J Stat Phys (2010) 140: 1130–1156
DOI 10.1007/s10955-010-0033-6

A Fast Algorithm for the First-Passage Times of Gauss-Markov Processes with Hölder Continuous Boundaries

Thibaud Taillefumier · Marcelo O. Magnasco

Received: 23 March 2009 / Accepted: 22 July 2010 / Published online: 11 August 2010

Abstract Even for simple diffusion processes, treating first-passage problems analytically proves intractable for generic barriers and existing numerical methods are inaccurate and computationally costly. Here, we present a novel numerical method that is faster and has more tightly controlled accuracy. Our algorithm is a probabilistic variant of dichotomic search for the computation of first passage times through *non-negative homogeneously Hölder continuous* boundaries by Gauss-Markov processes. These include the Ornstein-Uhlenbeck process underlying the ubiquitous "leaky integrate-and-fire" model of neuronal excitation. Our method evaluates discrete points in a sample path *exactly*, and refines this representation recursively *only* in regions where a passage is rigorously estimated to be probable (e.g. when close to the boundary).

As a result, for a given temporal accuracy in the location of the first passage time, our method is orders of magnitude faster than direct forward integration such as Euler or stochastic Runge-Kutta schemata. Moreover, our algorithm rigorously bounds the probability that such crossings are not true *first*-passage times.

Keywords First-passage times · Gauss-Markov processes · Ornstein-Uhlenbeck process

1 Introduction

The time at which the continuous sample path of a stochastic process first reaches a given boundary is a deceptively simple problem with manyfold practical applications. The first attainment of the boundary can model the onset of a chemical reaction, the triggering of a

Electronic supplementary material The online version of this article (doi:10.1007/s10955-010-0033-6) contains supplementary material, which is available to authorized users.

T. Taillefumier (✉) · M.O. Magnasco
Laboratory of Mathematical Physics, The Rockefeller University, 10021 New York, NY, USA
e-mail: ttaillefum@rockefeller.edu

M.O. Magnasco
e-mail: magnasco@rockefeller.edu

limit order for a commodity or the firing of a neuronal action potential. Despite considerable effort, studying the first-passage time problem remains among the "tough" problems in the theory of stochastic processes [18, 32, 34]. For example, probability distributions of passage times are known analytically only for the most trivial situations, such as a Wiener process first crossing a straight boundary [19]. For more practical problems, first-passage times must be computed by numerical integration, simulating many sample paths of the process until it first crosses the boundary. Such an approach traditionally carries both practical and theoretical difficulties whether focusing on the computational cost of the method or on the accuracy of the returned times [3, 16, 27].

For a class of simple stochastic processes, the Gauss-Markov processes [15, 17], which includes the practically important Ornstein-Uhlenbeck process [37], numerical integration can be performed in a completely different way. As closed-form knowledge of the conditioning formula [26] permits error-free sample path constructions at large time steps [2, 6, 11], the only issue is estimating if the process has exceeded the boundary within a given timestep. Thus, the process may be path-wise simulated at high resolution only on regions where such a resolution is warranted [10, 33]. In particular, numerical effort can be devoted to accurately reconstruct the process only when close to the boundary (for example through a dichotomic refining procedure), but not when far from the boundary [12]. Such an approach has already been advantageously used to evaluate expectation [13, 14].

The Ornstein-Uhlenbeck process is the most studied of the Gauss-Markov processes [1, 29]. It is defined as the process U solution of the linear stochastic differential equation of type

$$dU_t = \alpha U_t\, dt + \Gamma\, dW_t \quad \text{with } U_0 = u_0. \tag{1}$$

An example of first-passage problems for the Ornstein-Uhlenbeck process is the "leaky integrate-and-fire" neuron [5, 20, 30, 31]. In this model, the process U_t represents the membrane potential of a neuron: any time the voltage crosses a given threshold value, the cell fires an action potential and resets its potential to a base value. More general Gauss-Markov processes can be seen as Ornstein-Uhlenbeck processes for which the parameters α and Γ become functions of time [4, 8, 21], such as in the case of the "leaky integrate-and-fire" neuron with variable conductance.

We shall study computationally the first-passage problem of such generic Gauss-Markov processes for continuous boundaries. It is known that two factors are important: the regularity of the barrier, and the regularity of the coefficients in (1). These determine the existence and regularity of a continuous density function for first-passage time [22], and prescribes the speed of convergence of first-passage times computation [28, 38, 39]. We therefore make the two following assumptions:

Coefficient Regularity Assumption: The Gauss-Markov processes are solution of a linear stochastic equation with time-dependent *non-positive*, *bounded* function α and with time-dependent *positive*, *homogeneously Hölder continuous* function Γ.

Barrier Regularity Assumption: The barrier function is assumed to be *homogeneously Hölder continuous* and *non-negative*.

Since our algorithm is probabilistic in nature, it can return erroneous approximate time values: in such cases, it always produces approximate crossing times that are not first-passage times. However, the probability of occurrences of such errors can be tightly controlled. Our algorithm is designed not to search for first-passages in time intervals where the probability of such crossings is known to be less than a parameter value $\epsilon > 0$. As a consequence, our algorithm has the following essential property:

Error Bound: Given a real $\eta > 0$ and a recursive depth N, choosing the parameter $\ln \epsilon$ to satisfy the criterion

$$\ln \epsilon < \ln \eta - (N-1) \ln 2,$$

ensures that our algorithm returns an approximate first-passage time τ^N with a resolution of 2^{-N}, and with an error tolerance

$$\mathcal{E}(N, \epsilon) = \mathbf{P}\left(\tau^N - \tau > 2^{-N-1}\right) \leq \eta,$$

where τ is the true first-passage.

The paper is organized in two parts. In the method part, we introduce the prerequisite analytical results. We then present the implementation of the algorithm and illustrate it on simple examples. In the analysis part, we apply our method to the seminal case of a Wiener process, for which an analytical treatment is possible. We conclude by discussing how the parameters should be set to achieve computational efficiency with controlled accuracy.

2 Methods

2.1 Background

In this preliminary section, we formally define the first-passage problems and introduce some important known results about stochastic processes. This section is rather technical and can be skipped by readers familiar to the subject. The important points to retain are the discrete construction of the Gauss-Markov processes as exemplified in Fig. 1 and the analytical result (9) about the probability of a classical Wiener process crossing a linear barrier.

2.1.1 Definitions and Notations

A continuous-time real stochastic process X is a collection of real random variables X_t defined for a continuous index set of t on some abstract underlying measurable space $(\Omega, \mathcal{F})$, where Ω is the sample space and $\mathcal{F}$ denotes its associated σ-field. Assuming the index set to be $[0, \infty)$, the process X takes values in the space of functions $[0, \infty) \times \mathbb{R}$. For every realization ω in Ω, we call the outcome function $t \mapsto X_t(\omega)$ a sample path or trajectory of the process X. The natural filtration $\mathcal{F}_t \subset \mathcal{F}$ is just the intersection of the σ-algebra $\sigma(X_s)$ generated by X_s for $0 \leq s < t$ and represents the past history of the process at time t.

In the following, we assume that the process X is continuous, i.e. that it has continuous pathways $t \mapsto X_t(\omega)$ for any ω in Ω. The state space of X is the set of continuous functions on $[0, \infty)$ that are zero-valued at zero. Such a space is called the Wiener space $C_0[0, \infty)$ and is naturally provided with the σ-field $\mathcal{B}(C_0[0, \infty))$ generated by the cylinder sets $\mathcal{C}_t(A) = \{x \in C_0[0, \infty) | x(t) \in A\}$, where A is a real Borelian in $\mathcal{B}(\mathbb{R})$. Then, it is always possible to equip $(\Omega, \mathcal{F})$ with a probability measure $\mathbf{P}$, so that X induces its law measure $\mathbf{P}_{X^{-1}}$ on $\mathcal{B}(C_0(0, \infty))$ defined on the generating cylinder sets $\mathcal{C}_t(A)$ by:

$$\mathbf{P}_{X^{-1}}(\mathcal{C}_t(A)) = \mathbf{P}(\{\omega \mid X_t(\omega) \in A\}) \stackrel{\text{def}}{=} \mathbf{P}(X_t \in A).$$

Thus specified on the probability space $(\Omega, \mathcal{F}, \mathbf{P})$, the process X is entirely characterized by its law $P_{X^{-1}}$.

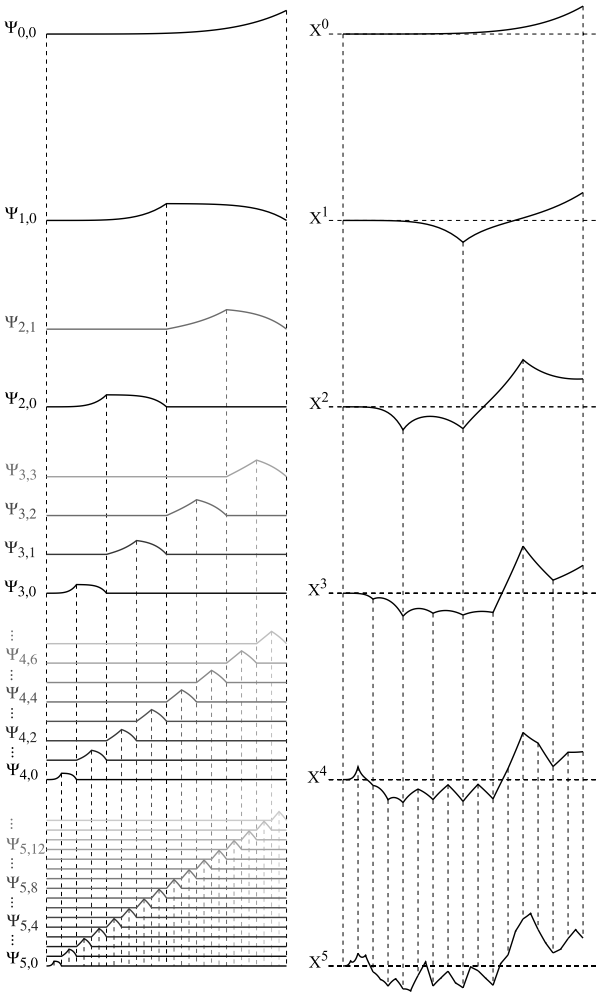

Fig. 1 In the *left column*, the elements of the basis $\Psi_{n,k}$ are represented for each rank n with $0 \le n < 6$ in the case of $\alpha(t) = -\beta/(t+\theta)$ with $\beta = 4$ and $\theta = 0.01$. In the *right column*, a sample path of the process $X^n = E(X \mid \{X_s\}$, $s \in D_n = \{k2^{-n} \mid 0 \le k \le 2^n\}$ with $D_n = \{k2^{-n} \mid 0 \le k \le 2^n\}$. Each element $\Psi_{n,k}$ has a compact support delimited by dyadic numbers in D_n and all $\Psi_{n',k}$ are zero on D_n for $n' < n$. At a fixed rank n, the time-dependence of α entails the varying shape of the basis elements $\Psi_{n,k}$. Note that this variation in shape disappears for vanishing scale when, for increasing n, the basis elements tends to the classical Schauder elements of the Wiener process

We are now in a position to formally define the Gauss-Markov processes. Given some probability space $(\Omega, \mathcal{F}, \mathbf{P})$, let X be a continuous stochastic process with natural filtration $\mathcal{F}_t$. Then X is a Gauss-Markov process if it satisfies both Gaussian and Markov properties:

1. X is a Gaussian process if, for any integers k and positive reals $t_1 < t_2 < \cdots < t_k$, the random vector $(X_{t_1}, X_{t_2}, \ldots, X_{t_k})$ has a joint Gaussian distribution.
2. X is a Markov process if, for any $s, t \ge 0$ and $A \in \mathcal{B}(\mathbb{R})$,

$$\mathbf{P}\left(X_{t+s} \in A \mid \mathcal{F}_s\right) = \mathbf{P}\left(X_{t+s} \in A \mid X_s\right),$$

which states that the conditional probability distribution of future states X_{t+s}, given the present state and all past states $\mathcal{F}_s$, depends only upon the present state X_s.

