0, \quad d(\mathbf{x} \in \Omega/\partial\Omega) < 0. \quad (13)$$

With the above convention that distances are negative inside Ω , it holds

$$\mathbf{N}(\mathbf{x}) = \nabla d(\mathbf{x}), \quad \mathbf{x}^{\partial\Omega} = \mathbf{x} - d(\mathbf{x}) \nabla d(\mathbf{x}). \quad (14)$$

Like in all schemes for bounded SDEs, the distance map must be calculated numerically prior to the Monte Carlo simulation. As proposed in [4], Sethian’s Fast Marching is the method of choice.

4.4 Construction of the tangent ellipsoid

At a given location \mathbf{X}_k sufficiently far from the boundary, WoE draws \mathbf{X}_{k+1} from the surface of the ellipsoid E_k^r , which is inscribed in the ball $B(\mathbf{X}_k, r\lambda_1^{1/2}(\mathbf{X}_k, T-t_k))$. When $|d_k| < r\lambda_1^{1/2}(\mathbf{X}_k, T-t_k)$, there is no guarantee that E_k^r does not stick out of $\partial\Omega$. Milstein’s WoE postulates that the tangent ellipsoid Σ_k be taken in that case (see [27, algorithm 3.7]). This prevents boundary overshoots while maximising the probability of \mathbf{X}_{k+1} being absorbed, thus leading to the least average number of hops—namely $O(1/h)$, see [27, section 6.4.3]—and hence to an optimally efficient algorithm. (Even though Σ_k might be larger than $E_r(\mathbf{X}_k)$, the asymptotic weak convergence rate is $O(h)$, as proven in [26].)

Therefore, the determination of $\Sigma(\mathbf{x})$ is needed for implementing Milstein’s WoE. This is now formally solved by Lemma 2 below—seemingly, a new result.

Lemma 2 *Let Ω be a closed domain in \mathbb{R}^D (not necessarily smooth). Then, $\Sigma_k = E_k^{\Psi(\mathbf{X}_k)} = E_k^{\bar{\rho}}$, where $\Psi(\mathbf{x})$ is the solution of the anisotropic Eikonal equation*

$$\|\nabla\Psi^T A_k \nabla\Psi\| = 1 \text{ and } \Psi > 0 \text{ in } \bar{\Omega}, \quad \Psi = 0 \text{ on } \partial\Omega. \quad (15)$$

Lemma 2 clarifies a fundamental issue of WoE. When A is a constant matrix, (15) need be solved just once before the simulation, and then $\Psi(\mathbf{x})$ will be evaluated in computing time—analogously to the Eikonal equation for the distance map. In fact, the solution to (15) is required only inside a narrow shell on the inner side of $\partial\Omega$. A Fast-Marching-like method for Lemma 2 can solve (15) only there without regard to the rest of Ω , thus cutting back on preprocessing overhead. There are, however, two caveats to using Lemma 2. Numerical methods for the anisotropic Eikonal equation are less developed (this is further commented on in Section 6). Moreover, in the event of a non-constant matrix $A(\mathbf{x}, T-t)$, solving (15) at every \mathbf{X}_k which needs it will in general be out of the question. For those reasons, we introduce the straightforward approximation of Lemma 3, which in Section 5 is shown to work very well. The idea is to replace Σ_k by the ellipsoid tangent to the closest tangent hyperplane, which we call $\tilde{\Sigma}_k$.

Lemma 3 *The ellipsoid centred at \mathbf{X}_k and tangent to the hyperplane tangent to $\partial\Omega$ at $\mathbf{X}_k^{\partial\Omega}$ is given by $\tilde{\Sigma}_k = \{\mathbf{y} \mid (\mathbf{y} - \mathbf{X}_k)^T A_k^{-1} (\mathbf{y} - \mathbf{X}_k) = \rho_{k+1}^2\}$, where*

$$\rho_{k+1} = \frac{|d_k|}{\|\sigma_k^T \mathbf{N}_k\|}. \quad (16)$$

When $(\partial\Omega \cap B(\mathbf{X}_k, |d_k|)) \subset \Pi_k$, the ellipsoid $E_k^{r_{k+1}}$ is fully inside Ω , where

$$r_{k+1} = \min(\rho_{k+1}, r). \quad (17)$$

Lemma 3 yields the required value r_{k+1} for the next hop (check lines 18-22 in Algorithm 1). (Note that A_k^{-1} is never used.) The cost of this construction is $O(D^2)$, due to the product $\sigma_k^T \mathbf{N}_k$.

When the tangency point between Π_k and $\tilde{\Sigma}_k$ lies outside of Ω , there is a nonzero probability that \mathbf{X}_{k+1} overshoots: this is the reason of the safeguard in Algorithm 1, line 26. If $\partial\Omega$ is smooth at $\mathbf{X}_k^{\partial\Omega}$, the portion of $\tilde{\Sigma}_k$ sticking out of $\partial\Omega$ tends to zero as $h \rightarrow 0^+$ (i.e. as $r \rightarrow 0^+$), since $\Pi_k \rightarrow \partial\Omega$ around $\mathbf{X}_k^{\partial\Omega}$. However, this may not be the case close to cusps or corners, specially if A_k is a very eccentric ellipsoid with the major semiaxis parallel to Π_k . In that worst case scenario, it is important that $\tilde{\Sigma}_k$ shrinks with h so that the overshooting probability tends to vanish. This is the point of the rule $r_{k+1} = \min(\rho_{k+1}, r)$ in Lemma 3: Algorithm 1 hops on $\tilde{\Sigma}_k$ only if it is smaller than E_k^r (i.e. if the tangency point is nearby); otherwise it sticks to E_k^r , knowing (by Lemma 3) that the probability of \mathbf{X}_{k+1} overshooting goes asymptotically to zero. We note, however, that this is an heuristic reasoning rather than a rigorous proof that the construction given by Lemma 3 preserves $\delta = 1$ from MM in presence of curved boundaries.

4.5 Change of coordinates close to the reflecting boundary

Notation. In this subsection, column vectors in \mathbb{R}^D and \mathbb{R}^{D-1} are respectively written in bold (like \mathbf{N}_k) and with arrows (like \vec{N}). Given a $D \times D$ matrix (like σ), a starred matrix (like σ_*) denotes the submatrix obtained by removing the first column and row from the former. Matrix (or vector) elements superfluous for the discussion are depicted by $*$ (like $\mathbf{N}_k^T = (*, \vec{N}^T)$). I_S stands for the identity matrix in dimension S ; and \mathbf{e}_n is the n^{th} column of I_D . We drop subindex k except for \mathbf{X}_k , $\mathbf{X}_k^{\partial\Omega}$, d_k , $\nabla\psi_{k,r}$ and \mathbf{N}_k . All numbers are real.

When the trajectory \mathbf{X}_k is closer than r to $\partial\Omega_R$, coordinates are locally changed so that the new origin is \mathbf{X}_k and the new first component points towards $-\mathbf{N}_k$ (i.e. inwards). The rotation is thus defined as

$$Q(\mathbf{X}_k^{\partial\Omega} - \mathbf{X}_k) = (-|d_k|, 0, \dots, 0)^T, \quad (18)$$

where Q is an orthogonal matrix, i.e. $Q^T = Q^{-1}$ (not unique, in general). In order to perform the rotations in an efficient way, we adapt the approach in [6]. The Givens matrix⁵ $G^\theta(i, j)$ (where $i < j$) is defined element-wise as

$$[G^\theta(i, j)]_{kl} = \delta_{kl}, \text{ except: } \begin{cases} [G^\theta(i, j)]_{ii} = \cos \theta, & [G^\theta(i, j)]_{ij} = \sin \theta, \\ [G^\theta(i, j)]_{ji} = -\sin \theta, & [G^\theta(i, j)]_{jj} = \cos \theta. \end{cases} \quad (19)$$

Givens matrices are orthogonal and have two properties of interest to us:

1. Let $\mathbf{v} = (v_1, \dots, v_D)$ be an arbitrary vector. There is an angle $\theta(\mathbf{v})$, given by

⁵Householder transformations could be used as well, see [6].

$$\begin{cases} \cos \theta(\mathbf{v}) = 1, & \sin \theta(\mathbf{v}) = 0 & \text{if } v_i = v_j = 0, \\ \cos \theta(\mathbf{v}) = \frac{v_i}{\sqrt{v_i^2 + v_j^2}}, & \sin \theta(\mathbf{v}) = \frac{v_j}{\sqrt{v_i^2 + v_j^2}}, & \text{otherwise,} \end{cases} \quad (20)$$

such that the action of $G^{\theta(\mathbf{v})}(i, j)$ on a vector zeroes its j^{th} element, may change the i^{th} one, and leaves the rest unchanged:

$$\mathbf{v}' = G^{\theta(\mathbf{v})}(i, j)\mathbf{v} = (v_1, \dots, v_{i-1}, v'_i, v_{i+1}, \dots, v_{j-1}, v'_j = 0, v_{j+1}, \dots, v_D)^T.$$

2. If $v_i = v_j = 0$, and $\mathbf{v}' = G^\beta(i, j)\mathbf{v}$, then $v'_i = v'_j = 0$ for any $\beta \in \mathbb{R}$.

