



Simulation of stopped diffusions [☆]

F.M. Buchmann *

Seminar for Applied Mathematics, Swiss Federal Institute of Technology, CH-8092 Zürich, Switzerland

Received 6 April 2004; received in revised form 13 July 2004; accepted 19 July 2004

Available online 26 August 2004

Abstract

In this work, we study standard Euler updates for simulating stopped diffusions. As an immediate application, we discuss the computation of first exit times of diffusions from a domain. We focus on one-dimensional situations and show how the ideas for the simulation of killed diffusions can be adapted to this problem. In particular, we give a fully implementable algorithm to compute the first exit time from an interval numerically. The Brownian motion case is treated in detail and extensions to general diffusions are given. Special emphasis is given to numerical experiments: For every ansatz, we include numerical experiments confirming the conjectured accuracy of our methods. Our algorithm is of order one in a weak sense. Comparisons with other algorithms are shown. Results that are superior to those obtained with other methods are presented. When approximating a first hitting time distribution, the results obtained with our algorithm are much better than those achieved with other methods.

© 2004 Elsevier Inc. All rights reserved.

PACS: 02.50.Ey; 02.60.Cb; 02.70.Lq; 07.05.Tp

MSC: 35J25; 60G40; 60H10; 65C05

Keywords: Stopped diffusions; Stochastic differential equations; Weak approximation; Dirichlet problems; Monte-Carlo simulations

1. Introduction

1.1. Motivation

The simulation of a stopped diffusion with high accuracy is of significant interest in many applications. Often, a good approximation of the first exit time of a stochastic process from a domain is needed to get good convergence in numerical simulation. A typical application is the probabilistic solution of Dirichlet

[☆] This work was supported in parts under the TH-Gesuch funds 0-20981-2.

* Tel.: +41 1 632 35 33.

E-mail address: fab@math.ethz.ch.

problems in bounded domains. There, applying the Feynman–Kac formula to get a probabilistic representation of the solution, the first exit time plays a crucial role. Roughly speaking, a simulation procedure works as follows: A path (a trajectory of a stochastic process) connected to the differential operator of the partial differential equation is simulated and one integrates along this path. The integration procedure has to be stopped when the path leaves the domain for the first time and the boundary condition is evaluated at this first exit point. Approximating the mathematical expectation by a finite mean over a (large) sample then yields the (point-wise) Monte-Carlo approximation to the solution of the Dirichlet problem.

We recall this formulation briefly and introduce some notation. Let D be a bounded domain in n -dimensional space with smooth boundary ∂D and consider the following boundary value problem (BVP). For simplicity, we focus on Poisson’s equation:

$$\frac{1}{2}\Delta u(x) + g(x) = 0, \quad x \in D, \quad u(x) = \psi(x), \quad x \in \partial D. \quad (1)$$

Consider the stochastic process

$$X_x(t) = x + \int_0^t dW(s), \quad x \in D, \quad (2)$$

where the integral is a stochastic integral in the sense of Itô and therefore $(X_x(t))_{t \geq 0}$ is a Brownian motion starting at x [1,2]. We introduce the first exit time of $(X_x(t))_{t \geq 0}$ from D :

$$\tau(x) = \inf\{t > 0 : X_x(t) \notin D\} = \inf\{t > 0 : X_x(t) \in \partial D\}. \quad (3)$$

The connection to the BVP (1) is given by the following version of the Feynman–Kac formula: The solution $u(x)$ has the stochastic representation (under some regularity and smoothness conditions on g , ψ and D , see [3])

$$u(x) = \mathbb{E} \left[\psi(X_x(\tau(x))) + \int_0^{\tau(x)} g(X_x(s)) ds \right]. \quad (4)$$

Sometimes, we find it more convenient to write $u(x) = \mathbb{E}_x[\psi(X(\tau)) + f(\tau)]$, where $df = g(X(t))dt$ with $f(0) = 0$. In this notation, the expectation is taken with respect to the measure \mathbb{P}_x connected to the solution of $dX = dW$ with $X(0) = x$ (and implicitly $\tau = \tau(x)$).

Clearly, the Feynman–Kac formulation (4) reveals its full strength in numerical simulations mainly (but not only) in high dimensions. Nevertheless, we concentrate on one-dimensional settings here, because: (i) the simple one-dimensional situation is already interesting in its own right and contains the main difficulties and (ii) we hope to be able to apply a big part of the ideas presented here also in higher dimensions. If n becomes large, the domains D are usually smooth with boundaries. Near to the boundary it looks flat. There, locally, the problem of a random walk approaching the boundary resembles to some extent that of the one-dimensional situation. However, for domains with corners the situation becomes more complex – but this topic will not be addressed here.

The algorithm we will construct exploits the fact that the stochastic differential equations (SDEs) need only be approximated numerically in a weak sense with a finite summation arithmetic mean approximating the expectation.

1.2. Difficulties in numerical simulation

At a first glance the numerical approximation of $u(x)$ using (2) and (4) involves only the numerical solution of SDEs and averaging over a large sample (Monte-Carlo method [4]). This is nowadays a standard procedure in many applications, see [5,6]. Nevertheless, if boundaries are involved, the situation is much more subtle.

The Euler scheme (or Euler–Maruyama scheme), due to its simplicity, is of great interest. Applied to above situation with a fixed time step of size h , it takes the form [5,6] $X_0 = x, f_0 = 0$ and

$$X_{k+1} = X_k + \Delta W_k \quad \text{and} \quad f_{k+1} = f_k + g(X_k)h \quad \text{for } k = 0, 1, \dots \quad (5)$$

Here, an n -vector $\Delta W_k = W(t_{k+1}) - W(t_k)$ of i.i.d. normal random variables with mean 0 and variance h (Gaussian random variables) is generated in each time step. We denote this distribution by the symbol $\mathcal{N}(0, h)$, $\Delta W_k \sim \mathcal{N}(0, h)$. The main difficulty presents itself: When should the (numerical) integration be stopped? In other words: How shall $X(\tau)$ and in particular τ be approximated? We shall concentrate on the approximation of τ in this paper, corresponding to a constant boundary condition ψ in (1).

For a simple exposition of the main concepts, we consider $D = (-\infty, b)$ with $x < b$ in what follows. The naive approach is to stop as soon as $X_k \geq b$ and to take as an approximation for the first hitting time of level b either $\tau \approx (k-1)h$, $\tau \approx kh$ or a certain value between these two values. The drawback of this approach is the loss of accuracy: Although the Euler scheme is of weak order one for a fixed final time T with $M+1$ discretization points (giving $h = T/M$ in our notation), the rate of convergence (even in the weak sense) in the presence of a boundary reduces to $\mathcal{O}(\sqrt{h})$, i.e., it is of weak order one half [7]. The use of exact Gaussian random variables for the increments of the Brownian motion in (5) causes the following important drawback: The resulting discrete time random walk is no longer restricted to the closure of the domain under consideration. In particular, $(X(t))_{t \geq 0}$ (which we try to approximate) might become larger than b within any temporal discretization subinterval: Although the discrete random walk resulting from the Euler approximation (5) is exact in distribution sense, it gives the process values only at discrete $t_k = kh$. In between, for $t_k < t < t_{k+1}$, we have no information on the behaviour of the continuous process $X(t)$ that we wish to approximate. It is well known [8,9,7] that in numerical simulation one has to take into account the fact that anywhere near the boundary the process might have left D and come back within step h : Even if both X_k and $X_{k+1} < b$, it is not unlikely that $X(t) \geq b$ for some $t \in (t_k, t_{k+1})$ – the process $X(t)$ might follow an excursion within h , implying $\tau < t_{k+1}$. Obviously, the trivial stopping procedure (stopping only if $X_k \geq b$) will overestimate τ , as no intermediate excursions are monitored.