After these preliminary definitions, let us introduce the first-passage problem. For an arbitrary continuous function L, impose the conditioning $X_{t_0} = x_0 \le L(t_0)$ and consider the

random variable τ_s^L on $(\Omega, \mathcal{F}, \mathbf{P})$

$$\tau_s^L(\omega) = \inf\left\{t \geq s \mid X_t(\omega) \geq L(t)\right\}.$$

The random variable τ_s^L thus defined is a stopping time with respect to $\mathcal{F}_t$ and is called the first passage time of X for a boundary condition L. The first-passage time problem consists in evaluating the distribution of τ_s^L on $[s, +\infty)$. Furthermore, as we specifically deal with barriers that are homogeneously Hölder continuous, we recall the definition of the set H of such functions on a given interval $[0, T]$

$$H = \left\{L \in C[0,T] \mid \exists \delta > 0, \sup_{0 \leq t,s \leq T} \frac{|L(t) - L(s)|}{|t-s|^\delta} < \infty\right\},$$

where $C[0, T]$ denotes the set of continuous function on $[0, T]$.

2.1.2 Doob's Integral Representation

If X is a centered Gauss-Markov process with $X_0 = 0$ on $(\Omega, \mathcal{F}, \mathbf{P})$ with $\mathcal{F}$ containing the natural Brownian filtration, there exists a positive non-zero function g and a function f in $L^2_{\text{loc}}(\mathbb{R}^+)$ such that

$$X_t = g(t) \int_0^t f(u)\, dW_u, \tag{2}$$

where W denote a classical Wiener process on $(\Omega, \mathcal{F}, \mathbf{P})$.[1] This representation of X in terms of a stochastic integral with respect to the Wiener process is called the Doob's representation [9]. If we introduce the non-decreasing function h as

$$h(t) = \int_0^t f^2(u)\, du,$$

then, for any $t > t_0 \geq 0$, the forward conditioning formula $p(x, t \mid x_0, t_0)\, dx = \mathbf{P}(X_t \in dx \mid X_{t_0} = x_0)$ is expressed in terms of functions h and g

$$p(x, t \mid x_0, t_0) = \frac{1}{g(t)\sqrt{2\pi(h(t) - h(t_0))}} \cdot \exp\left(-\frac{\left(\frac{x}{g(t)} - \frac{x_0}{g(t_0)}\right)^2}{2(h(t) - h(t_0))}\right). \tag{3}$$

Under the condition of derivability of g, the process X is solution of the stochastic equation

$$d\left(\frac{X_t}{g(t)}\right) = f(t)\, dW_t.$$

It is desirable to write X as a solution of the equation

$$dX_t = \alpha(t) \cdot X_t\, dt + \sqrt{\Gamma(t)}\, dW_t, \tag{4}$$

[1] If $(\Omega, \mathcal{F}, \mathbf{P})$ is not large enough to contain the natural Brownian filtration, we can always define W on an appropriate extension of the original probability space.

because α is easily interpreted in term of an elastic force and Γ in term of noise intensity. The functions α and Γ need satisfying the following relations

$$\alpha(t) = \frac{g'(t)}{g(t)} \quad \text{and} \quad \sqrt{\Gamma(t)} = g(t) f(t),$$

conditions that we can equivalently express under the reciprocal form

$$g(t) = e^{\int_0^t \alpha(u)\,du} \quad \text{and} \quad h(t) = \int_0^t \Gamma(u) \cdot e^{-2\int_0^u \alpha(v)\,dv}\,du.$$

In that respect, we remark that the centered Gauss-Markov processes notably includes the centered Ornstein-Uhlenbeck process U^α with $f(t) = \sqrt{\Gamma} \cdot e^{\alpha t}$ and $g(t) = e^{-\alpha t}$. We restrain the scope of our algorithm to the first-passage problem for processes X that are solution of (4) with α non-positive. This stems from the fact that for such α, it is possible to find under a simple analytical form an upper bound to the probability of a crossing a non-negative barrier within a given interval.

2.1.3 Discrete Construction of Gauss-Markov Processes

Given some reals $t_x < t_y < t_z$, the conditioning formula for Gaussian processes together with the Markov property yield the distribution of X_{t_y} knowing $X_{t_x} = x$ and $X_{t_z} = z$ [36]. It can be shown that the corresponding probability density is the normal law $\mathcal{N}(\mu(t_y), \sigma(t_y))$, where $\mu(t_y)$ denotes the time-dependent mean

$$\mu(t_y) = \frac{g(t_y)}{g(t_x)} \cdot \frac{h(t_z) - h(t_y)}{h(t_z) - h(t_x)} \cdot x + \frac{g(t_y)}{g(t_z)} \cdot \frac{h(t_y) - h(t_x)}{h(t_z) - h(t_x)} \cdot z, \tag{5}$$

and $\sigma(t_y)$ is the time-dependent standard deviation defined by

$$\sigma(t_y)^2 = g^2(t_y) \cdot \frac{(h(t_y) - h(t_x))(h(t_z) - h(t_y))}{h(t_z) - h(t_x)}. \tag{6}$$

We emphasize that whereas the mean depends on the conditioning $X_{t_x} = x$ and $X_{t_z} = z$, the variance only exhibits time-dependence.

Let $\xi_{n,k}$ with $n \geq 0$ and $0 \leq k < 2^{n-1}$ be Gaussian random variables of law $\mathcal{N}(0, 1)$ on the probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Under assumption that g and h are continuous on $[0, 1]$, we use relations (5) and (6) in [35] to show that there exists a basis of continuous functions $\Psi_{n,k}$ for $n \geq 0$ and $0 \leq k < 2^{n-1}$ such that the random variable

$$X_t^N = \sum_{n=0}^{N} \sum_{0 \leq k < 2^{n-1}} \Psi_{n,k}(t) \cdot \xi_{n,k} \tag{7}$$

follows the same law as the conditional expectation of X_t with respect to the filtration generated by $\{X_{k/2^N}\}$ for $0 \leq k \leq 2^N$ (see Appendix B). If D_N denotes the dyadic ensemble $\{k2^{-N} | 0 \leq k \leq 2^N\}$, it just means that the process X^N follows the same law as the conditional expectation $E(X_t \mid \{X_s\}, s \in D_N)$, considered as a function of the random variables X_s with s in D_N. By construction, the functions $\Psi_{n,k}$ thus defined have support in $S_{n,k} = [k \cdot 2^{-n+1}, (k+1)2^{-n+1}]$ for $n \geq 1$. Moreover, the path-wise limit $\lim_{N \to \infty} X^N$ defines almost surely a continuous process which is an exact representation of X.

The functions $\Psi_{n,k}$ admit simple analytical expression in terms of functions g and f, as shown in Appendix B. Unfortunately the mean $\mu(t_y)$ and the variance $\sigma(t_y)^2$ as well as the basis elements $\Psi_{n,k}$ display more complicated integral formulations in terms of α and Γ. Yet, these quantities can be expressed in closed form in favorable situations as exemplified in Fig. 1.

2.1.4 Analytical Results on First-Passage Times

Interesting analytical results are available for the distribution of first-passage times in the case of simple processes derived from the classical Wiener process W.

First, we mention the situation of a time-changed Wiener process X with a constant threshold Λ. By time-changed process, we mean that X is solution of (4) for α set to zero: indeed the function h is then equal to

$$h(t) = \int_0^t \Gamma(u)\, du$$

and it is known that the process X has the same law the time-changed process $\{W_{h(t)}, \mathcal{F}_{h(t)}; 0 \le t \le 1\}$, where $W = \{W_t, \mathcal{F}_t; 0 \le t \le h^{-1}(1)\}$ is a standard Wiener process. For such processes, the strong Markov property and the reflection principle allow us to compute the joint probability of the running maximum of X between any two times t_x and t_z and of the end value of X in t_z. Let assume that $X_{t_x} = x$, if we denote

$$M_{t_x,t_z} = \max_{t_x \le t \le t_z} X_t,$$

it is known [19] that we have for any $\lambda \ge x, z$

$$\begin{aligned}&\mathbf{P}\big(X_{t_z} \in dz, M_{t_x,t_z} \in d\lambda \mid X_{t_x} = x\big)\\ &= \frac{2 \cdot (2\lambda - (x+z))}{\sqrt{2\pi (h(t_z) - h(t_x))^3}} \cdot \exp\left(-\frac{(2\lambda - (x+z))^2}{2(h(t_z) - h(t_x))}\right) d\lambda\, dz.\end{aligned}$$

As shown in Appendix C1, direct manipulations of this expression yields the probability for the process X to reach a constant threshold Λ between the times t_x and t_z

$$\mathbf{P}\big(\tau_{t_x}^{\Lambda} < t_z \mid X_{t_x} = x,\ X_{t_z} = z\big) = \exp\left(-2 \cdot \frac{(\Lambda - x)(\Lambda - z)}{h(t_z) - h(t_x)}\right). \tag{8}$$

Second, for general Gauss-Markov processes, it is possible to chose a barrier interpolating two given points (t_x, L_x) and (t_z, L_z), so that a similar result holds. Indeed, for to the case of a standard Wiener process and an affine barrier

$$L(t) = \frac{t_z - t}{t_z - t_x} \cdot L_x + \frac{t - t_x}{t_z - t_x} \cdot L_z,$$

it is possible to show that we have

$$\mathbf{P}\big(\tau_{t_x}^{L} < t_z \mid X_{t_x} = x,\ X_{t_z} = z\big) = \exp\left(-2 \cdot \frac{(L_x - x)(L_z - z)}{t_z - t_x}\right),$$

assuming suitably that $x \le L_x$ and $z \le L_z$ (see Appendix C2). From there, we can use the well known fact that the Gauss Markov process X has the same law as $\{g(t) W_{h(t)}, \mathcal{F}_{h(t)}; 0 \le$

$t \leq 1\}$, with $W = \{W_t, \mathcal{F}_t; 0 \leq t \leq h^{-1}(1)\}$ a standard Wiener process [9], to express an equivalence between the first-passage problem of W with L and the first-passage problem of X with a well-chosen barrier. More precisely, we demonstrate in Appendix C3 that, if we define the barrier to be $\mu(t)$, the mean value of X_t conditioning to $X_{t_x} = L_x$ and $X_{t_z} = L_z$, we have

$$\mathbf{P}\left(\tau_{t_x}^{\mu} < t_z \mid X_{t_x} = x,\ X_{t_z} = z\right) = \exp\left(-\frac{2\,(L_x - x)\,(L_z - z)}{g(t_z)g(t_x)(h(t_z) - h(t_x))}\right). \tag{9}$$

Unfortunately, if there are other analytical results holding for some Gauss-Markov process X [23, 24], none of them are of relevance for first-passage problems with non-specified continuous barrier. Yet, expression (8) and (9) proves of particular interest to estimate the probability for a Gauss-Markov process X to hit the barrier L within a certain time interval $[t_x, t_z]$. We use extensively these results in our proposed algorithm.