Let $\mathbf{v}^{(1)}$ be a vector and $\mathbf{v}^{(2)} = G^{\theta_2}(1, 2)\mathbf{v}^{(1)}$, where $\theta_2 = \theta(\mathbf{v}^{(1)})$ according to (20). Then, by construction, the second element of $\mathbf{v}^{(2)}$ is zero and every other one is unchanged except for the first one. Similarly, the second and third elements of $\mathbf{v}^{(3)} = G^{\theta_3}(1, 3)\mathbf{v}^{(2)}$ with $\theta_3 = \theta(\mathbf{v}^{(3)})$ per (20) are zero. Iterating, it is clear that

$$G^{\theta_D}(1, D) \dots G^{\theta_3}(1, 3)G^{\theta_2}(1, 2)\mathbf{v}^{(1)} =: G\mathbf{v}^{(1)} = (\beta, 0, \dots, 0)^T. \quad (21)$$

Performing the rotations in the above order has the following properties:

Lemma 4 Let $\|\mathbf{v}^{(1)}\| > 0$, and G be the matrix product of the Givens rotations $G = \prod_{k=D}^2 G^{\theta_k}(1, k)$ as in (21), i.e. zeroing the vector elements from the second to the last, starting from $\mathbf{v}^{(1)}$. Then: (i) G_* is lower triangular; (ii) $\beta = \|\mathbf{v}^{(1)}\| > 0$.

In order for $Q\mathbf{N}_k = (-1, 0, \dots, 0)$, we set $\mathbf{v}^{(1)} = \mathbf{N}_k$, sequentially construct $G^{\theta_2}(1, 2), \dots, G^{\theta_D}(1, D)$, and set

$$Q = -G^{\theta_D}(1, D) \dots G^{\theta_2}(1, 2), \quad (22)$$

and therefore

$$\begin{aligned} Q^{-1} &= -[G^{\theta_2}(1, 2)]^{-1} \dots [G^{\theta_D}(1, D)]^{-1} \\ &= -[G^{\theta_2}(1, 2)]^T \dots [G^{\theta_D}(1, D)]^T \\ &= -[G^{-\theta_2}(1, 2)] \dots [G^{-\theta_D}(1, D)]. \end{aligned} \quad (23)$$

In order to premultiply a vector, Q (or Q^T) need not be formed, but the two sets $\{\cos \theta_2, \dots, \cos \theta_D\}$ and $\{\sin \theta_2, \dots, \sin \theta_D\}$ are calculated and stored in advance, and later used whenever needed. To signify that the sequence of Givens rotations is performed on \mathbf{v} , without ever constructing Q , we shall write

$$Q\mathbf{v} = -\text{ROTATIONS}(\mathbf{v}, \theta_2 \rightarrow \theta_D), \quad Q^T\mathbf{v} = -\text{ROTATIONS}(\mathbf{v}, -\theta_D \rightarrow -\theta_2). \quad (24)$$

The cost of carrying out the $D - 1$ Givens rotations sequentially is just $\mathcal{O}(D - 1)$, instead of $\mathcal{O}((D - 1)^2)$ as would be the case for an explicit matrix-vector multiplication. Several vectors in Algorithm 1 need to be rotated while computing the reflection. For instance, $(\hat{\psi}'_1, \dots, \hat{\psi}'_D)^T = Q\nabla\psi_k = -\text{ROTATIONS}(\nabla\psi_k, \theta_2 \rightarrow \theta_D)$.

Let us now focus on the following two parts of matrix A in the rotated frame:

$$\hat{A} = QAQ^T = \begin{pmatrix} \hat{A}_{11} = (QAQ^T)_{11} & * \\ * & \hat{A}_* = (QAQ^T)_* \end{pmatrix}. \quad (25)$$

Because A is positive definite, so are QAQ^T (since Q is orthogonal) and⁶ $(QAQ^T)_*$, and $(QAQ^T)_{11} > 0$. Then $\hat{A}_{11} > 0$ and $\hat{A}_{11}r^2 + d_k^2 > 0$ in Algorithm 1.

Updating the decomposition of the rotated submatrix. Let us now address another important computational aspect, not discussed in MM. It is required to compute the matrix-vector product $\Lambda_*\vec{v}$, where Λ_* is defined by

$$\hat{A}_* = \Lambda_*\Lambda_*^T. \quad (26)$$

Thanks to the positive-definiteness of \hat{A}_* , Λ_* could be obtained by Cholesky factorization at a cost $\mathcal{O}((D-1)^3)$. This is taxing if \mathbf{N}_k is not constant on $\partial\Omega_R$ and/or in high dimension. Fortunately, the factorization $A = \sigma\sigma^T$ can be efficiently recycled into Λ_* —although, as it will be shown next, not necessarily in the standard way. The following notation will be convenient:

$$Q = \begin{pmatrix} Q_{11} & \vec{q}_1^T \\ \vec{q}_2 & Q_* \end{pmatrix}, \quad \sigma = \begin{pmatrix} \sigma_{11} & \vec{0}^T \\ \vec{s} & \sigma_* \end{pmatrix}. \quad (27)$$

Since σ is nonsingular and lower triangular by construction (see Section 4.2), $\sigma_{11} \neq 0$, $\det \sigma_* \neq 0$, and σ_* is lower triangular. On the other hand, Q_* need not be orthogonal or even nonsingular. Writing out $\hat{A} = Q\sigma\sigma^TQ^T$ gives

$$\Lambda_*\Lambda_*^T = Q_*\sigma_*\sigma_*^TQ_*^T + \vec{w}\vec{w}^T, \quad \text{where } \vec{w} = \sigma_{11}\vec{q}_2 + Q_*\vec{s}. \quad (28)$$

Therefore, $\Lambda_*\Lambda_*^T$ can be regarded as the factorization $(Q_*\sigma_*)(Q_*\sigma_*)^T$ plus the rank-one update $\vec{w}\vec{w}^T$. Borrowing a standard ansatz from [13], we set

$$\Lambda_* = Q_*\sigma_*(I_{D-1} + \alpha\vec{z}\vec{z}^T) \quad (29)$$

and look for suitable α and \vec{z} . Inserting (29) into (28) and noting that $(I_{D-1} + \alpha\vec{z}\vec{z}^T)(I_{D-1} + \alpha\vec{z}\vec{z}^T)^T = (I_{D-1} + \alpha\vec{z}\vec{z}^T)^2 = I_{D-1} + \vec{z}(\alpha^2\vec{z}^T\vec{z} + 2\alpha)\vec{z}^T$, one has

$$(\alpha^2\vec{z}^T\vec{z} + 2\alpha)(Q_*\sigma_*\vec{z})(Q_*\sigma_*\vec{z})^T = \vec{w}\vec{w}^T, \quad (30)$$

which can be readily solved by letting $\alpha = \alpha_+$, where

$$\alpha^2\vec{z}^T\vec{z} + 2\alpha = 1 \Rightarrow \alpha_{\pm} = \frac{-1 \pm \sqrt{1 + \|\vec{z}\|^2}}{\|\vec{z}\|^2}, \quad (31)$$

(we choose α_+ for concreteness). Then, the vector \vec{z} is given by

⁶If M is positive definite, M_* too, since $0 < \mathbf{y}^T M \mathbf{y} = \vec{y}^T M_* \vec{y}$, for any $\mathbf{y} = (0, \vec{y}^T)^T$ such that $\|\vec{y}\| > 0$.

$$Q_*\sigma_*\vec{z} = \vec{w}. \quad (32)$$

Formulas (29)-(32) allow for calculation of $\Lambda_*\nu$ at $\mathcal{O}(D(D-1))$ cost. However, Lemma 5 below shows that this ansatz may fail.