1.3. An exit probability approach for killed diffusions

To overcome this problem, instead of the unbounded increments $\Delta W_k \sim \mathcal{N}(0, h)$, bounded approximations can be used [10,11], or a quantization approach is adequate, see [12] and references therein.

Nevertheless, applying the usual Euler scheme (with $\Delta W_k \sim \mathcal{N}(0, h)$) can have its advantages as well. To restore usual first-order convergence (in the weak sense), a simple hitting test was introduced by various authors, see [8,9] and references therein. This test has to be performed after each time step with $X_{k+1} < b$. It estimates the probability that an excursion occurred within (t_k, t_{k+1}) if both $X_k, X_{k+1} < b$ and leads to improved statistics.

We summarize the principal idea of this approach for *killed* diffusions: Let a *fixed* $T < \infty$ be given and suppose that we are interested in the approximation of $\mathbb{E}_x[F(X(T))\mathbf{1}_{T < \tau}]$ for some measurable F : Paths that reach level b up to (and including) time T are killed, that is, they do not contribute to the expectation. If we have, after an Euler step $t_k \rightarrow t_{k+1} = t_k + h < T$, that $X_{k+1} \geq b$ then, obviously, $\tau < T$ and the corresponding path is killed. To take into account a possible excursion across level b if $X_{k+1} < b$ one proceeds as follows: At the time the test is performed (*after* a step), X_{k+1} is known. Therefore, the bridge process [13, p. 67] pinned in time-space coordinates at (t_k, X_k) and at (t_{k+1}, X_{k+1}) has to be considered (and will be denoted by $X_{X_k, h, X_{k+1}}(s)$). To check for a possible excursion, an i.i.d. random number distributed uniformly in $(0, 1)$ (denoted by $u \sim \mathcal{U}$) is generated and the path is killed if

$$u \leq \mathbb{P} \left[\sup_{t_k \leq s \leq t_{k+1}} X_{X_k, h, X_{k+1}}(s) \geq b \right] = e^{-(2/h)(b-X_k)(b-X_{k+1})}, \quad u \sim \mathcal{U}. \tag{6}$$

Gobet proved that first-order weak convergence can be obtained for the Euler scheme when applying this test for killed diffusions in the presence of a boundary [7], see also [14].

1.4. Outline

The purpose of this work is to modify these ideas to the case of stopped (rather than killed) diffusions. In this case, we try to approximate expectations of the form $\mathbb{E}_x[F(X(\tau), \tau)]$ (see (4)). Our interest is hence in the actual value of τ rather than being satisfied by the assertion that (or if) $\tau < T$ for some predefined (deterministic) T . In other words, one wants to know (again in a statistical sense) *when* the first exit time actually took place – in contrast to asking only *if* the exit did already occur. To accomplish this task, we construct in a first stage the density of τ of the bridge process under consideration and sample in a later stage a random number from it. We show how a new interpretation of the exit probability of the bridge process (6) as a distribution leads to more accurate results (yet of the same order) for exactly the same computational cost. We then further improve our algorithm for the case that a discrete X_{k+1} falls outside D . In that case, clearly $\tau \leq t_{k+1}$. Nevertheless, we show how to find an approximation for $\tau \in (t_k, t_{k+1})$.

We start with the Brownian motion case in Section 2. The simplicity of this process will allow us to present our ideas precisely without obscuring details of notation. We then extend these ideas to general autonomous diffusions in Section 3. We always consider the two possible cases after a step: (i) $X_{k+1} \in D$ (in Sections 2.1 and 3.1, respectively) and (ii) $X_{k+1} \notin D$ (in Sections 2.2 and 3.2, respectively). We show results from numerical experiments in Section 4 where we first discuss a statistical study comparing various algorithms (Section 4.1) and later show results of some applications to the Feynman–Kac formulation (Section 4.2). We conclude in Section 5.

2. The Brownian motion case

To simplify notation, we write $y = X_k$ and $z = X_{k+1}$. Recall that for a Brownian motion application of the Euler scheme with step size $h > 0$ means that $z = y + \xi$ with $\xi \sim \mathcal{N}(0, h)$. In what follows, we denote the corresponding Brownian bridge pinned at (t_k, y) and at $(t_k + h, z)$ by $X_{y, h, z}(s)$ and its law by $\mathbb{P}_{y, h, z}[\cdot]$. Additionally, τ denotes the first hitting time of level b .

2.1. Inside: $y, z < b$ (test for an excursion)

Recalling (6) we find the distribution, F , of the first hitting time $\tau = \tau(y)$ w.r.t. Brownian bridge measure for $t > 0$ as

$$F(t) \equiv \mathbb{P}_{y, h, z}[\tau \leq t] = e^{-(2/t)(b-y)(b-z)}. \tag{7}$$

The idea is now to generate a random variable \mathcal{T}_1 with distribution (7). To this end, invert (7) [5, p. 12],

$$\mathcal{T}_1 = -\frac{2(b-y)(b-z)}{\log u}, \quad u \sim \mathcal{U}. \tag{8}$$

The path hit b between t_k and $t_k + h$ if $\mathcal{T}_1 \leq h$ (in a statistical sense). In that case, b was hit for the first time at $t = t_k + \mathcal{T}_1$ and we approximate $\tau \approx kh + \mathcal{T}_1$.

The application to the approximation of $f(\tau)$ using (5) (for example) is now straightforward and we show it only for this variant of our algorithm:

$$f(\tau) \stackrel{(5)}{\approx} h \sum_{i=0}^{k-1} g(X_i) + \mathcal{T}_1 g(X_k). \tag{9}$$

In what follows, we will not write down these approximations explicitly but only show how to generate the last summand (i.e., its length of integration).

2.2. *Outside: $y < b \leq z$ (compute first exit time)*

We first construct the needed density. Using absolute continuity of the measures \mathbb{P}_y and $\mathbb{P}_{y,h,z}$ we have [13, p. 67]

$$\mathbb{P}_{y,h,z}[\tau \in dt] = p(h; y, z)^{-1} p(h - t; b, z) \mathbb{P}_y[\tau \in dt],$$

where $p(t; x, y)$ denotes the Gaussian transition density: $p(t; x, y)dy = \mathbb{P}_x[W_t \in dy]$. Inserting [13, (1.2.0.2), p. 198] for $\mathbb{P}_y[\tau \in dt]$ gives with $y < b$

$$\mathbb{P}_{y,h,z}[\tau \in dt] = \frac{b - y}{\sqrt{2\pi t^3}} \sqrt{\frac{h}{h - t}} \exp\left(-\frac{(z - b)^2}{2(h - t)} + \frac{(z - y)^2}{2h} - \frac{(b - y)^2}{2t}\right) dt \tag{10a}$$

and after some algebra

$$\mathbb{P}_{y,h,z}[\tau \in dt] = (b - y) \sqrt{\frac{h}{2\pi t^3(h - t)}} \exp\left(-\frac{((b - y)h - t(z - y))^2}{2ht(h - t)}\right) dt. \tag{10b}$$

We remark that some simple manipulations in the exponent show that this formula reduces for $h = 1$ and $y = 0$ to [15, Formula (2.1) and Lemma 3].