2.2 Algorithm

We now turn to our algorithm that efficiently computes the distribution of first-passage times for a general class Gauss-Markov processes X and of continuous thresholds L. It consists in implementing recursively a dichotomic search for the first crossing of simulated sample paths with the boundary, assuming the following facts:

Assumption 1 The Gauss-Markov process X under scrutiny are solutions of (4) with $\sqrt{\Gamma}$ in the set of homogeneously Hölder continuous functions and with α being a non-positive function, bounded on every compact support $[0, T]$, $T > 0$.

Assumption 2 The barrier L is a non-negative homogeneously Hölder continuous function on every compact support $[0, T]$, $T > 0$.

Bearing in mind these restrictions under which we operate, we proceed to explain the algorithm in several stages.

First, we explicit the recursive scheme for simulating Gauss-Markov sample paths. Second, we give the plain dichotomic search algorithm for first-passages. Third, we elaborate the probabilistic version of the dichotomic search by adding a probabilistic screening at each recursive step: the recursive construction of sample paths is only further if the probability of a crossing within a given time interval, is estimated greater than some small parameter ϵ. Fourth, in order to check the previous probabilistic screening, we establish an upper bound to the conditional probability that a crossing happens in a given segment, knowing the value of the sample paths at the endpoints of that segment. Fifth, we detail the base case responsible for the termination of the algorithm at the resolution limit: we simulate a first-passage in an end segment with the exact same probability as the probability that the Gauss-Markov process crosses a particular continuous function that interpolates the barrier at the endpoints. Sixth, we summarize formally the whole algorithm by giving its condensed mathematical formulation. Finally, we illustrate the algorithm on two simple examples, one of which is analytically solvable.

2.2.1 Recursive Schema for the Sample Paths

Beforehand, let us introduce the following short notations to simplify the writing of forthcoming expressions

$$l_{n,k} = (2k)\,2^{-n}, \qquad m_{n,k} = (2k+1)\,2^{-n}, \qquad r_{n,k} = 2\,(k+1)\,2^{-n}.$$

Remark that the support $S_{n,k}$ can be written as $[l_{n,k}, r_{n,k}]$, that $D_{N-1} = \{l_{N,k}, r_{N,k} \mid 0 \leq k < 2^{N-1}\}$ and that $D_N \setminus D_{N-1} = \{m_{N,k} \mid 0 \leq k < 2^{N-1}\}$. Also notice that the support $S_{n,k}$ constitute a binary tree of nested compact supports: up to the point $m_{n,k}$, we have the partition $S_{n,k} = S_{n+1,2k} \cup S_{n+1,2k+1}$ and we can identify $m_{n,k}$ with $r_{n+1,2k}$ and $l_{n+1,2k+1}$. We say that a sample path is simulated up to depth N or with resolution 2^{-N} if the value of $t \mapsto X_t(\omega)$ for a given ω have been simulated for every time t in the dyadic set D_N.

Now, assume we can generate a collection of random variables $\xi_{n,k}$ simulating independent normal laws $\mathcal{N}(0,1)$ on some probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Then expressions (5) and (6) provide us with an exact iterative scheme to simulate a sample path of X appropriately defined on $(\Omega, \mathcal{F}, \mathbf{P})$ and up to a depth N. Indeed, supposing that we have simulated the values of X_t for two consecutive dyadic points $l_{n,k}$ and $r_{n,k}$ in D_{N-1}, we can simulate the outcome of X_t at the midpoint $m_{n,k}$ in $D_N \setminus D_{N-1}$ by drawing a random variable

$$X_{m_{n,k}} = \sigma(m_{n,k}) \cdot \xi_{n,k} + \mu(m_{n,k}),$$

where $\sigma(m_{n,k})$ and $\mu(m_{n,k})$ are given by the conditioning formula at $m_{n,k}$. If we posit $X_{l_{n,k}} = x_{n,k}$ and $X_{r_{n,k}} = z_{n,k}$, direct translation of (5) and (6) yields

$$\begin{aligned}
\mu(m_{n,k}) &= \frac{g(m_{n,k})}{g(l_{n,k})} \cdot \frac{h(r_{n,k}) - h(m_{n,k})}{h(r_{n,k}) - h(l_{n,k})} \cdot x_{n,k} \\
&\quad + \frac{g(m_{n,k})}{g(r_{n,k})} \cdot \frac{h(m_{n,k}) - h(l_{n,k})}{h(r_{n,k}) - h(l_{n,k})} \cdot z_{n,k}, \\
\sigma(m_{n,k}) &= g(m_{n,k}) \sqrt{\frac{(h(m_{n,k}) - h(l_{n,k}))(h(r_{n,k}) - h(m_{n,k}))}{h(r_{n,k}) - h(l_{n,k})}}.
\end{aligned}$$

The basis step of the recurrence, i.e. the simulation of a value $z_{0,0}$ at $r_{0,0}$ knowing $x_{0,0}$ at $l_{0,0}$, is immediate by direct application of the forward conditioning formula (3).

From there, the refinement of a simulated sample path $t \mapsto X_t(\omega)$ for a given ω satisfying $X_{l_{n,k}}(\omega) = x_{n,k}$ and $X_{r_{n,k}}(\omega) = x_{n,k}$ and up to depth $N \geq n$, is implemented recursively as follows:

Procedure 1 *subdivide*$(l_{n,k}, x_{n,k}, r_{n,k}, z_{n,k})$

```
if r_{n,k} − l_{n,k} = 2^{−N} then
    basecase(l_{n,k}, x_{n,k}, r_{n,k}, z_{n,k})
else
    simulate the value y_{n,k} of the sample path at m_{n,k}
    subdivide(l_{n+1,2k}, x_{n,k}, r_{n+1,2k}, y_{n,k})
    subdivide(l_{n+1,2k+1}, y_{n,k}, r_{n+1,2k+1}, z_{n,k})
end if
```

The function *basecase* in Procedure 1 merely implements the termination of the recursion. The interest of that scheme lies in the fact that it is constructed on a binary tree of nested supports, allowing us to refine the simulation of a sample path on any given $S_{n,k}$ independently of others disjoint supports $S_{n',k'}$.

2.2.2 Recursive Dichotomic Search of First-Passage Times

In order to compute first-passage times in a support $S_{n,k}$, a straightforward algorithm consists in constructing sample paths with depth N until a drawn value occurs above the barrier L, or until the constructed path reaches the endpoint $r_{n,k}$ of $S_{n,k}$ staying below L.

As observable in Procedure 2, we just have to alter the recursive scheme in Procedure 1 to integrate the following idea: if there is a time s in $S_{n,k}$ for which $X_s(\omega) \geq L(s)$, by continuity of the sample paths, we know that a crossing has occurred before s; therefore, we have to disregard continuing the simulation of the sample path $t \mapsto X_t(\omega)$ for time t following s.

Procedure 2 *passage*$(l_{n,k}, x_{n,k}, r_{n,k}, z_{n,k})$

```
if r_{n,k} − l_{n,k} = 2^{−N} then
    return basecase(l_{n,k}, x_{n,k}, r_{n,k}, z_{n,k})
else
    simulate the value y_{n,k} of the sample path at m_{n,k}
    if y_{n,k} ≥ L(m_{n,k}) then
        return passage(l_{n,k}, x_{n,k}, m_{n,k}, y_{n,k})
    else
        if time = passage(l_{n,k}, x_{n,k}, m_{n,k}, y_{n,k}) > 0 then
            return time;
        else
            return passage(m_{n,k}, y_{n,k}, r_{n,k}, z_{n,k})
        end if
    end if
end if
```

Procedure 2 will form the backbone of our algorithm and we refer to it as a dichotomic search algorithm. For any simulated sample path $t \mapsto X_t(\omega)$, the method relies on the recursive exploration of the binary tree of dyadic segments $S_{n,k}$, effectively investigating in the prefix order every segment for which $x_{n,k}$ is below the barrier L.

However, for a given sample path, it is obviously possible that no first-passage occurs within $S_{0,0}$, the root segment of the binary tree of supports. We then need to simulate such sample paths $t \mapsto X_t(\omega)$ for time t later than $r_{0,0}$. In other words, we have to simulate values of the sample path for every time step $d_0 = r_{0,0} - l_{0,0}$ by successive application of the forward conditioning formula (3), and then initiate a recursive search during any of these time steps.

Remark 1 It is possible that the expected first-passage time diverges (for instance, in the case of a Wiener process with a constant barrier). To circumvent this predicament, we have to limit the scope of the search to a given compact segment $[0, T]$.

2.2.3 Probabilistic Screening of First-Passage Times

Denote $P_{n,k}$ the probability for the process X to cross the boundary L between $l_{n,k} = k2^{-n+1}$ and $r_{n,k} = (k+1)2^{-n+1}$ knowing $X_{l_{n,k}}$ and $X_{r_{n,k}}$, the value of X on two successive points in D_{N-1}. Thus defined, $P_{n,k}$ is a random variable on the probability space $(\Omega, \mathcal{F}, \mathbf{P})$. More

Procedure 3 *searchfirstpassage*

$r_{0,0} = 0$
$z_{0,0} = 0$
time $= 0$
while time $= 0$ **do**
  $l_{0,0} \leftarrow r_{0,0}$
  $x_{0,0} \leftarrow z_{0,0}$
  $r_{0,0} \leftarrow r_{0,0} + d_0$
  simulate the value $z_{0,0}$ of the sample path at $r_{0,0}$
  time $=$ *firstpassage*$(l_{0,0}, x_{0,0}, r_{0,0}, z_{0,0})$
end while
return time

formally, we write $P_{n,k}$ as the conditional probability

$$P_{n,k} = \mathbf{P}\left(\tau^L_{l_{n,k}} \in S_{n,k} = [l_{n,k}, r_{n,k}] \mid X_{l_{n,k}}, X_{r_{n,k}}\right).$$

Still, as X is a Markov process, $P_{n,k}$ knowing $X_{l_{n,k}} = x_{n,k}$ and $X_{r_{n,k}} = y_{n,k}$, becomes a deterministic function of the times $l_{n,k}$ and $r_{n,k}$ and the corresponding values $x_{n,k}$ and $y_{n,k}$. We will simply write the outcome of $P_{n,k}$ as $P_{n,k}(\omega) = P_{n,k}(x_{n,k}, z_{n,k})$ making the dependence on $l_{n,k}$ and $r_{n,k}$ implicit.