Lemma 5 *Let $\mathbf{N}_k^T = (N_1, \vec{N}^T)$. The update formula $Q_*\sigma_*\vec{z} = \sigma_{11}\vec{q}_2 + Q_*\vec{s}$ (with Q_* , \vec{q}_2 , σ_{11} , σ_* , and \vec{s} from (27)) is inconsistent if and only if $N_1 = 0$. In that case, $\det Q_* = 0$. Otherwise, there is one unique \vec{z} , given by*

$$\sigma_*\vec{z} = \vec{s} - \frac{\sigma_{11}}{N_1}\vec{N}. \quad (33)$$

We stress that Lemma 5 is independent of the way in which the rotation is implemented. Swapping the first component for the first nonzero one of \mathbf{N}_k in (18) does not help, either: the permutation matrices involved induce structural changes (the analogous of σ_* is no longer lower triangular, for instance), with the result that Milstein's formulas in Algorithm 1 would have to be reworked. When the first component of \mathbf{N}_k is zero, Λ_* can still be computed at $\mathcal{O}((D-1)^2)$ cost by using Algorithm ZCHUD in LINPACK [10] (implemented as `cholupdate` in Matlab). This exploits Lemma 4: since $Q_*\sigma_*$ is lower triangular, Λ_* can be seen as a rank-one update of a Cholesky factorization. Despite the fact that $Q_*\sigma_*\sigma_*^T Q_*^T$ is only positive semidefinite if $N_1 = 0$, ZCHUD would also work in that case—thanks to the fact that one Cholesky matrix, namely $Q_*\sigma_*$, is available in the first place. However, using ZCHUD would involve forming the full matrices.

Instead, we put forward the following analytical approach, specifically tailored to the case $N_1 = 0$. It relies on the following decomposition:

$$\Lambda_*\Lambda_*^T = [Q_*\sigma_* \mid \vec{w}] \begin{bmatrix} \sigma_*^T Q_*^T \\ \vec{w}^T \end{bmatrix} = [Q_*\sigma_* \mid \vec{w}] Z Z^T \begin{bmatrix} \sigma_*^T Q_*^T \\ \vec{w}^T \end{bmatrix}, \quad (34)$$

where Z is an orthogonal matrix of order D . If it can be chosen such that

$$Z^T \begin{bmatrix} \sigma_*^T Q_*^T \\ \vec{w}^T \end{bmatrix} = \begin{bmatrix} H \\ \vec{0}^T \end{bmatrix}, \quad (35)$$

then $\Lambda_* = H^T$. It turns out that Z can be found analytically, leading to a closed formula for Λ_* .

Lemma 6 *Let $\mathbf{N}_k^T = (0, \vec{N}^T)$; Q_* , \vec{q}_2 , σ_{11} , σ_* , and \vec{s} from (27), and \vec{w} from (28). Then,*

$$Z = \begin{bmatrix} I_{D-1} - \frac{1}{\|\vec{r}\|^2} \vec{r} \vec{r}^T & \frac{\vec{r}}{\|\vec{r}\|} \\ \frac{\vec{r}^T}{\|\vec{r}\|} & 0 \end{bmatrix}, \quad (36)$$

where $\sigma_*\vec{r} = \vec{N}$. Furthermore, $\Lambda_* = Q_*\sigma_* + \frac{1}{\|\vec{r}\|} \vec{w} \vec{r}^T$.

Remark. Λ_* in Lemma 6 is not necessarily triangular—but note that this was not required, anyways (check [27, formula 6.19], where Λ_* is called λ).

Let \vec{v} be an arbitrary column vector in \mathbb{R}^{D-1} . Observe that

$$Q \begin{pmatrix} 0 \\ \vec{v} \end{pmatrix} = \begin{pmatrix} * \\ Q_* \vec{v} \end{pmatrix} = -\text{ROTATIONS} \left(\begin{pmatrix} 0 \\ \vec{v} \end{pmatrix}, \theta_2 \rightarrow \theta_D \right), \quad (37)$$

$$Q^T \begin{pmatrix} 0 \\ \vec{v} \end{pmatrix} = \begin{pmatrix} * \\ Q_*^T \vec{v} \end{pmatrix} = -\text{ROTATIONS} \left(\begin{pmatrix} 0 \\ \vec{v} \end{pmatrix}, -\theta_D \rightarrow -\theta_2 \right). \quad (38)$$

Let $[QM]_n$ be the n^{th} column of QM , where M is a square matrix of order n . Then, $[QM]_n$ can be calculated as $[QM]_n = (QM)\mathbf{e}_n = -\text{ROTATIONS}(M\mathbf{e}_n, \theta_2 \rightarrow \theta_D)$. Furthermore, the first column of matrix \hat{A} is $QAQ^T\mathbf{e}_1 = -Q(\sigma\sigma^T\mathbf{N}_k)$, i.e.

$$\begin{pmatrix} \hat{A}_{11} \\ \vdots \\ \hat{A}_{D1} \end{pmatrix} = \hat{A}\mathbf{e}_1 = \text{ROTATIONS}(\sigma\sigma^T\mathbf{N}_k, \theta_2 \rightarrow \theta_D). \quad (39)$$

The concrete calculation of $\Lambda_*\vec{v}$ with cost $\mathcal{O}(D(D-1))$ is listed as Algorithm 2. Complexity is dominated by the product $\sigma_*\vec{v}$ and the forward substitutions.

Algorithm 2 Fast computation of $\Lambda_*\vec{v}$ in Algorithm 1 (line 13)

- 1: **Data:** σ, \vec{v} , Givens angles $\{\cos \theta_2, \dots, \sin \theta_D\}$, $\mathbf{N}_k = (N_1, \vec{N})^T$, tiny $tol > 0$
 - 2: Let $\begin{pmatrix} * \\ \vec{q}_2 \end{pmatrix} = -\text{ROTATIONS}((1, 0, \dots, 0)^T, \theta_2 \rightarrow \theta_D)$
 - 3: Compute $Q_*(\sigma_*\vec{v})$ and $Q_*\vec{s}$ by (37) and let $\vec{w} = \sigma_{11}\vec{q}_2 + Q_*\vec{s}$
 - 4: **if** $|N_1| > tol$ **then**
 - 5: Compute \vec{z} from $\sigma_*\vec{z} = \vec{s} - \frac{\sigma_{11}}{N_1}\vec{N}$ by forward substitution
 - 6: $\Lambda_*\vec{v} = Q_*(\sigma_*\vec{v}) + \left(\sqrt{1 + \|\vec{z}\|^2} - 1 \right) \frac{\vec{z}\vec{v}}{\|\vec{z}\|^2} \vec{w}$
 - 7: **else**
 - 8: Compute \vec{r} from $\sigma_*\vec{r} = \vec{N}$ by forward substitution
 - 9: $\Lambda_*\vec{v} = Q_*(\sigma_*\vec{v}) + \frac{\vec{r}\vec{v}}{\|\vec{r}\|} \vec{w}$
 - 10: **end if**
-

5 Numerical experiments

In this section, we report numerical results obtained with Algorithm 1—henceforth, they are labelled as “MM”. The algorithm of Constantini *et al.* has also been used for comparison (“REF”). (Specifically, the version with pseudonormal variables, which has a weak convergence rate $\mathcal{O}(h^{1/2-\epsilon})$, $\epsilon > 0$ [8].) Both codes have been written in Matlab, and run on a laptop. In that case (even though

Algorithm 1 is written in sequential form), the codes should be fully vectorized, which is critical for speed (check appendix B).