We now show how to sample from (10). To simplify notation, we set $y = 0$ and $h = 1$. We say that a random variable X follows the *inverse Gaussian distribution* with parameters $\gamma > 0$, $\delta > 0$ (and write $X \sim \mathcal{IG}(\gamma, \delta)$) if it has the density [16]

$$\mathbb{P}[X \in dx] = \sqrt{\frac{\gamma}{2\pi x^3}} \exp\left(-\frac{\gamma(x - \delta)^2}{2\delta^2 x}\right) dx, \quad x > 0.$$

The basic observation is that if $X \sim \mathcal{IG}(b^2, b/(z - b))$ then $t = X/(1 + X)$ is a random variable with density (10) (with $h = 1$, $y = 0$).

To see this, define $p(t)$ for $0 < t < 1$ as $\mathbb{P}_{0,1,z}[\tau \in dt] = p(t)\mathbf{1}_{0 < t < 1} dt$ (see (10)). By the substitution $x = t/(1 - t) \geq 0$ with $dt = dx/(1 + x)^2$ we find

$$p(t)dt = \frac{b}{\sqrt{2\pi}} \frac{(1 + x)^2}{x\sqrt{x}} \exp\left(-\frac{(b - xz/(1 + x))^2}{2x/(1 + x)^2}\right) \frac{dx}{(1 + x)^2} = \frac{b}{\sqrt{2\pi x^3}} \exp\left(-\frac{b^2}{2x} \left(1 - \frac{z - b}{b} x\right)^2\right) dx.$$

The claim now follows immediately with $\gamma = b^2$ and $\delta = b/(z - b)$.

For the general bridge, we find analogously that if $X \sim \mathcal{IG}((b - y)^2/h, (b - y)/(z - b))$, then the random variable $t = hX/(1 + X)$ has density (10). Michael et al. presented an algorithm to generate random variables X s following the inverse Gaussian distribution [17, p. 89].

In the case that $z = X_{k+1} > b$, we therefore generate $X \sim \mathcal{IG}((b - y)^2/h, (b - y)/(z - b))$ using [17], set $\mathcal{T}_2 = hX/(1 + X)$ and stop integration at $t_k + \mathcal{T}_2$ (analogously to (9)).

3. Extension to general diffusions

We now expand the ideas presented for Brownian motion in Section 2 to general diffusions given by (compare with (2))

$$X_x(t) = x + \int_0^t \mu(X_x(s)) ds + \int_0^t \sigma(X_x(s)) dW(s), \quad \sigma(\cdot) > 0, \tag{11}$$

i.e., $X_x(t)$ solves the (autonomous) SDE $dX = \mu(X)dt + \sigma(X)dW$ with $X(0) = x$. The Euler approximation with step size $h > 0$ then reads (compare with (5))

$$X_0 = x \quad \text{and} \quad X_{k+1} = X_k + \mu(X_k)h + \sigma(X_k)\Delta W_k \quad \text{for } k = 0, 1, \dots \tag{12}$$

with corresponding continuous-time approximation (frozen coefficient approximation)

$$\bar{X}(t) = X_k + \mu(X_k)(t - t_k) + \sigma(X_k)(W(t) - W(t_k)), \quad t \in [t_k, t_{k+1}).$$

To derive our formulae, we therefore consider the constant coefficient diffusion

$$X_x^{\mu,\sigma}(t) = x + \mu t + \sigma W(t), \quad t > 0.$$

We further write $X_{y,h,z}^{\mu,\sigma}$ for the corresponding bridge pinned at y and z with length h and denote its law by $\mathbb{P}_{y,h,z}^{\mu,\sigma}$.

Remark 1. In the context of approximating killed diffusions, it was pointed out that this frozen coefficient approximation gives incorrect asymptotics and that more sophisticated approximations should be used [18,19]. The application of these ideas to stopped diffusions remains a topic of ongoing research.

3.1. Inside (test for an excursion)

Consider the function $f(u) = u/\sigma$ and define $D(t) = f(X_x^{\mu,\sigma}(t))$. By Itô’s formula [13,2], D satisfies $dD = \mu/\sigma dt + dW$ with $D(0) = x/\sigma$, i.e., $(D(t))_{t \geq 0}$ is a Brownian motion with drift $v = \mu/\sigma$ (starting at x/σ), see [13, I.IV.5. and II.2.]. As before let τ be the first hitting time of level $b > y$. Obviously,

$$\mathbb{P}_{y,h,z}^{\mu,\sigma}[\tau \leq h] = \mathbb{P} \left[\left(\sup_{0 \leq t \leq h} X_{y,h,z}^{\mu,\sigma}(t) \right) \geq b \right] = \mathbb{P} \left[\left(\sup_{0 \leq t \leq h} D_{f(y),h,f(z)}(t) \right) \geq f(b) \right],$$

where $D_{f(y),h,f(z)}$ denotes the version of D which is pinned at $f(y)$ and $f(z)$ and has length h . Now (we set $v = f(y)$, $w = f(z)$ and $c = f(b)$)

$$\mathbb{P} \left[\sup_{0 \leq t \leq h} D_{v,h,w}(t) \geq c \right] = \mathbb{P}_v \left[\sup_{0 \leq t \leq h} D(t) \geq c; D(h) \in dw \right] = \frac{\mathbb{P}_v \left[\sup_{0 \leq t \leq h} D(t) \geq c, D(h) \in dw \right]}{\mathbb{P}_v[D(h) \in dw]}.$$

Inserting [13, (2.1.0.6), p. 250 and (2.1.1.8), p. 251] yields

$$\mathbb{P}_{y,h,z}^{\mu,\sigma}[\tau \leq t] = e^{-(2/t\sigma^2)(b-z)(b-y)}, \tag{13}$$

which is equivalent to [7, p. 169]. Inverting (13) we get (compare with (8))

$$\mathcal{T}_1 = -\frac{2(b-y)(b-z)}{\sigma^2 \log u}, \quad u \sim \mathcal{U}, \tag{14}$$

where (we recall that) $\sigma = \sigma(y)$.

3.2. Outside (compute first exit time)

For $y < b < z$, we proceed similarly to Section 2.2. We have for $0 \leq t < h$ the density [20, (3.7), p. 371]

$$\mathbb{P}_{y,h,z}^{\mu,\sigma}[\tau \in dt] = \frac{p^{\mu,\sigma}(h-t; b, z)}{p^{\mu,\sigma}(h; y, z)} \mathbb{P}_y^{\mu,\sigma}[\tau \in dt], \quad (15)$$

where $p^{\mu,\sigma}(t; x, y)$ denotes the transition density of the solution to $dX = \mu dt + \sigma dW$. Solving Kolmogorov's forward equation one finds

$$p^{\mu,\sigma}(t; x, y) = \frac{\exp\left(-\frac{(y-\mu t-x)^2}{2t\sigma^2}\right)}{\sqrt{2\pi t\sigma^2}}.$$

To find the density of the first hitting time one solves for $\alpha > 0$ the differential equation [13, p. 18]

$$\left[\frac{\sigma^2}{2} \frac{d^2}{dx^2} + \mu \frac{d}{dx} - \alpha\right] u(x) = 0$$

and combines the increasing and decreasing solutions (denoted by u^\uparrow and u^\downarrow , respectively) to get the Laplace transform of τ

$$\mathbb{E}_y^{\mu,\sigma}[e^{-\alpha\tau}] = \begin{cases} u^\uparrow(y)/u^\uparrow(b), & y \leq b \\ u^\downarrow(y)/u^\downarrow(b), & y \geq b \end{cases} = \exp\left(\frac{\mu}{\sigma^2}(b-y) - \frac{\sqrt{\mu^2 + 2\alpha\sigma^2}}{\sigma^2}|b-y|\right).$$

Inverting it (using [13, Appendix]) we get the density (recall that $b > y$)

$$\mathbb{P}_y^{\mu,\sigma}[\tau \in dt] = \frac{b-y}{\sqrt{2\pi t^3\sigma^2}} \exp\left(-\frac{(b-\mu t-y)^2}{2t\sigma^2}\right) dt.$$

Inserting everything into (15) yields

$$\mathbb{P}_{y,h,z}^{\mu,\sigma}[\tau \in dt] = \frac{b-y}{\sqrt{2\pi\sigma^2 t^3}} \sqrt{\frac{h}{h-t}} \exp\left(-\frac{1}{2\sigma^2} \left(\frac{(z-b)^2}{(h-t)} - \frac{(z-y)^2}{h} + \frac{(b-y)^2}{t}\right)\right) dt.$$

Recalling (10), we thus generate $X \sim \mathcal{IG}((b-y)^2/(h\sigma^2), (b-y)/(z-b))$ and set $\mathcal{T}_2 = hX/(1+X)$ (where again $\sigma = \sigma(y)$).