Assuming that we are provided with an estimation of $P_{n,k}(x_{n,k}, z_{n,k})$, we can choose to refine the simulation of $t \mapsto X_t(\omega)$ only if the probability for a crossing to happen is larger than what is admissible: typically a small positive real ϵ chosen according to the total number of simulated paths and the desired level of accuracy. Formally stated, at a given depth n, we only investigate a sample path $t \mapsto X_t(\omega)$ between two consecutive dyadic times $l_{n,k}$ and $r_{n,k}$ if the outcome of $P_{n,k}$ for the particular occurrence ω is larger than ϵ.

The idea behind this probabilistic screening is to search for first-passage times by only simulating the process on the dyadic points where the outcomes happen close enough to the boundary L. Since exact computation of $P_{n,k}(x_{n,k}, z_{n,k})$ is impossible, we need a simple analytical upper bound $B_{n,k}$ to $P_{n,k}$: we can always discard recursive searches in supports $S_{n,k}$ for sample path satisfying $P_{n,k}(\omega) \leq B_{n,k}(\omega) \leq \epsilon$.

We refer to the underlying algorithm of Procedure 4 as a probabilistic dichotomic search algorithm. Two rules of exploration of the binary tree $S_{n,k}$ are implemented: we discard the branches of the tree issued from a root segment for which $B_{n,k}(\omega) \leq \epsilon$; we only explore the branches occurring before any dyadic point t for which $X_t(\omega)$ exceeds $L(t)$. We detail in the following section how to compute an upper bound to the probability $P_{n,k}$.

2.2.4 Upper Bound to the Probability of First-Passage Time

For an homogeneously Hölder continuous non-negative boundary, we introduce the binary tree of minima of L on the compact supports $S_{n,k}$ as

$$\underline{L}_{n,k} = \inf_{t \in S_{n,k}} L(t).$$

We stress that this structure needs to be computed for the implementation of the algorithm: if the threshold function L is of simple analytical expression, the values $\underline{L}_{n,k}$ can be evaluated dynamically; otherwise the tree of minima should be evaluated numerically once and for all.

Procedure 4 *passage*$(l_{n,k}, x_{n,k}, r_{n,k}, z_{n,k})$

```
if r_{n,k} − l_{n,k} = 2^{−N} then
    return basecase(l_{n,k}, x_{n,k}, r_{n,k}, z_{n,k})
else
    compute the upper bound B_{n,k}(x_{n,k}, z_{n,k})
    if B_{n,k}(x_{n,k}, z_{n,k}) > ε then
        simulate the value y_{n,k} of the sample path at m_{n,k}
        if y_{n,k} ≥ L(m_{n,k}) then
            return passage(l_{n,k}, x_{n,k}, m_{n,k}, y_{n,k})
        else
            if time = passage(l_{n,k}, x_{n,k}, m_{n,k}, y_{n,k}) > 0 then
                return time;
            else
                return passage(m_{n,k}, y_{n,k}, r_{n,k}, z_{n,k})
            end if
        end if
    end if
end if
```

We remark in Appendix D that, as long as the function α remains non-positive on $S_{n,k}$ and that $\underline{L}_{n,k} \geq \max(x, 0)$, we have

$$\begin{aligned}&\mathbf{P}\big(\tau^{L}_{l_{n,k}} < r_{n,k} \mid X_{l_{n,k}} = x,\ X_{r_{n,k}} = z\big)\\ &\quad\leq \mathbf{P}\left(\tau'^{\underline{L}_{n,k}}_{l_{n,k}} < r_{n,k} \,\middle|\, Y_{l_{n,k}} = x,\ Y_{r_{n,k}} = \frac{g(l_{n,k})}{g(r_{n,k})} z\right),\end{aligned}$$

with $\tau'^{\underline{L}_{n,k}}_{l_{n,k}} = \inf\{t > l_{n,k} \mid Y_t \geq \underline{L}_{n,k}\}$, where Y is a scaled time-changed Wiener process defined by

$$Y_t = \frac{g(t)}{g(t_x)} W_{h_{t_x}(t)}, \qquad h_{t_x}(t) = g^2(t_x)\big(h(t) - h(t_x)\big).$$

If we denote $B_{n,k}(x, z)$ the probability of Y to reach a constant threshold $\underline{L}_{n,k}$ knowing that $Y_{l_{n,k}} = x$ and $Y_{r_{n,k}} = g(l_{n,k})z/g(r_{n,k})$, expression (8) allows us to define a random variable $B_{n,k}$ on the probability space $(\Omega, \mathcal{F}, \mathbf{P})$

$$B_{n,k} = \exp\left(-2 \cdot \frac{(\underline{L}_{n,k} - X_{l_{n,k}})(\frac{g(r_{n,k})}{g(l_{n,k})}\underline{L}_{n,k} - X_{r_{n,k}})}{g(r_{n,k})g(l_{n,k})(h(r_{n,k}) - h(l_{n,k}))}\right), \tag{10}$$

satisfying the desired upper bound condition $P_{n,k} \leq B_{n,k}$.

We emphasize that the need to compute a binary tree of minimum $\underline{L}_{n,k}$ is a crucial step to implement our algorithm. Indeed, we can be in situations where the starting time t_0 of the process is not known in advance, such as for the simulation of a train of spikes for a noisy "leaky integrate and fire" neuron. In this model, after a first-passage time (or spiking event) with a given barrier (or firing threshold), the neuron is reset to its initial condition (or resting potential), its state (or electric potential) then evolves as an Ornstein-Uhlenbeck process until it reaches again the firing threshold. It is easy to show that the evaluation of consecutive first-passage events is actually equivalent to consider a first-passage problem

with a class of time-shifted barrier for varying initial condition [30]. In such cases, our construction scheme uses a *floating* binary tree of supports $S'_{n,k}$ which depends on the time shift $t_0 = l_{0,0}$. We cannot assume it will superimpose the binary tree of minimum $\underline{L}_{n,k}$, since for computational cost we compute it once and for all on the partition $S_{n,k}$ for a given shift, say $t_0 = 0$. This complication is overcome by defining

$$B_{n,k} = \exp\left(-2 \cdot \frac{(\underline{L}'_{n,k} - X_{l_{n,k}})(\frac{g(r_{n,k})}{g(l_{n,k})}\underline{L}'_{n,k} - X_{r_{n,k}})}{g(r_{n,k})g(l_{n,k})(h(r_{n,k}) - h(l_{n,k}))}\right)$$

with $\underline{L}'_{n,k} = \min(\underline{L}_{n,k_1}, \underline{L}_{n,k_2})$, where k_1, k_2 are computed as the only two admissible integers satisfying $S'_{n,k} \cap S_{n,k_1} \neq \emptyset$ and $S'_{n,k} \cap S_{n,k_2} \neq \emptyset$ up to the endpoints. Indeed the minimum of the barrier L on given *floating* supports $S'_{n,k}$ is always less than the minimum of L on $S_{n,k_1} \cup S_{n,k_2} \supset S'_{n,k}$.

2.2.5 *Treatment of the Base Case*

If N denotes the maximum depth of exploration, a naive idea is to consider that no crossing has occurred in a limit segment $S_{N+1,k}$ of length 2^{-N}, unless the last simulated value $z_{N+1,k}$ is larger than $L(z_{N+1,k})$. As the algorithm explores the tree of nested support $S_{n,k}$ in the prefix order,[2] any such crossing in a limit segment $S_{N+1,k}$ has to be a first-pasage, whose timing can be arbitrarily set to $m_{N+1,k}$. However, this approach is very unsatisfactory because it neglects the potential occurrence of a crossing between two consecutive points in D_N which are below the barrier L. This eventuality cannot be discarded and is a major source of error in the simulation of first-passage times. To circumvent this issue, it is much more preferable to assess $P_{N+1,k}(x_{N+1,k}, z_{N+1,k})$, the probability of occurrence of a crossing in a limit segment $S_{N+1,k}$.

Since the function $t \mapsto \underline{L}_{N+1,k}(t)$ is a piecewise constant function on $S_{0,0}$, the first-passage problem for this presumably discontinuous barrier is not well posed and the upper bound $B_{N+1,k}(x_{N+1,k}, z_{N+1,k})$ yields an inconvenient estimate of $P_{N+1,k}$. For some x and z such that $x \leq L(l_{N,k})$ and $z \leq L(r_{N+1,k})$, it is actually preferable to approximate $P_{N+1,k}(x, z)$ by

$$Q_{N+1,k}(x, z) = \exp\left(-\frac{2(L(l_{N+1,k}) - x)(L(r_{N+1,k}) - z)}{g(r_{N+1,k})g(l_{N+1,k})(h(r_{N+1,k}) - h(l_{N+1,k}))}\right). \tag{11}$$

Recalling expression (9), $Q_{N+1,k}(x, z)$ just gives the probability that the process X conditioned by $X_{l_{N+1,k}} = x$ and $X_{r_{N+1,k}} = z$, crosses the barrier μ^N, which is defined such that for any t in $S_{N+1,k}$, $\mu^N(t)$ is the expected value of X_t knowing $X_{l_{N+1,k}} = L(l_{N+1,k})$ and $X_{r_{N+1,k}} = L(r_{N+1,k})$. Thus, the algorithm approximates the barrier L by a piecewise continuous function μ^N interpolating L on the dyadic numbers D_N. With such a piecewise barrier, the first-passage problem for the Wiener process is well-posed and yields consistent estimates of $P_{N+1,k}$ at the limit depth when N tends to infinity. Moreover, notice that in the limit of large N, the interpolating approximation of the barrier rapidly converges to the linear piecewise interpolation.

[2] Remember that for any explored segment $S_{n,k}$, we necessarily have $x_{n,k} < L(l_{n,k})$.

Procedure 5 *basecase*$(l_{n,k}, x_{n,k}, r_{n,k}, z_{n,k})$

if $z_{n,k} \geq L(r_{n,k})$ **then**
  return $m_{n,k}$
else
  compute the estimate $Q_{n,k}(x_{n,k}, z_{n,k})$
  draw ξ a uniformly between 0 and 1
  if $\xi \leq Q_{n,k}(x_{n,k}, z_{n,k})$ **then**
    return $m_{n,k}$
  else
    return 0
  end if
end if

2.3 Formal Definition of the Algorithm

We now recapitulate formally the recursive procedure for the dichotomic search of a first-passage within the segment $[0, 1]$. We aim at computing approximate occurrences τ^N of the true first-passage time τ, with an resolution of 2^{-N}. In that perspective, assume we are provided with two families of independent identically distributed random variables $\{\xi_{n,k}\}$, $0 \leq n < N$, $0 \leq k < 2^{n-1}$ of normal law $\mathcal{N}(0, 1)$, and $\{\upsilon_k\}$, $0 \leq k < 2^N$, of uniform law $U(0, 1)$.