The Matlab code and data files used for this paper are available at the journal repository, and upon request from the author.

Let $\phi_h^{(1)}, \dots, \phi_h^{(N)}$ be N iid numerical approximations to ϕ in (5). The numerical approximation to the Feynman-Kac functional is

$$u_{h,N}(\mathbf{x}_0, t) = \frac{1}{N} \sum_{j=1}^N \phi_h^{(j)}. \quad (40)$$

Asymptotically (i.e. as $N \rightarrow \infty$ and $h \rightarrow 0^+$), $\varepsilon_{N,h} \sim \mathcal{N}(Ch^\delta, \mathbb{V}[\phi_h]/N)$, where $\varepsilon_{N,h}$ is the root mean-square (RMS) error of ϕ_h , $C > 0$ is a constant and $\delta > 0$ is the weak convergence rate of the scheme [27]. In all the ensuing experiments, the pointwise exact solution $u_{ex}(\mathbf{x}_0, T) = \mathbb{E}[\phi] \neq 0$ is known. Then, after setting an accuracy goal ε , N is chosen accordingly, i.e. such that $2\sqrt{\mathbb{V}[\phi_h]/N} = .20 \times \varepsilon$. The idea is that asymptotically the RMS error carries less than a 20% statistical error with a large probability. However, for better assessment and comparison across experiments, the relative error $\varepsilon_h^{rel} := |1 - u_{h,N}(\mathbf{x}_0, T)/u_{ex}(\mathbf{x}_0, T)|$ is reported instead.

Remark. In practice, one starts simulating trajectories $j = 1, 2, \dots$, replaces $\mathbb{V}[\phi_h]$ by the sample estimate so far, V_j , and stops as soon as $\sqrt{V_j/j} \leq \varepsilon/10$.

5.1 Example I (three-dimensional)

We take this from [14]. The coefficients of (1) are: $c = 0$, $\mathbf{b} = (y, z, x)^T$, $\varphi = -\|\mathbf{x}\|^2$,

$$\sigma(x, y, z) = \begin{pmatrix} \sqrt{1+|z|} & 0 & 0 \\ \frac{1}{2}\sqrt{1+|x|} & \sqrt{\frac{3}{4}}\sqrt{1+|x|} & 0 \\ 0 & \frac{1}{2}\sqrt{1+|y|} & \sqrt{\frac{3}{4}}\sqrt{1+|y|} \end{pmatrix}. \quad (41)$$

By choosing as exact solution $u_{ex}(\mathbf{x}, t) = xyz$ (independent of time), and $g = p = u_{ex}$, the remaining coefficients ψ and f are derived from (1), i.e. $\psi = \mathbf{N}^T \nabla u_{ex} - \varphi u_{ex}$, and $f = y^2z + z^2x + x^2y + \frac{1}{2}\sqrt{1+|z|}\sqrt{1+|x|} + \frac{\sqrt{3}}{2}x\sqrt{1+|x|}\sqrt{1+|y|}$.

As in [14], we set $\Omega = B(\mathbf{0}, 1)$ and $\mathbf{x}_0 = (.56, .52, .30)^T$; thus $u_{ex}(\mathbf{x}_0, T) \approx .08736$. We consider three different sets of BCs: purely absorbing, purely reflecting, and mixed. In the latter case, the hemisphere with $z < 0$ is absorbing and the other one, reflecting. (Note that the reflection is not conormal.) The convergence of ε_h^{rel} w.r.t. h is shown in Table 1, along with the results with REF. Regardless of the BCs, MM is both more accurate and has a faster convergence rate (estimated by least-squares regression), in fact very close to the theoretical value $\delta = 1$. In many applications boundaries are not so smooth as on a sphere; we also solved this problem in the box $\Omega = [-\sqrt{2}/2, \sqrt{2}/2]^3$. We only show results for

h	absorbing BCs ($T = \infty$)		mixed BCs ($T = \infty$)		reflecting BCs ($T = 1$)	
	MM	REF	MM	REF	MM	REF
.0128	.07594	.33949	.29351	.27006	.26478	.21171
.0064	.04022	.25511	.15755	.20827	.13380	.16075
.0032	.02101	.18721	.08050	.15630	.07619	.12005
.0016	.00930	.13048	.04348	.11840	.03377	.09089
.0008	.00490	.10203	.01737	.08428	.01424	.06866
.0004	.00232	.07380	.01016	.06603	.00839	.05175
.0002	.00140	.04984	.00405	.04455	.00327	.04189
.0001	.00070	.04219	.00223	.03169	.00213	.03023
δ	0.98	0.44	1.02	0.44	1.03	0.40

Table 1: Convergence of ϵ_h^{rel} for Example I in $\Omega = B(\mathbf{0}, 1)$ at $\mathbf{x}_0 = (.56, .52, .30)^T$. (Entries carry about $\pm 20\%$ statistical error.)

the purely absorbing and purely reflecting cases; see Table 2. Computational times have been included.

Since the boundary is nonsmooth now (due to the corners), accuracy and convergence rate are bound to deteriorate. In particular, theoretical rates no longer apply. Nonetheless, MM is still the more efficient integrator. For instance—according to Table 2—REF took 340 s. to attain a relative error of .12722, while MM took just 65.7 s. for .12573 (with the same confidence interval). We stress that this is the case where MM performs worst (purely reflecting nonsmooth boundary).

h	absorbing BCs ($T = \infty$)				reflecting BCs ($T = 1$)			
	MM	time	REF	time	MM	time	REF	time
.0128	.03400	0.85	.08591	0.67	.45085	5.78	.70420	5.34
.0064	.01605	7.34	.06711	1.14	.27996	14.3	.53835	12.6
.0032	.00770	50.7	.05659	1.97	.18981	28.9	.40311	21.2
.0016	.00488	279	.04218	3.66	.12573	65.7	.30115	40.1
.0008	.00221	2957	.03130	13.7	.10147	115	.23036	88.0
.0004	.00106	25715	.02441	24.2	.07350	245	.15882	178
.0002	.00062	145236	.01930	72.6	.05324	987	.12722	340
.0001	.00027	849117	.01086	520	.04752	3076	.08534	644
δ	0.98		0.40		0.47		0.43	

Table 2: Convergence of ϵ_h^{rel} (within about $\pm 20\%$ statistical error) for Example I in $\Omega = [-\sqrt{2}/2, \sqrt{2}/2]^3$ at $\mathbf{x}_0 = (.56, .52, .30)$. Times in s. on a single processor.

5.2 Example II (arbitrary-dimensional)

This problem features crossed second derivatives, oscillating coefficients and solution, and can be posed in any dimension. The diffusion matrix is

$$\sigma = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 1 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \dots & 1 \end{pmatrix} \Rightarrow \sigma\sigma^T = A = \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 2 & \dots & 2 \\ 1 & 2 & 3 & \dots & 3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & \dots & D \end{pmatrix}. \quad (42)$$

The exact solution is $u_{ex} = \cos \sum_{i=1}^D x_i$. The coefficients are: $\mathbf{b} = (\sin \pi x_1, \dots, \sin \pi x_D)^T$, $c = 0$, $\varphi = -1$, and $g = u_{ex}$ (f and ψ can be derived from (1)). We take $T = \infty$ (an elliptic PDE, so that no initial condition is needed), and $\Omega = B(\mathbf{0}, 1)$ in \mathbb{R}^D .

h	MM (D=4)	MM (D=5)	REF (D=5)	MM (D=6)	MM (D=9)
.0064	.00829	.02507	.19833	.07132	.45630
.0032	.00358	.01320	.14175	.04110	.20576
.0016	.00132	.00655	.10550	.01727	.11527
.0008	.00067	.00371	.08002	.00914	.06331
.0004	.00032	.00162	.04708	.00512	.02679
.0002	—	.00087	.03566	.00207	.01468
.0001	—	.00043	.02916	.00131	.00642
.00005	—	—	.01832	.00061	.00399
.000025	—	—	.01248	—	.00169
δ	1.18	0.98	0.49	0.99	1.00

Table 3: Convergence of ϵ_h^R (within about $\pm 20\%$ statistical error) for the Example II in increasing dimension D . Solved in $\Omega = B(\mathbf{0}, 1)$ at $T = \infty$ (i.e. it is an elliptic problem) and at $\mathbf{x}_0 = (-\pi/4, 0, \dots, 0)$. BCs of mixed type (absorbing if $z < 0$). Missing entries were not calculated because they would take unacceptably long.