4. Numerical experiments

We show results of extensive tests performed with the algorithm derived in the previous sections. For weak approximations, path-wise convergence is not required, but a good approximation of the distribution is important. In our case, special emphasis is on the computation of first exit times. We thus start (in Section 4.1) with a statistical test where we compare the numerically obtained density of a simple first hitting time (i.e., a histogram) with the known analytical density. We compare our algorithm with a variety of other approaches. At a later stage (in Section 4.2), we show the performance of our algorithm when applied to the numerical solution of some one dimensional Dirichlet problems via the stochastic representation of the solution.

4.1. Approximating the density of the first hitting time: a statistical comparison

We compute numerically the first hitting time (denoted by τ) of level $b = 1$ of a Brownian motion. The corresponding density is

$$\mathbb{P}_0[\tau \in dt] = \frac{e^{-(1/2t)}}{\sqrt{2\pi t^3}} dt, \quad t > 0. \tag{16}$$

It has its maximum at $t = 1/3$, where it forms a non-symmetrically shaped peak, and it has a very long tail, see Fig. 1. We performed two tests checking the approximation of the peak and of the tail, respectively (*peak test* and *tail test*).

4.1.1. Setup

We briefly describe the precise setup for the statistical tests. As $\mathbb{E}_0[\tau] = \infty$ and $\mathbb{P}_0[\tau \in dt] \approx 0$ for $t \approx 0$ we compute a histogram only for $t \in [T_0, T_1]$ with $0 \leq T_0 < T_1 < \infty$ fixed. If the size of the bins is very small, we choose $T_0 > 0$ such that every bin is hit with sufficiently high probability. If a simulated path ran longer than T_1 it was stopped and thus contributes only to the tail of the corresponding histogram (which was not included into the χ^2 -test). To measure the quality of the approximations, we performed a χ^2 -test over two different sets of (equidistant) time intervals (the bins of the histogram). As a measure of approximation, we computed [21]

$$\bar{\chi}^2 = \sum_{i=0}^{N_b-1} \frac{(N_i - Np_i)^2}{Np_i}, \tag{17}$$

where N_b denotes the number of bins, N is the sample size, N_i is the number of trials that fell in bin i and p_i is the relative expected frequency of bin i . Asymptotically, $\bar{\chi}^2$ has a χ^2 -distribution with $N_b - 1$ degrees of freedom (DOFs).

For the peak (tail) test, we have chosen the bins 0.05, 0.06, ..., 1.0 (0, 1, ..., 250) and the sample size $N = 1e5$ (1e6).

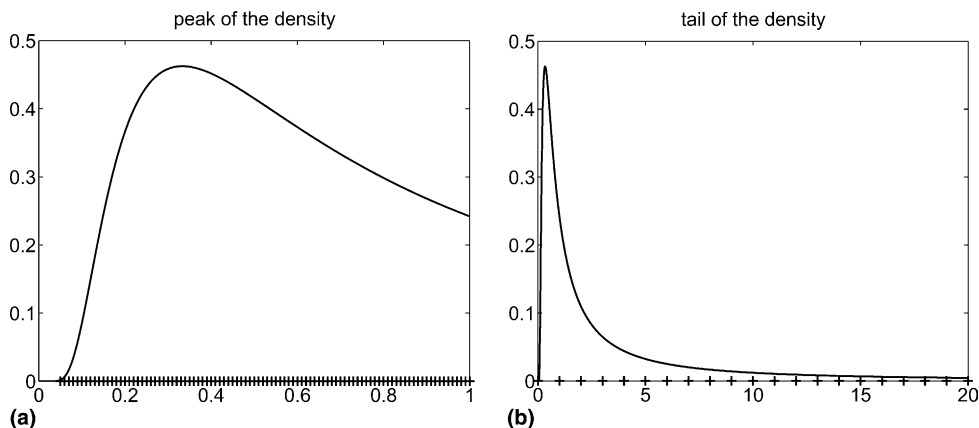


Fig. 1. Plot of the density (16) together with the bins chosen for the statistical test. Note (i) the different scalings used on the t -axis ($0 \leq t \leq 1$ on the left, whilst $0 \leq t \leq 20$ on the right), and (ii) that for the *tail test* bins up to $t = 250$ were used (see Section 4.1.1) whereas $t \leq 20$ in the plot.

4.1.2. Results

To get an impression of what can be expected with the chosen setup of the test and the random number generator used, we started by sampling N random numbers directly from the density (16). Inspired by the Box–Muller-method to generate two normally distributed i.i.d. random numbers [5, p. 13] we define for $u, v \sim \mathcal{U}$ and $b \neq 0$

$$s(u, v) = \left(\frac{b}{\sqrt{-2 \ln u \sin(2\pi v)}} \right)^2, \quad t(u, v) = \left(\frac{b}{\sqrt{-2 \ln u \cos(2\pi v)}} \right)^2 \quad (18)$$

with $0 < s, t < \infty$. It is easy to see that for $b = 1$ the random variables s, t are i.i.d. with density (16).

We obtained the results shown in Table 1 when running the two tests.

From the results in Table 1, we see that a $\bar{\chi}^2$ per DOF of the order of unity can be expected.

In order to thoroughly motivate the need for an exit test we start with the results for the Euler method without any exit test, i.e., we stop (only) if $X_{k+1} \geq 1$ and set $\tau = t_k$ in this case.

From Table 2, we see that the simple Euler scheme gives very poor results. In particular, it completely fails to resolve the peak at $t = 1/3$. In addition, the approximation of the tail is far from satisfactory. The poor resolution becomes especially apparent, if we compare the results from Table 2 with the results we obtained when applying our algorithm to the same test problem, see Table 3.

From Table 3, we see that our algorithm gives very good results: The tail of (16) is approximated perfectly, independent of the chosen step size h , and the approximation of the peak of this same density becomes better and better as h is reduced (k increases). Note that the barrier is at $b = 1$ whereas the peak of the density is at $t = 1/3$. Therefore, good approximation can only be expected when: (i) $\sqrt{h} \ll 1$ and (ii) $h \ll 1/3$ – and from Table 3 we see, that a step size as big as $h = 1/8$ already gives reasonable results.

We next compare our algorithm with other approaches. To demonstrate the superiority of our algorithm we show the ratio obtained when dividing the $\bar{\chi}^2$ -value of the alternative approaches by the corresponding value of the advocated algorithm (the larger a value the poorer the corresponding result).

We start with algorithms applying a *killing* test. These perform the test (6) if $z = X_{k+1} < 1$ and stop (if the test evaluates successfully or if $z \geq 1$) at $t_k, t_{k+1/2} = t_k + h/2$ or at t_{k+1} . Results from the peak test are in Table 4 and those from the tail test in Table 5.