Definition 1 The definition of the approximate sample path τ^N of our algorithm proceeds as follows:

1. The values of a sample path X at the endpoints of $[0, 1]$ are given by

$$\mathcal{I}_0 = \{0, 1\}, \qquad X_0^0 = 0, \qquad X_1^0 = g(1)\sqrt{h(1)} \cdot \xi_{0,0}.$$

By induction, for $1 \leq n \leq N$, we define:

(a) The set of indices

$$\mathcal{K}_n = \left\{0 \leq k < 2^{n-1} \mid B_{n,k}(X_{l_{n,k}}^{n-1}, X_{r_{n,k}}^{n-1}) > \epsilon\right\},$$

indicating the segments in which a crossing is to be investigated. The supports $S_{n,k}$, with k in $\mathcal{K}_n$, is thus inductively defined as the set of supports with dyadic endpoints in D_{n-1}, for which the corresponding upper bound $B_{n,k}$ is not small enough for the probability of a crossing to be neglected.

(b) Corresponding to the supports $S_{n,k}$, for k in $\mathcal{K}$, we define the set of dyadic times

$$\mathcal{L}_n = \{l_{n,k} \mid k \in \mathcal{K}_n\}, \qquad \mathcal{M}_n = \{m_{n,k} \mid k \in \mathcal{K}_n\}, \qquad \mathcal{R}_n = \{r_{n,k} \mid k \in \mathcal{K}_n\},$$

which consists of the left points, midpoints and right points.

(c) The values of X_t^n, for dyadic times in $\mathcal{L}_n \cup \mathcal{M}_n \cup \mathcal{R}_n$ by

$$\forall t \in \mathcal{L}_n \cup \mathcal{R}_n, \quad X_t^n = X_t^{n-1},$$
$$\forall t \in \mathcal{M}_n, \quad X_t^n = \sigma(t) \cdot \xi_{n,2^N t} + \mu(t)$$

from the values of $\{X_t^{n-1}\}_{t \in D_{n-1}}$ and the random drawings of $\xi_{n,k}$.

2. For each $1 \leq n \leq N$, the algorithm disregard any time following the occurrence of a value above the barrier L, that is, defining

$$t^n = \inf\left\{t \in \mathcal{L}_n \cup \mathcal{M}_n \cup \mathcal{R}_n \mid X_t^n \geq L(t)\right\},$$

the algorithm investigates for a first-passage time in dyadic segments $S_{n,k}$ delimited by endpoints in the set

$$\mathcal{I}_n = \left\{t \in \mathcal{L}_n \cup \mathcal{M}_n \cup \mathcal{R}_n \mid t \leq t^n\right\}.$$

In the previous definition of t^n, we observe the convention that $\inf \emptyset = 1$ to account for the possibility of a first-passage, even if all the simulated value of $\{X_t^n\}_{t \in D_n}$, are below the barrier.

3. Finally, we define formally the approximate first-passage time τ^N as

$$\tau^N = \inf\left\{t \in \mathcal{I}_N \mid Q_{N+1,2^N t}(X_t^N, X_{t+2^N t}^N) \geq \upsilon_{2^N t}\right\} + 2^{-N-1}.$$

The time τ^N is then the midpoint of the first support $S_{N,k}$, k in $\mathcal{I}_N$ on which a simulated sample path of X interpolating $\{X_t^N\}_{t \in \mathcal{I}_N}$ would cross $\mu^N(t)$, the piecewise continuous approximation of L interpolating its value on the dyadic points D_N.

2.4 Examples

We illustrate the use of the algorithm in the simple case of an Ornstein-Uhlenbeck process U with an elastic coefficient $\alpha = -1$ and a noise intensity $\Gamma = 1$.

Considering the case of a constant barrier at $\Lambda = 1$, we represent in Fig. 2 the simulation of a sample path up to a first-passage time. Remark that the sample path is simulated with increasing precision close to the barrier. The zooming operations in time regions where the path is about to cross the threshold underscore this fact. Figure 2 exemplifies a rather unfavorable situation: the first-passage occurs relatively late in time, about 4 times later than its expected value, and the sample path wanders three times in the close vicinity of the barrier. Notice that, despite this, the algorithm only needs computing 683 sample points, when the simulation of a sample path at full resolution would have required to compute more than 8×10^6 sample points.

In the particular case of a constant barrier set to $\Lambda = 0$, the probability density of the first-passage time is known analytically. If we assume the initial condition $U_0 = -1$, we actually have [1]

$$\mathbf{P}({}_U\tau_{-1}^0 \in dt) = \frac{1}{\sqrt{2\pi}} \left(\frac{1}{\sinh(t)}\right)^{\frac{3}{2}} \exp\left(-\frac{e^{-t}}{2\sinh t} + \frac{t}{2}\right) dt.$$

In Fig. 3, we compare the inferred distribution function obtained from binning the first-passages of our algorithm with this true probability density function. As apparent in logarithmic coordinates, the agreement is excellent.

3 Analysis

In the present part, we analyze the behavior of the algorithm. First, due to its analytical tractability, we review the properties of the probabilistic screening in the seminal case of a Wiener process with a constant threshold. Second, we study in detail the computational efficiency and accuracy of our algorithm in the general case set by Assumptions 1 and 2

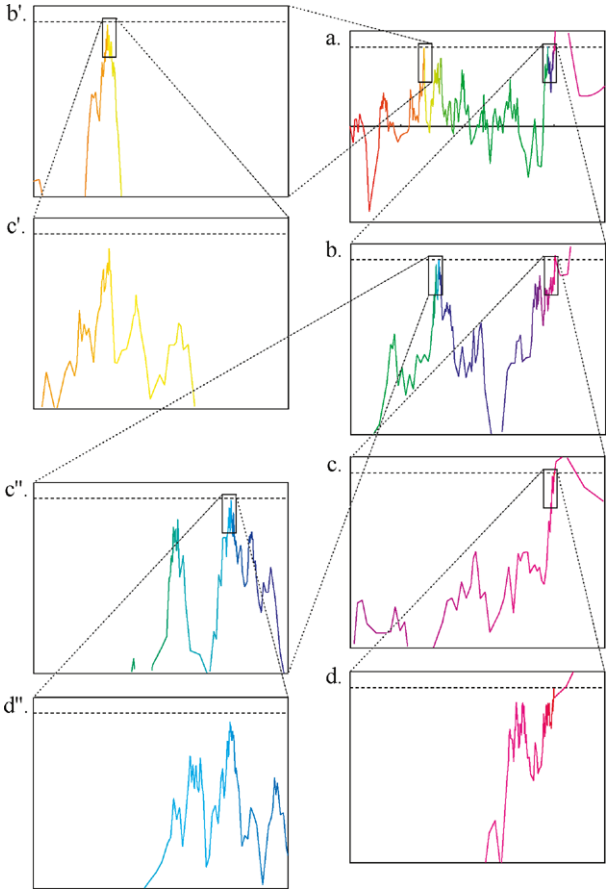

Fig. 2 (Color online) We consider the first-passage problem for an Ornstein-Uhlenbeck process U with an elastic coefficient $\alpha = -1$ and a noise intensity $\Gamma = 1$, the barrier is constant $\Lambda = 1$ and the initial condition is $U_0 = 0$. We represent at different scales a realization $U_t(\omega)$ for which the algorithm returns a first-time passage $\tau(\omega) = 8.00469684$ with a resolution $\delta t = 2^{-21} = 5 \times 10^{-7}$ after 5 recursive calls. The whole sample path is represented in (**a**.) and a series of zooms is carried out around $\tau(\omega)$ in (**b**.), (**c**.), (**d**.). The dilation coefficients are set according to the scale invariance of a Wiener process with the time scale being expanded by 20 and the distance scale by $4\sqrt{5}$ during each dilation. The series (**b**′.), (**c**′.) and (**c**″.), (**d**″.) zoom on regions where the sample path gets close to the barrier. The simulation of the sample path has required 683 subdivisions and illustrates an unfavorable situation since the expected number of divisions is approximatively 284 for this particular setting

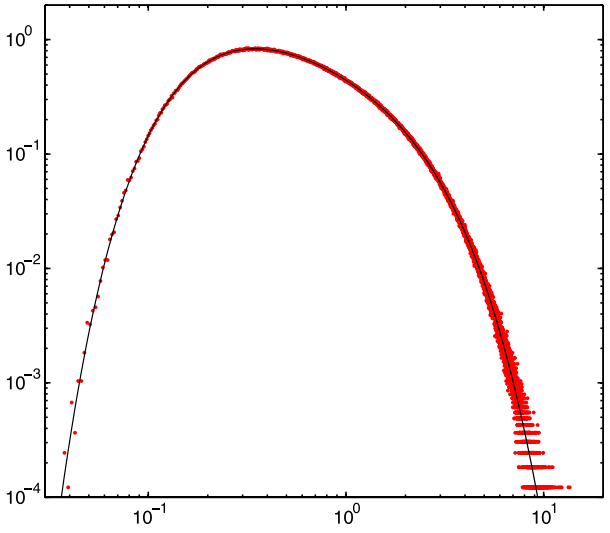

Fig. 3 (Color online) We consider the first-passage problem for an Ornstein-Uhlenbeck process U with an elastic coefficient $\alpha = -1$ and a noise intensity $\Gamma = 1$, the barrier is constant $\Lambda = 0$ and the initial condition is $U_0 = -1$. The *black thin line* represents the exact distribution of first-passage times, which is known analytically in the specific case of a barrier set to zero. The *red dots* are obtained by populating an histogram of 10^3 bins with 10^6 first-passage times simulations

3.1 Statistical Screening for the Wiener Process with Constant Threshold

The statistical screening of the algorithm consists in checking whether a recursion call terminates before the limit resolution, in which case it finds $B_{n,k} \leq \epsilon$ if the support $S_{n,k}$ is under scrutiny. For each search, the algorithm constructs a binary tree of visited supports $S_{n,k}$ constructed by recursive statistical screening. It generally proves intractable to study the statistics of such binary trees because the random variables $B_{n,k}$ exhibits non-trivials dependences across and within levels of the binary tree.

Yet, the simple situation of a standard Wiener process W and a constant barrier Λ is more propitious to characterize the statistical screening. Indeed, in this specific case, the upper bound $B_{n,k}$ becomes rigorously equal to the probability of crossing $P_{n,k}$, which simplifies the interpretation of the results.

In the following, we study the statistical properties of such screening rules. First, we explicit the conditional law of $B_{n+1,2k}$ and $B_{n+1,2k+1}$ knowing $X_{l_{n,k}} = x_{n,k}$ and $X_{r_{n,k}} = z_{n,k}$. Second, in light of this result, we describe when halting because $B_{n,k} \leq \epsilon$, leads to an erroneous approximation of a first-passage time is given. Third, we use the statistical property of the conditional probability to characterize the asymptotic behavior of the bound $B_{n,k}$ at vanishing scale.