In Table 3, we consider the same mixed BCs as before (i.e. reflecting on the upper hemisphere). The solution with MM at $\mathbf{x}_0 = (-0.7854, 0, \dots, 0)^T$ is evaluated for increasing D . For comparison, the solution with REF in $D = 5$ is also included. Times are not included, but let us give an example: REF took 3119 s. for a relative accuracy of .01248, while MM (always in $D = 5$) took 50.9 s. for $\epsilon_h^{rel} = .01320$. As usual, the difference grows with increasing accuracy (as $\mathcal{O}(1/\epsilon_h^{rel})$, see [12]) due to the better δ of MM.

We use Example II to compare the cost per iteration of MM and REF. Specifically, we solve it in $B(\mathbf{0}, 1)$ with $T = 1$, same \mathbf{x}_0 , $N = 10^5$ (fixed), $h = .001$ and purely reflecting BCs (so that the number of steps is roughly constant). The total times taken by MM and REF are shown on the left of Figure 1, for $D = 3, \dots, 8$. This indicates that the mean cost per iteration of MM is less than twice as with REF.




Figure 1: (Left) Relative average cost per iteration of MM and REF. (Right) Contours of the solution to (45). Three tangent ellipses are overlaid, for which the major semiaxis is proportional to $\Psi(\mathbf{x})$.

All numerical experiments so far have made use of Gershgorin’s circles theorem (Section 4.1) for estimating λ_1 . We checked that doing this does not have a discernible impact on accuracy (compared with using the exact λ_1), while it cuts back on computational overhead (by about 15% – 20%, in these examples).

5.3 Example III (two-dimensional with sensitive corners)

The purpose of this problem is to investigate the connection between MM and the anisotropic Eikonal equation established in this paper. We pick $T = \infty$, $\mathbf{x}_0 = (.823, -.875)^T$ and purely absorbing BCs in $\Omega = [-1, 1]^2$. The coefficients are:

$$A = \begin{pmatrix} 8 & -2.71 \\ -2.71 & 1 \end{pmatrix} = \sigma\sigma^T \quad \Rightarrow \quad \sigma \approx \begin{pmatrix} 2.8284 & 0 \\ -0.9581 & 0.2863 \end{pmatrix}, \quad (43)$$

$c = 0$, $\mathbf{b} = (\sin \pi x, \sin \pi y)^T$, $g = u_{ex}$, and f is derived from the exact solution,

$$u_{ex}(x, y) = \frac{1}{x^2 + y^2 - R}. \quad (44)$$

The point of (44) is that, as $R \rightarrow 2$, the exact solution tends to infinity on the four corners, substantially compounding the difficulty of integrating the corresponding SDEs in the domain with nonsmooth boundary. Problem III is solved with MM and REF, as before, plus a new version of MM where, instead of using Lemma 3 for $\tilde{\Sigma}_k$, Σ_k is calculated exactly according to Lemma 2—we call this integrator “MM+”. In order to do so, the anisotropic Eikonal equation

$$\|\nabla\Psi^T A \nabla\Psi\| = 1 \text{ and } \Psi > 0 \text{ in } \bar{\Omega}, \quad \Psi = 0 \text{ on } \partial\Omega, \quad (45)$$

is solved before the Monte Carlo simulation. (Importantly, note that thanks to the fact that A is constant, only one such equation must be solved.) We do




Figure 2: Convergence of ϵ_h^{rel} (within $\pm 20\%$ statistical error) for Problem III and increasing difficulty at the corners. The dashed and solid segments (with slope 1 and $1/2$, respectively) are meant as reference. With $R = 2.0001$ neither MM nor MM+ have yet attained the asymptotic linear regime.

so using the Fast Sweeping scheme described in [32], and store the numerical solution in a 251×251 grid, from where $\Psi(\mathbf{X}_k)$ is later interpolated “on the fly”. (Lines 19–21 in Algorithm 1 must then be replaced by $r_{k+1} = \Psi(\mathbf{X}_k)$.)

The contours of Ψ are sketched on Figure 1 (right). Figure 2 shows the convergence of ϵ_h^{rel} for increasing difficulty: $R = \{2.1, 2.01, 2.0001\}$. Note that MM+ converges faster than MM, and both versions can tackle this problem much better than REF—which in the hardest case $R = 2.0001$, essentially cannot cope.

6 Conclusions

Numerical experiments show a very satisfactory performance of Algorithm 1, in terms of accuracy, speed, and weak convergence rate. We stress that it has a favourable complexity, rendering it suitable for high-dimensional problems.

The connection, established in this paper, with the anisotropic Eikonal equation, paves the way for the efficient numerical integration of general absorbed diffusions in domains with corners. (The only other method with this capability known to the author is the Walk on Rectangles algorithm [9], which is nonetheless restricted to the Brownian motion and some variations thereof with constant coefficients.) On the other hand, solving the anisotropic Eikonal equation is nontrivial and the object of intense current research. Tailoring existing solvers for the anisotropic Eikonal equation to the particular features of the

algorithm presented here is left as future work.

Finally, we point out that the most promising direction for further development of this algorithm is producing a Multilevel formulation, along the lines of [12].

Acknowledgements

The author thanks J. A. Acebrón for proposing this work and helpful discussions throughout. Portuguese FCT funding is gratefully acknowledged.

A Proofs of Lemmas in Section 4

Proof of Lemma 1. This is a direct application of Gershgorin's circles theorem: for any eigenvalue λ of any $D \times D$ matrix M , there is $1 \leq i \leq D$ such that

$$|\lambda - M_{ii}| \leq \sum_{i \neq j} |M_{ij}|. \quad (46)$$

Since A is positive definite, $\lambda > 0$. By the triangle inequality, $|\lambda - a_{ii}| \geq |\lambda| - |a_{ii}| = \lambda - |a_{ii}|$, and hence

$$\lambda_1(\mathbf{x}, T - t) \leq \max_{i=1, \dots, D} \sum_{j=1}^D |a_{ij}(\mathbf{x}, T - t)|. \quad (47)$$

The cost is that of adding up D elements for each of the D rows, and comparing those D values to find the largest one. \square

Proof of Lemma 2. Let M be a $D \times D$ symmetric positive definite matrix, with spectrum $M\mathbf{m}_i = \mu_i\mathbf{m}_i$, where $\mu_1 \geq \dots \geq \mu_D > 0$ and $\|\mathbf{m}_i\| = 1$. The (positive) viscosity solution of the anisotropic Eikonal equation

$$\|\nabla \Psi^T M \nabla \Psi\| = 1 \text{ in } \bar{\Omega}, \quad \Psi = 0 \text{ on } \partial\Omega, \quad (48)$$

at $\mathbf{x} \in \Omega$, namely $\Psi(\mathbf{x})$, can be interpreted as the arrival time at \mathbf{x} of a monotonic front which marches anisotropically with velocity μ_i along the direction \mathbf{m}_i (see [31] and references therein).

By symmetry, it is also the arrival time on $\partial\Omega$ of a front marching outwards from \mathbf{x} with the same velocities in the opposite directions (i.e. μ_i along $-\mathbf{m}_i$).

Then, at time $t' > 0$, the front stemming from \mathbf{x} is the ellipsoid with semiaxes of length $\mu_1 t', \dots, \mu_D t'$ along the orthogonal directions $\mathbf{m}_1, \dots, \mathbf{m}_D$. Since M and M^{-1} have the same eigenvectors and inverse eigenvalues, that ellipsoid is the locus of $(\mathbf{y} - \mathbf{x})^T M^{-1} (\mathbf{y} - \mathbf{x}) = t'^2$.