From Table 4, we see that our algorithm gives much better results in the peak test. We further note that for smaller step sizes (larger k) the other algorithms show very similar (yet poor) results.

From Table 5, we see that for the approximation of the tail, stopping at t_k or at $t_k + h/2$ is satisfactory whereas stopping at t_{k+1} deteriorates results especially for large step sizes.

We next want to show that our test (see Section 2.1) already resolves the most severe problems in approximating the peak of the density (16). We therefore applied our test (Section 2.1) and stopped at $t_k + \mathcal{T}_1$ if $\mathcal{T}_1 \leq h$. In the case that $z = X_{k+1} \geq 1$ we stopped at t_k or at $t_{k+1/2} = t_k + h/2$, respectively. For comparison, we include further results when applying our test in conjunction with Mannella's approach [9] when $z \geq 1$. We limit ourselves to the results of the peak test, see Table 6.

From Table 6, we see that results are improved significantly when the test from Section 2.1 is applied, yet none of the approximations when $X_{k+1} \geq 1$ reaches the quality of the algorithm that samples \mathcal{T}_2 (as derived in Section 2.2).

Table 1
 $\bar{\chi}^2$ per DOF obtained when sampling directly (using (18)) from (16) for various sample sizes N

	$N = 1e4$	$N = 1e5$	$N = 1e6$	$N = 1e7$
Peak test	0.945	1.01	1.08	1.03
Tail test	1.10	1.12	0.897	1.00

Table 2

$\bar{\tau}^2$ per DOF obtained with the Euler method without any exit test for various step sizes $h = 1/2^k$

	$k = 0$	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$	$k = 6$
Peak test	7.62e3	5.89e3	4.15e3	2.51e3	1.29e3	5.57e2	1.60e2
Tail test	6.29e2	3.85e2	2.31e2	1.35e2	7.56e1	4.18e1	2.22e1

Table 3

$\bar{\tau}^2$ per DOF obtained with our algorithm for various step sizes $h = 1/2^k$

	$k = 0$	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$	$k = 6$
Peak test	6.38e1	9.38e1	1.99e1	2.16	1.58	1.34	1.14
Tail test	1.02	1.03	0.973	1.05	1.02	0.920	0.964

Table 4

Comparison of results of the peak test between various algorithms applying a killing test and our algorithm (for various $h = 1/2^k$)

	$k = 0$	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$	$k = 6$
$\tau = t_k$	1.81e2	1.11e2	3.90e2	2.09e3	1.32e3	5.69e2	1.70e2
$\tau = t_{k+1/2}$	3.96e2	1.54e2	4.14e2	1.85e3	1.13e3	5.43e2	1.68e2
$\tau = t_{k+1}$	6.83e2	1.84e2	4.15e2	1.79e3	1.12e3	5.38e2	1.66e2

Table 5

Comparison of results of the tail test between various algorithms applying a killing test and our algorithm (for various $h = 1/2^k$)

	$k = 0$	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$	$k = 6$
$\tau = t_k$	1.02	1.07	0.971	1.03	0.919	1.22	1.06
$\tau = t_{k+1/2}$	1.02	1.07	0.971	1.03	0.919	1.22	1.06
$\tau = t_{k+1}$	2.25e3	5.93e2	1.15e2	2.30e1	6.58	2.73	1.55

Table 6

Comparison of results of the peak test between our algorithm and various algorithms that generate \mathcal{T}_1 to test for an excursion (for various $h = 1/2^k$)

	$k = 0$	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$	$k = 6$
$\tau = t_k$ if $z \geq 1$	4.48e1	2.91e1	9.82e1	5.19e2	3.28e2	1.41e2	42.1
$\tau = t_{k+1/2}$ if $z \geq 1$	1.00e2	3.86e1	1.06e2	4.67e2	2.79e2	1.37e2	4.17e1
τ from [9] if $z \geq 1$	2.14	1.20	2.82	1.07e1	9.59	6.53	2.22

Table 7

Comparison of results of the peak test between our algorithm and exponential time stepping methods (for various $\lambda = 2^k = 1/h$)

	$k = 0$	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$	$k = 6$
$\tau = k/\lambda$	2.00e2	9.81e1	3.08e2	1.75e3	1.60e3	5.91e2	1.74e2
$\tau = (k+1)/\lambda$	3.98e2	1.34e2	3.38e2	1.59e3	1.10e3	5.26e2	1.64e2
$\tau = (k+u)/\lambda$	4.38	3.42	1.42e1	4.60e1	2.25e1	1.02e1	6.82

We conclude this section by showing the comparison with results obtained when applying the method of exponential time stepping [22,23]. This method allows a killing test similar to (6). We approximated τ by k/λ , $(k+1)/\lambda$ or by $(k+u)/\lambda$ where the random time step is exponentially distributed with parameter $\lambda > 0$ (the expected length of a time step is $1/\lambda$) and $u \sim \mathcal{U}$. The last approach is motivated by the fact that the first hitting time is independent of the length of a time step. To compare with our algorithm we set $\lambda = 1/h$. See Table 7 for the peak test and Table 8 for the tail test.

Table 8

Comparison of results of the tail test between our algorithm and exponential time stepping methods (for various $\lambda = 2^k = 1/h$)

	$k = 0$	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$	$k = 6$
$\tau = k/\lambda$	1.22e2	4.51e1	1.60e1	5.27	2.00	1.46	1.20
$\tau = (k + 1)/\lambda$	2.16e3	6.97e2	2.02e2	4.76e1	1.24e1	4.26	1.97
$\tau = (k + u)/\lambda$	1.22e2	4.51e1	1.60e1	5.27	2.00	1.46	1.20

From Table 7, we see that the exponential time stepping method has problems in approximating the peak of (16) if τ is approximated by the expected value of either the beginning or the end of the time step. We speculate that this is due to the smearing (around this expectation). Including additional randomness, however, leads to much better results.

From Table 8, we see that all the variants of the exponential time stepping method in discussion fail in approximating the tail of (16) for large step sizes (small λ). Increasing λ , however, gives results that are comparable to the ones obtained with our algorithm.

4.2. Application to the Feynman–Kac representation

We show results when applying our method to the solution of $((\cdot)' = d/dx(\cdot))$

$$\frac{\sigma(x)^2}{2}u''(x) + \mu(x)u'(x) + g(x) = 0, \quad x \in D = (a, b), \quad u(a) = u(b) = 0, \quad (19)$$

with $-\infty < a < b < \infty$, $\sigma(\cdot) > 0$, using the Feynman–Kac formulation (4) with $X_x(t)$ given by (11).

We apply the usual half-space approximation [7,24], i.e., an excursion test is applied only to the closest boundary. If $y(z)$ denotes the Euler approximation at the beginning (end) of a step, and $z \in D$, this is the boundary point which minimizes the sum of the distances to y and z . If this choice is not unique, we simply choose b . If $z \notin D$, we choose b (a) to be the closest boundary if $z \geq b$ ($z \leq a$) in order to sample \mathcal{T}_2 .

Compared methods: We always compare the results obtained with the following methods (recall that $z = X_{k+1}$):

- T:** The trivial Euler method which stops integration only if $z \notin D$ and approximates then $\tau \approx t_k$.
- K:** The method which tests for an excursion with the killing test (6). We show results from three variants which differ by the choice of the approximation for τ if an excursion is detected or if $z \notin D$:
- Kb:** $\tau \approx t_k$ (beginning of the corresponding time step).
- Km:** $\tau \approx t_k + h/2$ (middle of the corresponding time step).
- Ke:** $\tau \approx t_{k+1}$ (end of the corresponding time step).
- S:** Our algorithm, which samples \mathcal{T}_1 (see (8), (14)) to test for an excursion and \mathcal{T}_2 if $z \notin D$ (based on the inverse Gaussian distribution). We stop integration at $t_k + \mathcal{T}_1$ if an excursion is detected and at $t_k + \mathcal{T}_2$ if $z \notin D$.