3.1.1 Statistics of the Upper Bounds

Suppose we explore a sample path $t \mapsto W_t(\omega)$ for which $W(l_{n,k}) = x_{n,k} < \Lambda$ and $W(r_{n,k}) = z_{n,k} < \Lambda$ and posit the corresponding probability $P_{n,k}(x_{n,k}, z_{n,k}) = B_{n,k}(x_{n,k}, z_{n,k}) = p_{n,k}$. The rule of exploration of the binary tree is the following. If $p_{n,k} > \epsilon$, the algorithm proceeds to a refinement of the path by simulating a sample value $y_{n,k}$ at $m_{n,k}$. Incidentally, such an operation gives rise to the evaluation of the two probabilities $B_{n+1,2k}(x_{n,k}, y_{n,k}) = p_{n+1,2k}$ and $B_{n+1,2k+1}(y_{n,k}, z_{n,k}) = p_{n+1,2k+1}$. Next, the algorithm continues the exploration of $S_{n,k}$ by investigating the left support $S_{n+1,2k}$ if $p_{n+1,2k} > \epsilon$ or by investigating the right support if $p_{n+1,2k} \leq \epsilon$ and $p_{n+1,2k+1} > \epsilon$.

Here, we express the conditional law of the variables $B_{n+1,2k}$ and $B_{n+1,2k+1}$ knowing $X_{l_{n,k}} = x_{n,k}$ and $X_{r_{n,k}} = z_{n,k}$. Beforehand, in order to underline the dependence on $B_{n,k}$, we reformulate the previous conditioning under the equivalent form $B_{n,k} = p_{n,k}$ and $\Delta_{n,k} = \delta_{n,k}$, with

$$\Delta_{n,k} = \frac{\Lambda - X_{l_{n,k}}}{\Lambda - X_{r_{n,k}}} \quad \text{and} \quad \delta_{n,k} = \frac{\Lambda - x_{l_{n,k}}}{\Lambda - x_{r_{n,k}}}.$$

The result is the following:

Property 1 *Assuming $B_{n,k} = p_{n,k}$ and $\Delta_{n,k} = \delta_{n,k}$, the conditional variables $\{B_{n+1,2k} \mid B_{n,k}, \Delta_{n,k}\}$ and $\{B_{n+1,2k+1} \mid B_{n,k}, \Delta_{n,k}\}$ follow respectively the lognormal laws*

$$\{B_{n+1,2k} \mid B_{n,k}, \Delta_{n,k}\} \sim \log\mathcal{N}\Big((1+\delta_{n,k})\ln p_{n,k}, \sqrt{-2\ln p_{n,k}}\Big), \tag{12}$$

$$\{B_{n+1,2k+1} \mid B_{n,k}, \Delta_{n,k}\} \sim \log\mathcal{N}\Big((1+1/\delta_{n,k})\ln p_{n,k}, \sqrt{-2\ln p_{n,k}}\Big). \tag{13}$$

Proof Definition (11) allow us to write the random variables $B_{n+1,2k}$ and $B_{n+1,2k+1}$ knowing $W_{l_{n,k}} = x_{n,k}$ and $W_{r_{n,k}} = z_{n,k}$ as

$$B_{n+1,2k} = e^{-\frac{(\Lambda - x_{n,k})(\Lambda - W_{m_{n,k}})}{\Gamma \cdot 2^{-(n+1)}}} \quad \text{and} \quad B_{n+1,2k+1} = e^{-\frac{(\Lambda - z_{n,k})(\Lambda - W_{m_{n,k}})}{\Gamma \cdot 2^{-(n+1)}}}.$$

The natural scale invariance of the problem suggests to introduce the random variables $U_{n,k} = (\Lambda - W_{l_{n,k}})/\sqrt{\Gamma \cdot 2^{-n}}$ and $V_{n,k} = (\Lambda - W_{r_{n,k}})/\sqrt{\Gamma \cdot 2^{-n}}$, as well as their hyperbolic counterparts $\Pi_{n,k} = U_{n,k} \cdot V_{n,k}$ and $\Delta_{n,k} = U_{n,k}/V_{n,k}$. The conditioning $W_{l_{n,k}} = x_{n,k}$ and $W_{r_{n,k}} = y_{n,k}$ is equivalent to $\Pi_{n,k} = -\ln p_{n,k}$ and $\Delta_{n,k} = \delta_{n,k}$ with $p_{n,k} = B_{n,k}(x_{n,k}, z_{n,k})$ and $\delta_{n,k} = (\Lambda - x_{n,k})/(\Lambda - z_{n,k})$. As exponentials of Gaussian random variables, $B_{n+1,2k}$ and $B_{n+1,2k+1}$ knowing $W_{l_{n,k}} = x_{n,k}$ and $W_{r_{n,k}} = z_{n,k}$ follow log-normal distributions. Assuming $B_{n,k} = p_{n,k}$ and $\Delta_{n,k} = \delta_{n,k}$, the usual conditioning formulas (5) and (6) specify the conditional law of $B_{n+1,2k}$ and $B_{n+1,2k+1}$ as (12) and (13). □

We underline that the outcomes $p_{n+1,2k}$ and $p_{n+1,2k+1}$ can be greater than one: it just means that $W_{m_{n,k}}$ happens above the threshold in which case the interpretation of $p_{n+1,2k}$ and $p_{n+1,2k+1}$ in terms of probability does not hold.

Due to the skewness of log-normal distributions, the "central trend" of the conditional laws of $B_{n+1,2k}$ and $B_{n+1,2k+1}$ is best outlined by their medians

$$\mathrm{Med}\big(B_{n+1,2k} \mid B_{n,k} = p_{n,k}, \Delta_{n,k} = \delta_{n,k}\big) = p_{n,k}^{1+\delta_{n,k}},$$
$$\mathrm{Med}\big(B_{n+1,2k+1} \mid B_{n,k} = p_{n,k}, \Delta_{n,k} = \delta_{n,k}\big) = p_{n,k}^{1+1/\delta_{n,k}}.$$

The average absolute deviations to the median quantify the statistical dispersion of the log-normal distributions around their central values. Actually the median is the central point which minimizes the mean absolute deviation and we show in Appendix E1 that

$$\mathrm{Dev}\big(B_{n+1,2k} \mid B_{n,k} = p_{n,k}, \Delta_{n,k} = \delta_{n,k}\big) = p_{n,k}^{\delta_{n,k}} \cdot \mathrm{erf}\left(\sqrt{-\ln p_{n,k}}\right),$$
$$\mathrm{Dev}\big(B_{n+1,2k+1} \mid B_{n,k} = p_{n,k}, \Delta_{n,k} = \delta_{n,k}\big) = p_{n,k}^{1/\delta_{n,k}} \cdot \mathrm{erf}\left(\sqrt{-\ln p_{n,k}}\right),$$

where erf denotes the error function.

From there, notice the important following facts:

- The medians are always smaller than the conditioning value $p_{n,k}$, showing that the probabilities $p_{n+1,2k}$ and $p_{n+1,2k+1}$ of finding a crossing after refinement in $S_{n+1,2k}$ and $S_{n+1,2k+1}$ decrease *on average*.
- The deviation expressions reveal that the log-normal distributions increasingly concentrate their statistics around their medians for smaller $p_{n,k}$.

It means that, when refining the scale of sample paths, the outcomes of $B_{n,k}$ tends to accumulate on ever smaller values and the proportion of sample paths satisfying $p_{n,k} \geq \epsilon$ diminishes accordingly.

3.1.2 Omission Due to Large Increments

The probabilistic screening constitutes a source of error since it overlooks to investigate regions for which the probability of a crossing is less than ϵ. Here, we specify these deficient situations with the help of the analytical results available for the Wiener process with a constant barrier. First, remark that the quantity $\delta_{n,k}$ is the ratio of the distance to the barrier at time $l_{n,k}$ with the distance to the barrier at time $r_{n,k}$. For any given $p_{n,k}$, if $x_{n,k}$ is closer to the barrier than $z_{n,k}$, it is obvious that a crossing is more likely to happen in the left child support $S_{n+1,2k}$ than in the right child support $S_{n+1,2k+1}$. The expressions of the medians and absolute deviations quantify this effect: they describe how large values of $\delta_{n,k}$ cause the

statistics of $B_{n+1,2k}$ to be distributed on smaller values and the corresponding opposite effect for the statistics of $B_{n+1,2k+1}$.

In the deficient situation for which the algorithm neglects a first-passage time, the sample path is very close to the barrier at one of the endpoints but the overall probability of a crossing in the support $S_{n,k}$ remains inferior to ϵ, meaning that the other endpoint is relatively far from the barrier. In other words, it means that the sample path exhibits a large variation between $l_{n,k}$ and $r_{n,k}$ and that $\delta_{n,k}$ is either negligible or very large. Consequently, even though $p_{n,k} \leq \epsilon$ is very small, the value of $\delta_{n,k}$ can be extremal enough so that the absolute deviation in one of the child supports is almost one: the eventuality of a crossing in one of the child supports after refinement cannot be neglected.

3.1.3 Speed of Decay of the Upper Bounds

Along its way to a putative first-passage time, the exploration of a branch of the tree stops as soon as $B_{n,k} \leq \epsilon$. Bearing in mind that the parameter ϵ is responsible for the accumulation of errors, we require it to be set to an extremely small value. At the same time, this requirement is detrimental to the computational cost of our method since lowering ϵ increases the probability to refine sample paths at high resolution. In view of this, for the algorithm to be efficient, we hope the following: Given a nested sequence of support S_{n,k_n}, the sequence B_{n,k_n} decreases fast enough for it to be smaller than ϵ before the resolution limit case.

Only a few results can be established analytically about the speed of decay of B_{n,k_n} without making specific assumptions about the sample path. In particular, using the law of increments of the Wiener process, we prove in Appendix E2 that:

Property 2 *Knowing $B_{n,k} = p_{n,k}$, the medians of the conditional variables $\{B_{n+1,2k} \mid B_{n,k}\}$ and $\{B_{n+1,2k+1} \mid B_{n,k}\}$ are equal and satisfy*

$$\mathrm{Med}\big(B_{n+1,2k} \mid B_{n,k} = p_{n,k}\big) = \mathrm{Med}\big(B_{n+1,2k+1} \mid B_{n,k} = p_{n,k}\big) = p_{n,k}^2.$$

This means that, irrespective of the particular shape of sample paths, the decay of the bounds is doubly exponential at any depth *on average*. This result is reminiscent of the scale invariance of the Wiener process around the constant threshold. Unfortunately, such a simple result cannot be extended to the absolute deviation.

However, we can strengthen the previous result under additional assumption about the shape of the sample path. More precisely, if we assume that the sample path remain strictly below Λ on a given support $S_{p,q}$, we can see that:

Property 3 *If $\max_{t\in S_{p,q}} W_t < \Lambda$, for all $n > p$, $q2^{n-p} \leq k < (q+1)2^{n-p}$, we have $B_{n,k} < {C_{p,q}}^{2^n}$ where $C_{p,q}$ is defined as*

$$C_{p,q} = \exp\left(-\frac{(\Lambda - \max_{t\in S_{p,q}} W_t)^2}{\Gamma}\right) < 1.$$

The previous results directly stems from the expression of the upper bound B_{n,k_n}. The expression of the constant $C_{p,q}$ is determined by the distance of the maximum of the sample path to the threshold. Thus, any nested sequence of bounds B_{n,k_n} vanishes at least doubly exponentially in supports $S_{p,q} \supset S_{n,k_n}$ for which no crossing happens. Moreover, the further from the barrier, the fastest is the doubly exponential decay, which underlines the computational advantage of the algorithm.