In particular, at time $t'' = \Psi(\mathbf{x})$, the ellipsoidal front arrives on the boundary for the first time, meaning that it is tangent to it. Therefore,

$$\Sigma_M(\mathbf{x}) = \{\mathbf{y} \mid (\mathbf{y} - \mathbf{x})^T M^{-1} (\mathbf{y} - \mathbf{x}) = \Psi^2(\mathbf{x})\}. \quad (49)$$

Setting $M = A_k$ and $\Sigma_k = E_{A(\mathbf{X}_k, T-t_k)}^{|\Psi(\mathbf{X}_k)|}(\mathbf{X}_k)$ yields the desired tangent ellipsoid. \square

Proof of Lemma 3. Without loss of generality, let us take the origin at \mathbf{X}_k . Set $m(\mathbf{y}, \rho) = \mathbf{y}^T A_k^{-1} \mathbf{y} - \rho^2$. The sought-for ellipsoid is then $\widetilde{\Sigma}_k = \{\mathbf{y} \mid m(\mathbf{y}, \rho_{k+1}) = 0\}$. The tangency point \mathbf{y}_0 belongs both to $\widetilde{\Sigma}_k$ and to Π_k , so that for some $q \neq 0$,

$$\nabla m(\mathbf{y}_0, \rho_{k+1}) = 2A_k^{-1} \mathbf{y}_0 = q \mathbf{N}_k. \quad (50)$$

Since $\det A_k \neq 0$, (50) implies

$$\mathbf{y}_0 = \frac{q}{2} A_k \mathbf{N}_k \Rightarrow \begin{cases} \frac{q^2}{4} \mathbf{N}_k^T A_k^T A_k^{-1} A_k \mathbf{N}_k = \rho_{k+1}^2, \text{ and} \\ \mathbf{N}_k^T \mathbf{y}_0 = \frac{q}{2} \mathbf{N}_k^T A_k \mathbf{N}_k. \end{cases} \quad (51)$$

Moreover, $A_k = A_k^T$ and hence

$$\rho_{k+1} = \frac{\mathbf{N}_k^T \mathbf{y}_0}{\sqrt{\mathbf{N}_k^T A_k^T \mathbf{N}_k}} = \frac{|d_k|}{\|\sigma^T \mathbf{N}_k\|}. \quad (52)$$

The condition $(\partial\Omega \cap B(\mathbf{X}_k, |d_k|)) \subset \Pi_k$ is equivalent to $\partial\Omega$ being planar inside $B(\mathbf{X}_k, |d_k|)$. Let us assume it is. Then, E_k^r sticks out of $\partial\Omega$ iff it sticks out of Π_k . If it doesn't, then $\rho_{k+1} \geq r$. If it does, then $\widetilde{\Sigma}_k = E_k^{\rho_{k+1}}$ is concentric to and inside of E_k^r , so that $\rho_{k+1} < r$. In both cases, $E^{r_{k+1}}(\mathbf{X}_k)$ given by rule (17) is inside Ω . \square

Proof of Lemma 4. In preparation, let $\theta \in \mathbb{R}$, $1 < j, k \leq D$, and $\mathbf{w} = G^\theta(1, k) \mathbf{e}_j$. Then,

$$\text{if } k \neq j \text{ and } j > 1 \Rightarrow G^\theta(1, k) \mathbf{e}_j = \mathbf{e}_j. \quad (53)$$

To see this, simply note that $w_1 = (\mathbf{e}_j)_1 \cos \theta + (\mathbf{e}_j)_k \sin \theta$ and $w_k = -(\mathbf{e}_j)_1 \sin \theta + (\mathbf{e}_j)_k \cos \theta$ are both zero if $j > 1$ and $j \neq k$, while the other elements of \mathbf{w} are unaffected by the rotation and thus are the same as in \mathbf{e}_j .

Assume now $1 < i < j \leq D$. Then,

$$\begin{aligned} G_{ij} &= \mathbf{e}_i^T G \mathbf{e}_j \\ &= \mathbf{e}_i^T G^{\theta_D}(1, D) \cdots G^{\theta_{i+1}}(1, i+1) G^{\theta_i}(1, i) \cdots G^{\theta_2}(1, 2) \mathbf{e}_j \\ &= \left([G^{\theta_{i+1}}(1, i+1)]^T \cdots [G^{\theta_D}(1, D)]^T \mathbf{e}_i \right)^T \left(G^{\theta_i}(1, i) \cdots G^{\theta_2}(1, 2) \mathbf{e}_j \right) \\ &= \left(G^{-\theta_{i+1}}(1, i+1) \cdots G^{-\theta_D}(1, D) \mathbf{e}_i \right)^T \mathbf{e}_j \\ &= \mathbf{e}_i^T \mathbf{e}_j = 0. \end{aligned} \quad (54)$$

In (54), we have used the facts that: $[G^\theta(m, n)]^T = G^{-\theta}(m, n)$; that $\Pi_{k=i}^2 G^{\theta_k}(1, k) \mathbf{e}_j = \mathbf{e}_j$ because the interval $2, \dots, i$ does not include j , by assumption; and analogously $i+1, \dots, D$ does not include i . Since $G_{1 < i \leq D, 1 < j < i} = 0$, G_* is lower triangular, thus proving (i).

On the other hand, the upper row of G is not zero, in general. In fact, working out those entries (which is tedious but straightforward), one has:

$$\begin{aligned} G_{11} &= \cos \theta_2 \cos \theta_3 \cos \theta_4 \cdots \cos \theta_D, \\ G_{1k} &= \sin \theta_k \cos \theta_{k+1} \cos \theta_{k+2} \cdots \cos \theta_D, \quad (2 \leq k \leq D). \end{aligned} \quad (55)$$

Let us now prepare for (ii). First, let \mathbf{v} be a vector and $\mathbf{v}' = G^{\theta(\mathbf{v})}(1, k)\mathbf{v}$. By (20),

$$v'_1 = \sqrt{(v_1^{(1)})^2 + (v_k^{(k)})^2} \geq 0. \quad (56)$$

Consequently, in the sequence $\mathbf{v}^{(k)} = G^{\theta_k}(1, k)\mathbf{v}^{(k-1)}$ with $k = 2, \dots, D$, it holds that $v_1^{(2)} \geq 0, \dots, v_1^{(D)} \geq 0$. By (20), $\cos \theta_3, \dots, \cos \theta_D$ are thus all nonnegative. (Only $\cos \theta_2 < 0$ iff $v_1^{(1)} < 0$).

Clearly, $\beta^2 = (\mathbf{v}^{(1)})^T G^T G \mathbf{v}^{(1)} = (\mathbf{v}^{(1)})^T \mathbf{v}^{(1)} = \|\mathbf{v}^{(1)}\|^2 > 0$, by hypothesis. Also, it is always possible to pick n such that $v_n^{(1)} \neq 0$. Then, $v_n^{(1)} = \mathbf{e}_n^T \mathbf{v}^{(1)} = \mathbf{e}_n^T G^T G \mathbf{v}^{(1)} = \mathbf{e}_n^T G^T (\beta, 0, \dots, 0)^T = \beta G_{1n}$. This prevents G_{1n} from being zero, so that $\beta = v_n / G_{1n}$. Moreover, by the previous discussion and (55), $G_{1n} \neq 0$ implies that $\cos \theta_{n+1} \cdots \cos \theta_D > 0$.

Therefore, it is clear that, if $n = 1$, $\text{sign}(G_{11}) = \text{sign}(\cos \theta_2) = \text{sign}(v_1^{(1)})$, by (20) and (55).