Symbols in plots:

In our plots (Figs. 2–4) we use the symbols summarized in Table 9. Individual results are connected with a dotted line to guide the eye.

Parameters in simulations:

For $D = (a, b)$ we always evaluate numerically $u(x)$ at $x = 0$ (top), at $x = 0.9 \cdot b$ (middle) and at $x = 0.99 \cdot b$ (bottom). Due to the small magnitude of the errors we had to take very large sample sizes in order to observe the convergence of the systematic errors: We show plots of the relative errors versus step size h for the two sample sizes $N = 1.6e7$ (left) and $N = 6.4e7$ (right).

4.2.1. Brownian motion case

We start with the Brownian motion case, i.e., we set $\sigma(x) \equiv 1$ and $\mu(x) \equiv 0$. Then (19) reduces to Poisson’s equation (1) and $X_x(t)$ is given by (2).

We show results for two variants, namely

$$D = (-1, 1), \quad g(x) \equiv 1 \Rightarrow u(x) = \mathbb{E}_x[\tau] \quad (\text{exit problem}) \tag{20a}$$

and

$$D = \left(-\frac{\pi}{2}, \frac{\pi}{2}\right), \quad g(x) = \cos(x) \Rightarrow u(x) = 2 \cos(x). \tag{20b}$$

The results are shown in Fig. 2 for (20a) and in Fig. 3 for (20b).

From the plots in Fig. 2, we see that for the constant coefficient case the proposed method **S** (symbol ‘o’) indeed gives very accurate results. On the other hand, it is again obvious that the Euler method without any corrections (**T**, symbol ‘*’) gives very poor results (and this remark carries over to the other examples whose results are in Figs. 3 and 4).

Compared to the other methods tested, we consider method **S** overall most satisfactory: The resulting errors are always among the smallest ones obtained. Although, for example method **Km** (symbol ‘+’) shows a comparable behavior for $x = 0$ and $x = 0.9$, this method is much less accurate for $x = 0.99$ (very close to the boundary). There, method **Kb** (symbol ‘□’) gives small errors for relatively large step sizes h albeit at the price of a lower convergence order.

4.2.2. General test problem

In this section, we consider (19) with

$$\sigma(x) = 2 + \sin(x), \mu(x) = -\cos(x) \left(2 + \frac{\sin(x)}{2}\right) \quad \text{and} \quad g(x) = 2 \cos(x). \tag{21}$$

For $D = (-\pi/2, \pi/2)$ the analytical solution is $u(x) = \cos(x)$. Results are shown in Fig. 4.

From the plots in Fig. 4, we see that all the methods that apply an a posteriori test of some kind to test for a possible excursion yield very similar results. Sometimes errors resulting from method **S** are smaller than those obtained with all the other methods we tested, whilst in other tests methods **Km** or **Ke** show the best results. Note that method **S** is *never* considerably worse than any of the other methods. Therefore this method is our preferred choice for the simulation of stopped diffusions.

5. Summary

In this work, we presented an algorithm which leads itself to an efficient implementation for the simulation of stopped diffusions. Our approach used standard Euler updates and it was based on a method for the simulation of killed diffusions. Instead of simply checking if a path has reached a certain level within or at the end of a time step, we constructed a true stopping time to stop the integration. To achieve this goal, we sampled random numbers having approximatively the right distributions. In the case of diffusions with constant coefficients, these distributions are by construction exact. This allowed us to add a final Euler step of corresponding length to the simulated path and connected integrals. We think that this is the right approach for approximations in the weak sense.

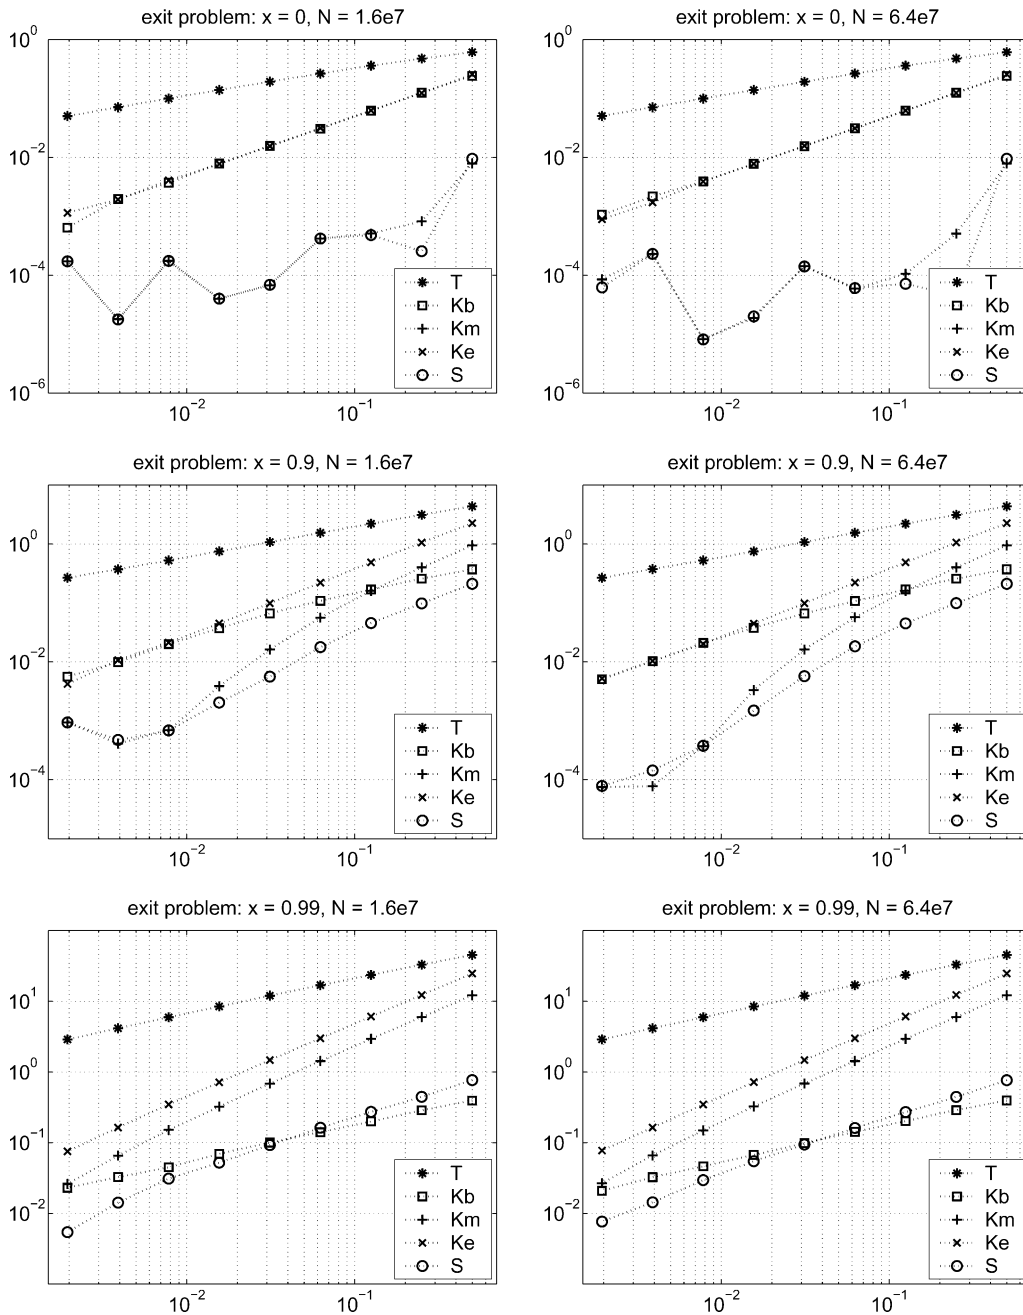


Fig. 2. Relative error vs. step size h for test problem (20a).