3.2 Properties of the Algorithm

We show that the essential features of the probabilistic screening similarly hold in the general case of a Gauss-Markov process X with a continuous barrier L satisfying Assumptions 1 and 2. In particular, we verify the fast decay of the upper bounds $B_{n,k}$ for finer scale. Such a property results in two forms of advantage for computational accuracy and efficiency of our algorithm:

Computational Efficiency: while exploring sample paths at finer scale, the condition $B_{n,k} \leq \epsilon$ is satisfied after a few steps of the recursive exploration, saving us the computational cost of simulating the path at finer resolution.

Computational Accuracy: The limit probability ϵ of $B_{n,k}$ for which we neglect to continue the search for a passage, can be set all the smaller as $B_{n,k}$ vanishes fast for increasing n, thus lowering the overall probability to overlook a first-passage.

The present section is organized as follows. First, we establish an upper bound to the probability of returning an erroneous approximate first-passage time. Second, we justify the fast decay of the upper bounds $B_{n,k}$ used in the algorithm and we give a criterion to measure the algorithm efficacy. Finally, we explain the strategy to set the values of the algorithm parameters.

3.2.1 Algorithmic Accuracy

In the erroneous case, the algorithm always returns a crossing time that is not the true first-passage time. This happens when, while exploring sample paths, the algorithm dismisses regions where a first-passage actually occurs against the odds. As a result, the algorithm delays the first-passage time.
It is possible to naively estimate an upper bound to the probability of occurrence of such errors when the algorithm returns a putative first-passage time.

Property 4 *Given a parameter value $\epsilon > 0$ and a time resolution of 2^{-N}, the probability of error*

$$\mathcal{E}(N, \epsilon) = \mathbf{P}(\tau^N > \tau + 2^{-N-1}),$$

i.e. the probability that a simulated first-passage in $[0, 1]$ *does not approximate a true first-passage, satisfies*

$$\mathcal{E}(N, \epsilon) \leq \epsilon 2^{N-1}.$$

Proof False timings possibly occur when the algorithm halts the search for a first-passage in a support $S_{n,k}$ because the simulated values $x_{n,k}$ and $z_{n,k}$ are such that $B_{n,k}(x_{n,k}, z_{n,k}) \leq \epsilon$. For every such halting situations, the algorithm possibly neglects the occurrence of a crossing, the probability of such an error is then

$$\begin{aligned}&\mathbf{P}\big(\exists t \in S_{N,k}, X_t \geq L(t) \mid B_{n,k} \leq \epsilon\big)\\ &\quad = \mathrm{E}\left(P_{n,k}(X_{l_{n,k}}, X_{r_{n,k}}) \mid B_{n,k}(X_{l_{n,k}}, X_{r_{n,k}}) \leq \epsilon\right).\end{aligned}$$

By definition of $B_{n,k}$ as an upper bound to the probability of a crossing $P_{n,k}$ within $S_{n,k}$, the probability that the algorithm disregards a crossing is clearly dominated above by ϵ, i.e.

$$\mathbf{P}\big(\exists t \in S_{N,k}, X_t \geq L(t) \mid B_{n,k} \leq \epsilon\big) \leq \epsilon.$$

If we set the limit depth to be N, there are at most 2^{N-1} halts during each call of the recursive search procedure. Indeed the algorithm explores the binary tree of supports $S_{n,k}$ for $n \geq 1$, such that a given segment $S_{n,k}$ admits children supports if $B_{n,k} \leq \epsilon$. Such a tree has a depth of at most N levels and a support $S_{n,k}$ is a leaf if $n = N+1$ or $B_{n,k} < \epsilon$. There are at most 2^{N-1} leaves for such a binary tree structure.

This corresponds to the worst-case scenario for which the algorithm explores the sample path up to depth N on the whole segments and find $B_{N,k} \leq \epsilon$ for every k. The we have

$$\begin{aligned}
\mathcal{E}(N,\epsilon) &= \mathbf{P}\big(\exists k, 0 \leq k < 2^{N-1} \mid \exists t \in S_{N,k}, X_t \geq L(t), B_{N,k} \leq \epsilon\big) \\
&\leq \sum_{k=0}^{2^{N-1}-1} \mathbf{P}\big(\exists t \in S_{N,k}, X_t \geq L(t), B_{N,k} \leq \epsilon\big) \\
&\leq 2^{N-1} \sup_{0 \leq k < 2^{N-1}} \mathbf{P}\big(\exists t \in S_{N,k}, X_t \geq L(t), B_{N,k} \leq \epsilon\big) \\
&\leq \epsilon 2^{N-1}.
\end{aligned}$$

In other words, if we look for first-passage times with a resolution of 2^{-N}, the probability of an error $\mathcal{E}(N,\epsilon)$ per recursive call is inferior to $\epsilon 2^{N-1}$. □

3.2.2 Algorithmic Efficiency

Under Assumptions 1 and 2, we ideally wish to established an upper bound to the complexity of the algorithm. Since the complexity increase linearly with the number of dichotomy operations, we see that the worse case scenario consists of the situation when every sample points needs to be simulated, which yields the complexity of the classic Runge-Kutta method. Unfortunately, there appears to be no direct way to establish analytically the complexity of the algorithm. Nevertheless, because of its obvious dichotomic structure, our algorithm substantially outperforms Euler and Runge-Kutta.

We then prefer to measure the computational efficiency of the algorithm indirectly. The exploration of a branch of the tree of supports $S_{n,k}$ stops as soon as $B_{n,k} \leq \epsilon$. The value of $n \leq N$ at such a halt defines the local depth of exploration for times in the support $S_{n,k}$.

Definition 2 For all s in $[0,1]$, the local depth of exploration d_s is the discrete random variable

$$d_s = \sup\{n \leq N \mid s \in S_{n,k}, B_{n,k}(X_{l_{n,k}}, X_{r_{n,k}}) > \epsilon\}.$$

The smaller the typical depth of exploration, the fewer sample points are simulated in regions where a sample path occurs strictly below the barrier, which is desirable for computational efficiency.

We thus want to estimate the local depth of exploration d_s around s to measure the algorithmic efficiency. This is determined by the speed at which $B_{n,k}$ concentrates its distribution on vanishing values for finer resolution. In that respect, bearing in mind the case of a standard Wiener process and a constant threshold, we now show that the doubly exponential speed of decay still holds in the general case.

To be more precise, assume that for some s, the sample path occurs strictly below the continuous barrier $X_s = y < L(s)$, so that we can find a neighboring of s for which the sample points are strictly below L. There exists a unique sequence of indices k_n such

that $0 \leq k_n \leq 2^{n-1}$ and s satisfy $k_n 2^{-n+1} \leq s < (k_n+1)2^{-n+1}$. It defines a sequence of nested supports S_{n,k_n} such that we have $\bigcap_n S_{n,k_n} = \{s\}$ and by continuity of L, we have $\lim_{n\to\infty} \underline{L}_{n,k_n} = L(s)$. In view of this, if the increasing sequence $\underline{L}_{n,k_n}$ is strictly above y and if the noise intensity function Γ is continuous, we show in Appendix F that:

Property 5 *If L and $\sqrt{\Gamma}$ are respectively homogeneously δ_L- and δ_Γ-Hölder continuous, then conditionally to $X_s = y < L(s)$, for every ω in Ω, we have the following asymptotic equivalent when n tends to infinity: if $\min(\delta_L, \delta_\Gamma) < 1/2$, we have*

$$\ln B_{n,k_n}(\omega) = -2^{n+1}\frac{(L(s)-y)^2}{\Gamma(s)} + O\left(2^{n(1-\min(\delta_L,\delta_\Gamma))}\right),$$

and otherwise for every $0 < \delta < 1/2$, we have

$$\ln B_{n,k_n}(\omega) = -2^{n+1}\frac{(L(s)-y)^2}{\Gamma(s)} + O\left(2^{n(1-\delta)}\right).$$

Moreover, if P_{n,k_n} denotes the conditional probability of a crossing, we have $\ln B_{n,k_n} \sim \ln P_{n,k_n}$ when n tends to infinity.

We stress that we only account for the certainty to reach an asymptotic regime without information about the time when this regime is reached. The speed at which such a regime is attained is set by the constant bounds implicitly present in the Landau notations of Proposition 5. Introducing the (ϵ,δ)-modulus $C_f(\epsilon,\delta)$ of a δ-Hölder function f as

$$C_f(\epsilon,\delta) = \sup_{|t-s|<\epsilon} \frac{|f(t)-f(s)|}{|t-s|^\delta},$$

Appendix F shows that the previous constants are directly set by the value of the moduli $C_L(2^{-N},\delta_L)$, $C_\Gamma(2^{-N},\delta_\Gamma)$ and $C_{X(\omega)}(2^{-N},\delta')$, $\delta' > 1/2$, for every ω in Ω. Large values of these moduli delay the onset of the decaying behavior. The modulus $C_X(2^{-N},\delta')$, $\delta' > 1/2$ plays a specific part here since it clearly appears as random variable. However, the Lévy's modulus of continuity theorem ensures that $\lim_{N\to\infty} C_X(2^{-N},\delta')$ is bounded by 1.

Keeping in mind the previous limitations, the asymptotic behavior in Property 5 is similar to the case of the Wiener process and the upper bound approximation becomes exact at vanishing scale. It suggests to approximate the local depth of exploration d_s by the following quantity:

Definition 3 For all s in $[0,1]$, we define

$$N(\epsilon, y) = \log_2(-\ln\epsilon) + \log_2\left(\frac{\Gamma(s)}{(L(s)-y)^2}\right)$$

as an approximation of the local depth of exploration d_s.

This quantity results from the sum of two contributions: one from a statistical term expressing the stringency of the statistical screening and a geometrical term stemming from the interplay of the distance to the barrier and the noise intensity. Notice that the dependence on ϵ through an iterated logarithm indicates that the statistics of d_s should vary very weakly with the parameter ϵ which can be set very small. A straightforward criterion to measure

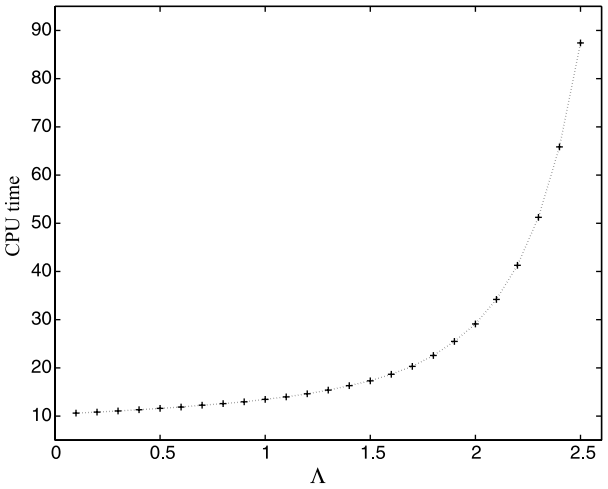

Fig. 4 Computational cost of our algorithm for the first-passages of an Ornstein-Uhlenbeck process U (with elastic coefficient $\alpha = -1$ and noise intensity $\Gamma = 1$) through a constant boundary $\Lambda = 0.1, 0.2, \ldots, 2.5$ and initial condition $U_0 = 0$. The termination of the recursion is set by the limit resolution $\delta t = 2^{-21} = 5 \times 10^{-7}$, and by the tolerance for a false first-passage $\mathcal{E} \leq 10^{-10}$. The computational cost blows up when the threshold exceeds the persistence length of the Ornstein-Uhlenbeck process $2/|\alpha| = 2$. CPU time is quoted in microseconds per first-passage computed for a single core in a 2.66 GHz Xeon processor

the algorithm computational advantage, is then to compare the typical value $N(\epsilon) = N(\epsilon, 0)$ with the total number of recursions allowed by the resolution. We will assess such a criterion in the following section, after describing how to choose the parameters values. Beforehand, we illustrate the computational cost of our algorithm numerically in Fig. 4.