If $1 < n \leq D$, $\text{sign}(G_{1n}) = \text{sign}(\sin \theta_n) = \text{sign}(v_n^{(n-1)})$, by (20). Let us show that $v_n^{(n-1)} = v_n^{(1)}$. This occurs because $\mathbf{v}^{(k)} = G^{\theta_k}(1, k)\mathbf{v}^{(k-1)}$ preserves all but the first and k^{th} elements of $\mathbf{v}^{(k-1)}$, so that if $n > k$, then $v_n^{(k)} = v_n^{(k-1)}$. Iterating, it is clear that $\mathbf{v}^{(n-1)} = G^{\theta_{n-1}}(1, n-1) \cdots G^{\theta_2}(1, 2)\mathbf{v}^{(1)}$ has preserved the elements $v_n^{(n-1)}, \dots, v_D^{(n-1)}$ from $\mathbf{v}^{(1)}$. Then, by (20), $\text{sign}(\sin \theta_n) = \text{sign}(v_n^{(n-1)}) = \text{sign}(v_n^{(1)})$.

Summing up, we have proved that $\text{sign}(\beta) = \text{sign}(v_n^{(1)} / G_{1n}) > 0$. In fact, $\beta = +\|\mathbf{v}^{(1)}\|$, which is (ii). \square

Proof of Lemma 5. We recall that $\det \sigma_* \neq 0$ and $\sigma_{11} \neq 0$ because $\sigma \sigma^T$ is positive definite. By definition $Q(N_1, \vec{N}^T)^T = (-1, \vec{0}^T)^T$ and hence

$$Q_{11}N_1 + \vec{q}_1^T \vec{N} = -1, \quad (57)$$

$$N_1 \vec{q}_2 + Q_* \vec{N} = \vec{0}. \quad (58)$$

Writing out $QQ^T = I_D$ and $Q^T Q = I_D$ in terms of the blocks defined in (27) yields

$$Q_{11}^2 + \vec{q}_1^T \vec{q}_1 = 1, \quad (59)$$

$$Q_{11}^2 + \vec{q}_2^T \vec{q}_2 = 1, \quad (60)$$

$$-Q_{11} \vec{q}_2 = Q_* \vec{q}_1, \quad (61)$$

$$-Q_{11} \vec{q}_1 = Q_*^T \vec{q}_2, \quad (62)$$

$$\vec{q}_2 \vec{q}_2^T + Q_* Q_*^T = I_{D-1}, \quad (63)$$

$$\vec{q}_1 \vec{q}_1^T + Q_*^T Q_* = I_{D-1}. \quad (64)$$

- Let us start by the ‘if’ part. Assume that $N_1 = 0$.

Since $\|\mathbf{N}_k\| = 1$, then $\vec{N}_k \neq \vec{0}$. From (58), $Q_*\vec{N} = \vec{0}$, and therefore $\det Q_* = 0$. Applying Sylvester’s determinant identity to (63) gives:

$$\det(Q_*Q_*^T) = \det(I_{D-1} - \vec{q}_2\vec{q}_2^T) = 1 - \vec{q}_2^T\vec{q}_2 \Rightarrow \vec{q}_2^T\vec{q}_2 = 1. \quad (65)$$

Then, by (60), $Q_{11} = 0$. By (61), this means that $\vec{q}_2^T Q_* = \vec{0}^T$, so that \vec{q}_2 is not in the range of Q_* (i.e. it is orthogonal to the column space of Q_*^T). On the other hand, $Q_*(\sigma_*\vec{z})$ has a projection on \vec{q}_2 , since $\vec{q}_2^T Q_*(\sigma_*\vec{z}) = \sigma_{11} \neq 0$. Consequently, there cannot be such $\sigma_*\vec{z}$ if $N_1 = 0$.

- Assume now that the system is inconsistent. Since $\sigma_{11} \neq 0$ and $Q_*\vec{s}$ in \vec{w} clearly belongs to the range of Q_* , it follows that $\vec{q}_2 \neq 0$ does not. This implies $\vec{q}_2^T Q_* = \vec{0}^T$, and by (58), $N_1 = 0$. This proves the “only if” part.
- Finally, assume there is a solution. It has been proved that then $N_1 \neq 0$ and $Q_{11} \neq 0$. By Sylvester, $\det(Q_*Q_*^T) = 1 - \vec{q}_2^T\vec{q}_2 = Q_{11}^2$, whence $\det Q_* \neq 0$. This means that $\sigma_*\vec{z}$, and thus \vec{z} , is unique.

If $N_1 \neq 0$, then $-\frac{\sigma_{11}}{N_1}Q_*\vec{N} = \sigma_{11}\vec{q}_2$ by (58). It follows that $Q_*(\vec{s} - \frac{\sigma_{11}}{N_1}\vec{N}) = \vec{w}$. \square

Proof of Lemma 6. Let us try to fit (35) with

$$\begin{bmatrix} \Phi & \vec{r} \\ \vec{r}^T & 0 \end{bmatrix} \begin{bmatrix} \sigma_*^T Q_*^T \\ \vec{w}^T \end{bmatrix} = \begin{bmatrix} \Phi\sigma_*^T Q_*^T + \vec{r}\vec{w}^T \\ \vec{r}^T\sigma_*^T Q_*^T \end{bmatrix}. \quad (66)$$

In order for the bottom row of the rightmost matrix above to vanish, $Q_*\sigma_*\vec{r} = \vec{0}$. Since $\det\sigma_* \neq 0$, this has a nontrivial solution iff $\det Q_* = 0$, i.e. iff $N_1 = 0$, by Lemma 5. In that case, by (58), $\sigma_*\vec{r} = \lambda\vec{N}$, with $\lambda \in \mathbb{R}$. (For definiteness, we pick $\lambda = 1$.)

The second requirement for a candidate Z is that it be orthogonal. By symmetry of the rightmost matrix in (66), it suffices to check that $Z^T Z = I$:

$$Z^T Z = \begin{bmatrix} \Phi\Phi^T + \vec{r}\vec{r}^T & \Phi\vec{r} \\ \vec{r}^T\Phi^T & \vec{r}^T\vec{r} \end{bmatrix} = \begin{bmatrix} I_{D-1} & \vec{0} \\ \vec{0}^T & 1 \end{bmatrix} \Leftrightarrow \begin{cases} \Phi\Phi^T = I_{D-1} - \vec{r}\vec{r}^T & (i) \\ \Phi\vec{r} = \vec{0} & (ii) \\ \|\vec{r}\| = 1 & (iii) \end{cases} \quad (67)$$

Assume for the time being that $\|\vec{r}\| = 1$. Let us take $\Phi = I_{D-1} - \vec{r}\vec{r}^T$. Then $\Phi = \Phi^T$ and $\Phi\Phi^T = \Phi^2 = I_{D-1} - 2\vec{r}\vec{r}^T + \vec{r}(\vec{r}^T\vec{r})\vec{r}^T = I_{D-1} - \vec{r}\vec{r}^T = \Phi$, meeting (i). For (ii), note that $\Phi\vec{r} = (I_{D-1} - \vec{r}\vec{r}^T)\vec{r} = \vec{r} - (\vec{r}^T\vec{r})\vec{r} = \vec{0}$. Finally, condition (iii) is met simply by substituting $\vec{r}/\|\vec{r}\|$ for \vec{r} (this makes the result independent of λ).