Our numerical tests showed evidence that the resulting distributions and thereof constructed weak approximations are of very high quality. From the results of the numerical experiments we draw the following conclusions:

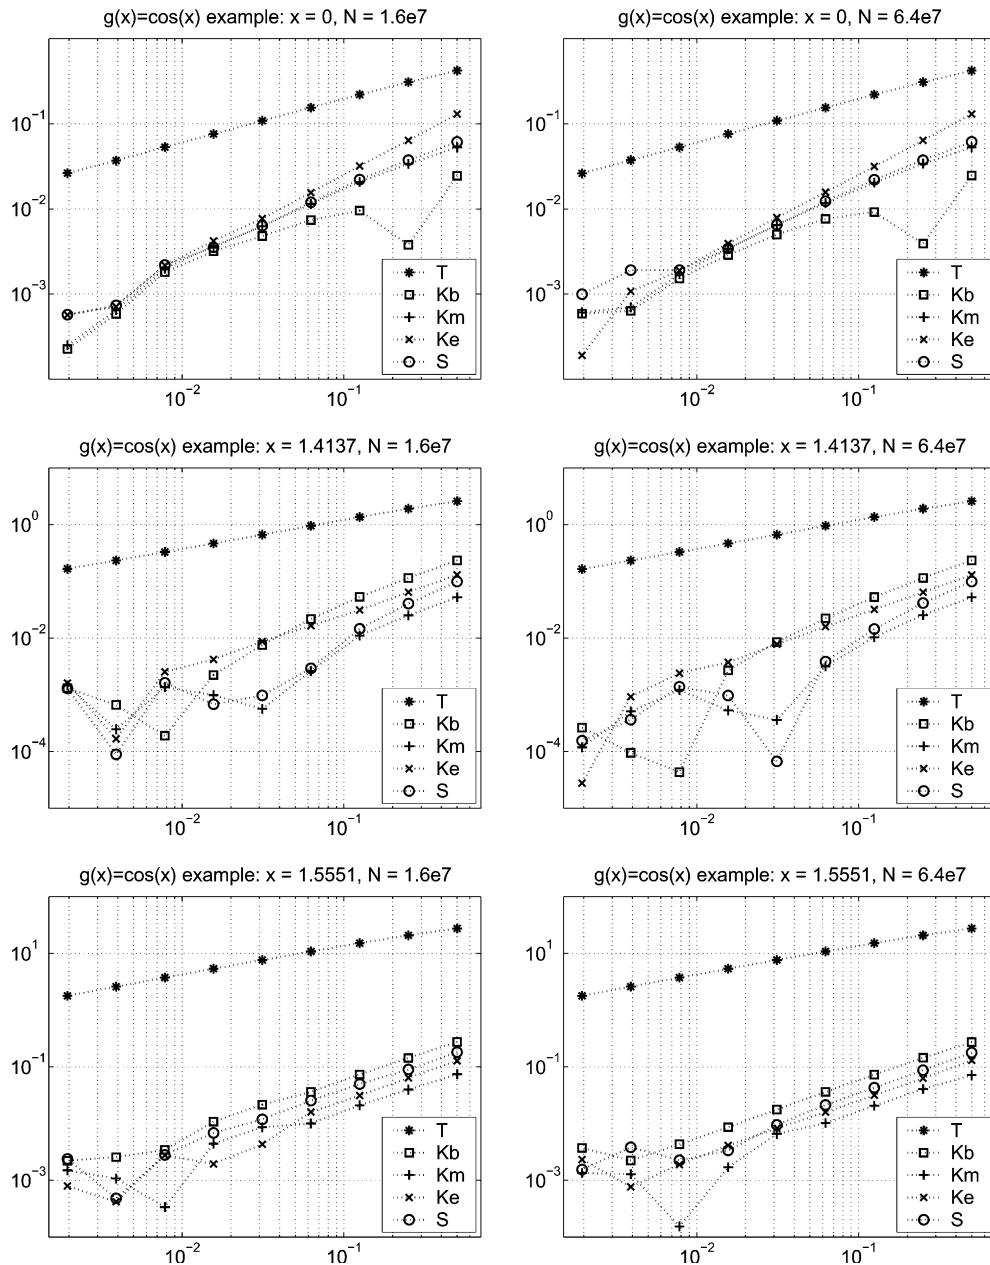


Fig. 3. Relative error vs. step size h for test problem (20b).

- If the quality of the distribution itself is important, our method S is clearly superior to all the others tested (Section 4.1).
- If the main interest lies in the approximation of (only) an expectation (Section 4.2), other methods deliver comparable results.