Therefore, Z in (36) is orthogonal, and yields $\Lambda_* = Q_*\sigma_*(I_{D-1} - \frac{1}{\|\vec{r}\|^2}\vec{r}\vec{r}^T) + \frac{1}{\|\vec{r}\|}\vec{w}\vec{r}^T$. Replacing \vec{r} by $\sigma_*^{-1}\vec{N}$ and using the fact that $Q_*\vec{N} = \vec{0}$ finishes the proof. \square

B Matlab code snippets

Vectorized functions are given below for ROTATIONS (Section 4.5), forward substitution with a lower triangular system, and distribution \mathcal{B} (Algorithm 1).

```
function varargout= rotations(opcion,varargin)
% [C,S]= rotations('init',vector) meaning that Givens(C,S)*vector=[1,0,...0]
% Q*vector= -rotations('forth',vector,C,S)
% Q'*vector= -rotations('back',vector,C,S)
X1= deal(varargin{1}); [N,dim]= size(X1);
switch opcion
case 'init'
C= zeros(N,dim-1); S= zeros(N,dim-1); %stored in same order as they are defined: ORD
for k=2:dim
wk= sqrt(X1(:,1).^2+X1(:,k).^2); C(:,k-1)= X1(:,1)./wk; S(:,k-1)= X1(:,k)./wk;
cero= find(wk==0); if ~isempty(cero), C(cero,k-1)= 1; S(cero,k-1)= 0; end
X2= X1; %and now overwrite two components:
X2(:,1)= C(:,k-1).*X1(:,1) + S(:,k-1).*X1(:,k);
X2(:,k)= -S(:,k-1).*X1(:,1) + C(:,k-1).*X1(:,k);
X1= X2; clear X2
end %vector X's only nonzero entry is the first one now.
varargout= {C,S}; return
case 'forth', rotations= [2:1:dim]; s1= +1; s2= -1; %order ORD
case 'back', rotations=[dim:-1:2]; s1= -1; s2= +1; %order ORD inverse
end
%Implement Givens rotations:
[C,S]= deal(varargin{2:3});
for k=rotations
X2= X1; %and now overwrite two components:
X2(:,1)= C(:,k-1).*X1(:,1) + s1*S(:,k-1).*X1(:,k);
X2(:,k)= s2*S(:,k-1).*X1(:,1) + C(:,k-1).*X1(:,k);
X1= X2;
end
varargout= {X2}; %rotated X1

function x= forsuvec(t,X,b) %vectorized forward substitution L*x=b
global PDEfile
[N,dim]= size(X); x= NaN(N,dim-1); %result
for k=1:dim-1
fila= feval(PDEfile,t,X,'sigmarow',k+1); fila= fila(:,2:end);
if k>1, x(:,k)= ( b(:,k) - sum(fila(:,1:k-1).*x(:,1:k-1),2) )./fila(:,k);
else x(:,k)= b(:,1)./fila(:,1); end
end

function [Eta]= distribucionboluda(N,dim)
Eta= randn(N,dim); R= sqrt(sum(Eta.^2,2)); Eta= Eta./(R*ones(1,dim));
```

References

- [1] C. Bayer, A. Szepessy and R. Tempone, *Adaptive weak approximation of reflected and stopped diffusions*. Monte Carlo Methods App. **16**, 1-67 (2010).
- [2] F. Bernal and J.A. Acebrón, *A multigrid-like algorithm for probabilistic domain decomposition*. Comput. Math. App. **72**(7) 1790–1810 (2016).

- [3] F. Bernal and J.A. Acebrón, *A comparison of higher-order weak numerical schemes for stochastic differential equations in bounded domains*. *Comm. Comput. Phys.* **20**(3) 703–732 (2016).
- [4] F. Bernal, J.A. Acebrón and I. Anjam, *A Stochastic Algorithm Based on Fast Marching for Automatic Capacitance Extraction in Non-Manhattan Geometries*. *SIAM J. Imaging Sci.* **7**(4), 2657–2674 (2014).
- [5] M. Bossy, E. Gobet, and D. Talay, *A symmetrized Euler scheme for an efficient approximation of reflected diffusions*. *J. App. Probability* **41**(3) 877–889 (2004).
- [6] F.M. Buchmann and W.P. Petersen, *An Exit Probability Approach to Solving High Dimensional Dirichlet Problems*. *SIAM J. Sci. Comput.* **28**, 1153–1166 (2006).
- [7] N. Cao, M. Ortner, A. Nehorai, *Solutions for diffuse optical tomography using the Feynman-Kac formula and interacting particle method*, *Proc. SPIE BIOS'07, San Jose, CA (USA) Vol. 6434* (2007) 643402–64313.
- [8] C. Constantini, B. Pacchiarotti and F. Sartoretto, *Numerical approximation for functionals of reflecting diffusion processes*. *SIAM J. Appl. Math.* **58**, 73–102 (1998).
- [9] M. Deaconu and A. Lejay, *A Random Walk on Rectangles Algorithm*, *Methodol. Comput. Appl. Probab.* **8**, 135–151 (2006).
- [10] J. J. Dongarra, J. R. Bunch, C. B. Moler, and G. W. Stewart. *LINPACK Users' Guide*. SIAM, Philadelphia (1979).
- [11] M. Freidlin, *Functional Integration and Partial Differential Equations*. *Annals of Mathematics Studies*, vol. 109, Princeton University Press (1985).
- [12] M.B. Giles and F. Bernal, *Multilevel simulations of expected exit times and other functionals of stopped diffusions*. *SIAM / ASA J. Uncertainty Quantification*, **6**(4), 1454–1474 (2018).
- [13] P.E. Gill, G.H. Golub, W. Murray, and M. Saunders, *Methods for modifying matrix factorizations*. *Math. Comput.*, **28** 505–535, (1974).
- [14] E. Gobet and S. Menozzi, *Stopped diffusion processes: overshoots and boundary correction*. *Stoch. Proc. Appl.* **120**, 130–162, (2010).
- [15] E. Gobet, *Euler schemes and half-space approximation for the simulation of diffusions in a domain*. *ESAIM: Prob. Stat.* **5**, 261–297 (2001).
- [16] E. Gobet, *Monte Carlo Methods and Stochastic Processes*. Chapman & Hall/CRC (2016).
- [17] O.A. Ladyženkaja, V.A. Solonnikov, N. Ural'ceva, *Linear and Quasi-linear Equations of Parabolic Type*. *Trans. Math. Monographs* 23, AMS, Providence, RI (1968).
- [18] V. Lemaire and G. Pagès, *Multilevel Richardson-Romberg extrapolation*. *Bernoulli* **23**(4A), 2643–2692 (2017).
- [19] D. Lépingle, *Euler scheme for reflected stochastic differential equations*, *Math. Comput. Simul.* **38** 119–126 (1995).
- [20] N.A. Licata, S.W. Grill, *The first passage problem for diffusion through a cylindrical pore with sticky walls*, *Euro. Phys. J. E* **30**(4) 439–447 (2009).
- [21] G.M. Lieberman, *Mixed boundary value problems for elliptic and parabolic differential equations of second order*, *J. Math. Anal. App.* **113**(2), 422–440 (1986).
- [22] S. Mairé and M. Simon, *A partially reflecting random walk on spheres algorithm for electrical impedance tomography*, *J. Comput. Phys.*, **303** 413–430 (2015).

-
- [23] S. Mancini, F. Bernal and J.A. Acebrón, *An efficient algorithm for accelerating Monte Carlo approximations of the solution to boundary value problems*. J. Sci. Comput., **66**(2), 577–597 (2016).
- [24] M. Mascagni, N.A. Simonov, *Monte Carlo methods for calculating some physical properties of large molecules*, SIAM J. Sci. Comput. **26**, 339-357 (2004).
- [25] G.N. Milstein, *Application of the numerical integration of stochastic equations for the solution of boundary value problems with Neumann boundary conditions*, Theor. Prob. Appl. **41**, 170-177 (1996).
- [26] G.N. Milstein, *Weak Approximation of a Diffusion Process in a Bounded Domain*. Stoch. Stoch. Rep., **62** 147-200 (1997).
- [27] G.N. Milstein and M.V. Tretyakov, *Stochastic Numerics for Mathematical Physics*. Springer, Berlin (2004).
- [28] C. Miranda, *Partial Differential Equations of Elliptic Type*, 2nd ed. Springer-Verlag, Berlin/Heidelberg/New York (1970).
- [29] S. Redner, *A Guide to First-Passage Processes*. Cambridge University Press (2001).
- [30] J.T.C. Schwabedal and A. Pikovsky, *Phase Description of Stochastic Oscillations*, Phys. Rev. Lett. **110**(20), 204102 (2013).
- [31] M. Schwenke, A. Hennemuth, B. Fischer, and O. Friman, *A Novel Anisotropic Fast Marching Method and Its Application to Blood Flow Computation in Phase-Contrast MRI*. Methods Inf. Med. **51**(5) 423–428 (2012).
- [32] Y.R. Tsai, L.T. Cheng, S. Osher, and H.K. Zhao, *Fast sweeping algorithms for a class of Hamilton-Jacobi equations*. SIAM J. Numer. Anal. **41**(2), 673–694 (2003).