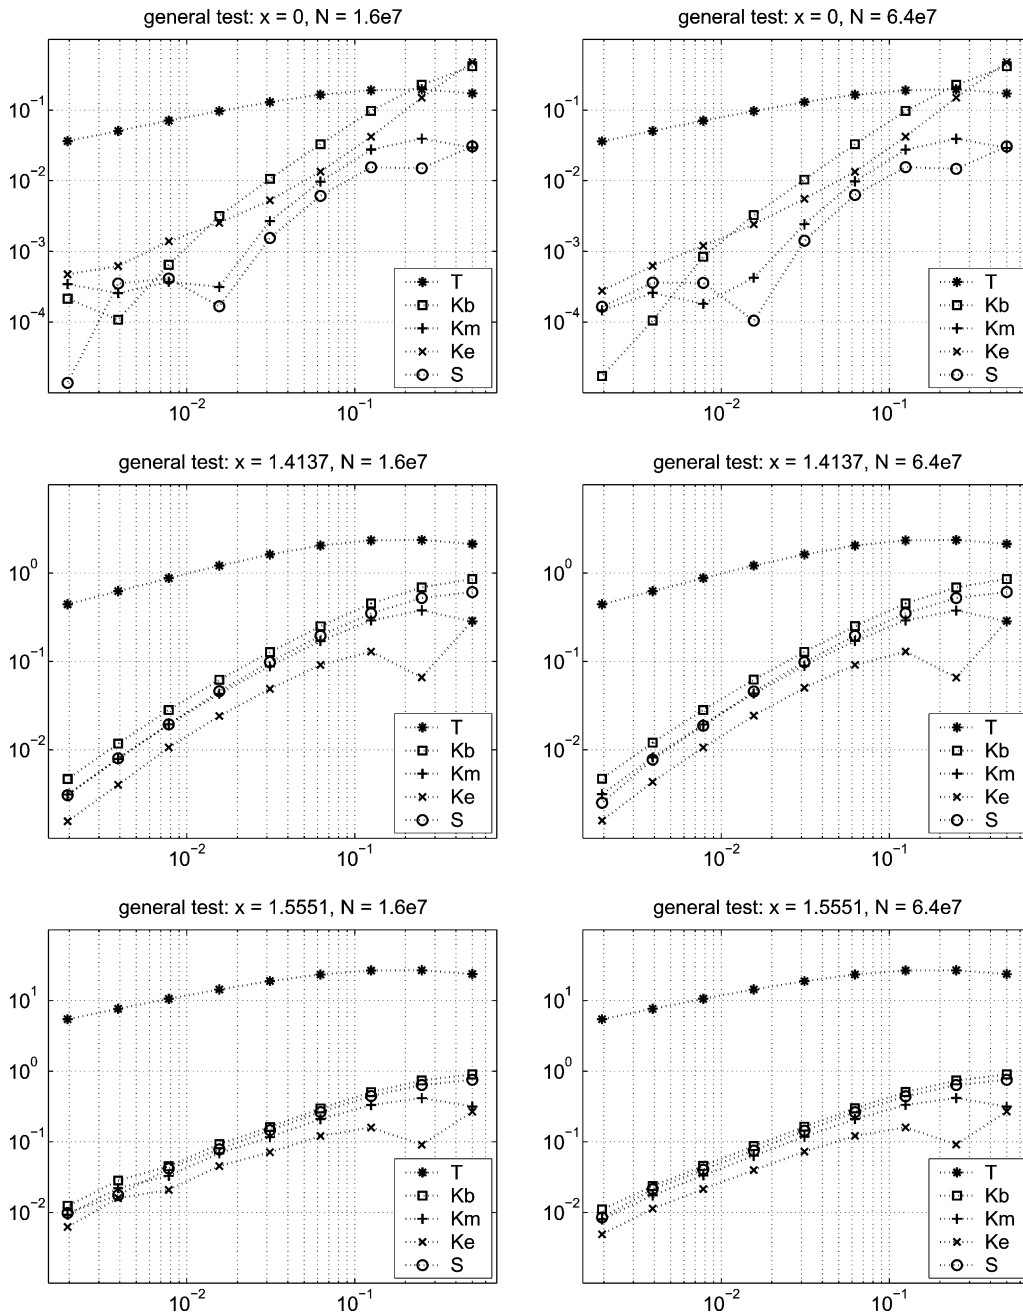


Fig. 4. Relative error vs. step size h for test problem (21).

Table 9

Symbols used in the plots in Figs. 2–4 for the different methods tested

Method	T	Kb	Km	Ke	S
Symbol	'*'	'□'	'+'	'x'	'o'

- It is, however, very important to include a test for intermediate excursions – the trivial method **T** is clearly never satisfying.
- Applying some sort of a posteriori corrections (either in form of a killing test as do methods **Kb**, **Km**, **Ke** or in form of the test presented in this work and which resulted in method **S**) remedies the main deficiencies of method **T** and yields similar results.
- The previous observation is (from a heuristical point of view) clear, as for example method **Km** (which stops integration always in the middle of a time step) adds *on average* just as much to the integral as does method **S**.
- For the approximation of the simplest possible expectation (see the *exit problem (20a)* in Section 4.2.1, where the numerical integration is exact in the weak sense and the only systematic errors are due to the stopping at first exit), though, our method **S** was again found to be most accurate (Fig. 2).

Acknowledgements

I thank Wesley P. Petersen for continuous support throughout the preparation of this paper. A special note of thanks goes to Alain-Sol Sznitman for providing me with Ref. [20] and explaining to me the details contained therein.

References

- [1] L.C.G. Rogers, D. Williams, Diffusions, Markov Processes, and Martingales, vol. 1, Cambridge Mathematical Library, Cambridge University Press, Cambridge, 2000. Foundations, Reprint of the second (1994) edition.
- [2] L.C.G. Rogers, D. Williams, Diffusions, Markov Processes, and Martingales, vol. 2, Cambridge Mathematical Library, Cambridge University Press, Cambridge, 2000. Itô calculus, Reprint of the second (1994) edition.
- [3] M. Freidlin, Functional Integration and Partial Differential Equations, Annals of Mathematics Studies, vol. 109, Princeton University Press, Princeton, NJ, 1985.
- [4] N. Madras, Lectures on Monte Carlo methods, Fields Institute Monographs, vol. 16, American Mathematical Society, Providence, RI, 2002.
- [5] P.E. Kloeden, E. Platen, Numerical Solution of Stochastic Differential Equations, Applications of Mathematics (New York), vol. 23, Springer, Berlin, 1992.
- [6] G.N. Milstein, Numerical Integration of Stochastic Differential Equations, Mathematics and its Applications, vol. 313, Kluwer Academic Publishers Group, Dordrecht, 1995 (translated and revised from the 1988 Russian original).
- [7] E. Gobet, Weak approximation of killed diffusion using Euler schemes, Stochastic Process. Appl. 87 (2) (2000) 167–197.
- [8] P. Baldi, Exact asymptotics for the probability of exit from a domain and applications to simulation, Ann. Probab. 23 (4) (1995) 1644–1670.
- [9] R. Mannella, Absorbing boundaries and optimal stopping in a stochastic differential equation, Phys. Lett. A 254 (5) (1999) 257–262.
- [10] G.N. Milstein, M.V. Tretyakov, The simplest random walks for the Dirichlet problem, Teor. Veroyatnost. i Primenen. 47 (1) (2002) 39–58.
- [11] F. Buchmann, W. Petersen, Solving Dirichlet problems numerically using the Feynman–Kac representation, BIT 43 (3) (2003) 519–540.
- [12] G.N. Milstein, Weak approximation of a diffusion process in a bounded domain, Stochastics Stochastics Rep. 62 (1–2) (1997) 147–200.
- [13] A.N. Borodin, P. Salminen, Handbook of Brownian motion – Facts and Formulae, second ed. Probability and its Applications, Birkhäuser Verlag, Basel, 2002.
- [14] E. Hausenblas, Monte Carlo simulation of killed diffusion, Monte Carlo Methods Appl. 6 (4) (2000) 263–295.
- [15] H.R. Lerche, D. Siegmund, Approximate exit probabilities for a Brownian bridge on a short time interval, applications, Adv. Appl. Probab. 21 (1) (1989) 1–19.
- [16] J.L. Folks, R.S. Chhikara, The inverse Gaussian distribution and its statistical application – a review, J. Roy. Statist. Soc. Ser. B 40 (3) (1978) 263–289 With discussion.

- [17] J.R. Michael, W.R. Schucany, R.W. Haas, Generating random variates using transformations with multiple roots, *Am. Stat.* 30 (1976) 88–90.
- [18] M.T. Giraud, L. Sacerdote, An improved technique for the simulation of first passage times for diffusion processes, *Comm. Statist. Simulation Comput.* 28 (4) (1999) 1135–1163.
- [19] P. Baldi, L. Caramellino, Asymptotics of hitting probabilities for general one-dimensional pinned diffusions, *Ann. Appl. Probab.* 12 (3) (2002) 1071–1095.
- [20] A.-S. Sznitman, A limiting result for the structure of collisions between many independent diffusions, *Probab. Theory Relat. Fields* 81 (3) (1989) 353–381.
- [21] M.G. Kendall, A. Stuart, third ed. *The Advanced Theory of Statistics*, vol. 2, Hafner Publishing Co., New York, 1973, Inference and relationship.
- [22] K.M. Jansons, G.D. Lythe, Efficient numerical solution of stochastic differential equations using exponential timestepping, *J. Statist. Phys.* 100 (5-6) (2000) 1097–1109.
- [23] K.M. Jansons, G.D. Lythe, Exponential timestepping with boundary test for stochastic differential equations, *SIAM J. Sci. Comput.* 24 (5) (2003) 1809–1822 (electronic).
- [24] E. Gobet, Euler schemes and half-space approximation for the simulation of diffusion in a domain, *ESAIM Probab. Statist.* 5 (2001) 261–297 (electronic).