3.2.3 Choice of the Parameters

Here, we explain the strategy to set the value of the various parameters intervening in the implementation in a segment $[0, T]$, where T is a positive integer. Focusing on reliability, we request that the probability of an erroneous result $\mathcal{E}_T$ is inferior to a fixed parameter $\eta > 0$. Since we simulate first-passage times to describe a full-statistics, we require η to be very small, typically of order 10^{-10}. It is straightforward to see that $\mathcal{E}_T$ as a function of the parameter ϵ and the limit depth N, is inferior to $T\mathcal{E}(N, \epsilon)$. It is then clearly enough to chose ϵ satisfying

$$\ln \epsilon \leq \ln \eta - (N-1)\ln 2 - \ln T,$$

and we consequently set $\epsilon(N, \eta, T) = \ln \eta - (N-1)\ln 2 - \ln T$. With such $\epsilon(N, \eta, T)$, the algorithm necessarily returns an erroneous first-passage with probability less than $\mathcal{E}_T \leq \eta$. Concretely, if $\eta = 10^{-10}$, let us set the limit depth to $N = 21$ so that the accuracy is $2^{-21} \simeq 5 \times 10^{-7}$. Taking a time window of total length $T = 100$, we need to chose $-\ln \epsilon(N, \eta)$ close to 40 to satisfy $\mathcal{E}_T \leq \eta$. Thus setting the parameters ensures that each returned time is a valid first-passage with probability superior to $1 - \eta$.

The question is then to inquire wether such values of ϵ actually saves us a considerable computational cost, using our previously stated criterion. The answer to that question depends on the particular nature of the first-passage problem, but a prototypical answer can

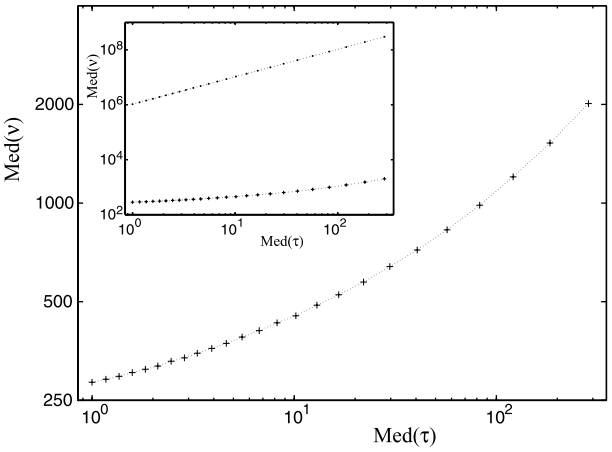

Fig. 5 Parametric plot of the median number of Gaussian draws and the median of the first-passage time for the same conditions as in Fig. 4 with $\Lambda = 0.1, 0.2, \ldots, 2.5$. Notice the logarithmic scale of both axes. As the boundary is moved away from zero, the mean first-passage time increases rapidly. *Inset*: We represent the number of calls to the Gaussian random generators for our method (+) and the classical Euler method (•) on the same log-log plot. In an Euler scheme, the number of random numbers drawn is equal to the first passage time divided by the time resolution δt. Given our time accuracy, the number of draws is many orders of magnitude smaller with our algorithm. Moreover, it is strongly sublinear in the length elapsed when increasing the time accuracy

be given in the situation of constant α and Γ for an Ornstein-Uhlenbeck process. Under the constraint that the limit resolution satisfies $2^{-N} \ll 1/|\alpha|$, if the local height of the barrier and the noise intensity are of same order, the main contribution to the local depth of exploration d_s in (3) is determined by the screening term $\log_2(-\ln \epsilon(N, \eta))$. As an example, if $\eta = 10^{-10}$ and $\alpha = -1$ for a resolution of 2^{-21}, the typical depth of exploration is $d = \log_2(-\ln \epsilon(N, \eta)) \simeq 6$. This has to be compared with the depth of the recursion necessary to simulate exhaustively the sample path up to the limit scale, i.e. $N = 21$. Clearly, the typical recursive exploration of a time region halts notably before the limit depth and, because of the iterated logarithm, this behavior is very slowly varying with the parameters ϵ or η.

We illustrate this approach for the first-passage time problem of the Ornstein-Ulhenbeck U process with $\alpha = -1$ and $\Gamma = 1$ for a constant boundary of value $\Lambda > 0$ and an initial condition $U_0 = 0$. The computational efficiency of the algorithm is essentially determined by the number of calls to the Gaussian random generator per simulated first-passage. For the Euler method, the number of Gaussian draws is just the number of time steps necessary to reach a first-passage and is therefore linear with the typical first-passage time. Equivalently stated, this number is linearly increasing with the inverse of the limit time step $1/\delta t = 2^N$ as opposed to our method which is strongly sub-linear with the inverse of the time resolution 2^N. Given the parameters of the simulation, Fig. 5 demonstrates that the performance of the algorithm is four to five orders of magnitude faster than an Euler implementation with a comparable temporal resolution, while keeping a strong guarantee that the passage determined is indeed the first one.

4 Conclusion

We have presented a probabilistic version of the classical dichotomic search algorithm, tailored to computing first-passage times of a Gauss-Markov process through continuous

boundaries. Instead of simulating all sample points up to a given resolution, our method recursively refines sample paths only in time regions $S_{n,k}$ where the conditional probability of a crossing $P_{n,k}$ cannot be neglected. This approach relies on the possibility to compute an upper-bound to $P_{n,k}$ that has a simple analytical solution. If the elastic coefficient is non-positive and if the barrier is non-negative, it is obtained by locally substituting the first-passage problem for the Gauss-Markov process and the varying barrier, with a simpler first-passage problem for a time-changed Wiener process and a constant threshold. Then, the analytically-known probability $B_{n,k}$ of crossing for the latter problem rigorously provides us with an upper bound to $P_{n,k}$. Since the contribution of the elastic force becomes negligible at vanishing scale,[3] and since the time-varying barrier can be faithfully approximated by its infinimum on intervals of vanishing length,[4] the identification of the preceding first-passage problems becomes exact in the asymptotic limit of large n.

While maintaining tight control on the probability of error, our algorithm achieves far greater computational efficiency than Euler and Runge-Kutta stochastic schemata. Traditional stochastic numerical methods simulate sample paths at a given resolution until the barrier is crossed, which overlooks the possibility that a crossing happens between two consecutive sample points. At the resolution limit case, our method solves this predicament by exactly simulating first-passages for a piecewise continuous functions that interpolates the barrier sample points. When it halts before the resolution limit case, our method can still produce crossing times that are erroneous first-passage approximations. Yet, by definition of our probabilistic screening, it is straightforward to compute an upper bound to the overall probability of error, which guarantees controlled accuracy. At the same time, the fast decay of the upper bound ensures that the halting condition is rapidly attained when sample paths are strictly below the barrier. Therefore, even if very fine resolution is enforced, our algorithm saves us considerable computational time by representing sample paths at poor resolution when far from the barrier, while rigorously controlling the error that introduces.

Before discussing potential applications, we underscore that our method should not be confused with an adaptive time-step method. In the latter, the time-step is changed heuristically according to an error budget, in an attempt to devote less numerical effort to pieces of the solution which are better behaved. Our algorithm does not have such an error budget because the reconstruction of the sample path is exact. The method inherits its recursive, adaptive nature from the dichotomic search algorithm and not from an adaptive time-step control.

We expect our algorithm to be advantageously applied whenever it is desirable to compute first-passages without simulating the entire sample paths in detail. This is particularly relevant to stochastic modeling of both neuronal firing in biology [5, 20, 30, 31] and certain derivatives in finance [7, 25, 40]. We hope this relatively modest setting is a springboard for generalizations to multidimensional processes passing through surfaces. Indeed, the recursive scheme that our algorithm uses as its backbone, can be generalized to Gauss-Markov processes of arbitrary dimensions. Moreover, similar ideas have been applied successfully to estimate the rate of convergence of the continuous Euler scheme for monitored barrier options in the multidimensional case [13, 14]. Finally, our approach was motivated by the need to produce first-passage times even when the barrier is δ-Hölder continuous with $\delta < 1/2$ [22]. In such situations, there are no guarantees that the first-passage time admits a density

[3] The contribution of the elastic force is of order $\alpha 2^{-n}$ at depth n to be compared with Γ.

[4] The Hölder continuity of the barrier prescribes that there exists $\delta, C > 0$ such that for all $n > 0$, we have $\sup_k(\sup_{t\in S_{n,k}} L(t) - \underline{L}_{n,k}) < C2^{-\delta n}$.

probability, rather the probability measure of such first-passages is strongly singular (as we shall soon communicate).

Acknowledgements We thank John Agapiou and Jonathan Touboul for their insightful comments on the manuscript. We are also grateful to the anonymous reviewers for their careful reading and their invaluable remarks.

References

1. Alili, L., Patie, P., Pedersen, J.L.: Representations of the first hitting time density of an Ornstein-Uhlenbeck process. Stoch. Models **21**(4), 967–980 (2005)
2. Bao, J., Abe, Y., Zhuo, Y.: An integral algorithm for numerical integration of one-dimensional additive colored noise problems. J. Stat. Phys. **90**(3), 1037–1045 (1998). doi:10.1023/A:1023201725795
3. Bouleau, N., Lépingle, D.: Numerical methods for stochastic processes. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. Wiley, New York (1994)
4. Burkitt, A.: A review of the integrate-and-fire neuron model: Ii. inhomogeneous synaptic input and network properties. Biol. Cybern. **95**(2), 97–112 (2006). doi:10.1007/s00422-006-0082-8
5. Burkitt, A.N.: A review of the integrate-and-fire neuron model: I. homogeneous synaptic input. Biol. Cybern. **95**(1), 1–19 (2006). doi:10.1007/s00422-006-0068-6
6. Cecchi, G.A., Magnasco, M.O.: Negative resistance and rectification in Brownian transport. Phys. Rev. Lett. **76**(11), 1968–1971 (1996). doi:10.1103/PhysRevLett.76.1968
7. DiCesare, J., Mcleish, D.: Simulation of jump diffusions and the pricing of options. Insur., Math. Econ. **43**(3), 316–326 (2008). doi:10.1016/j.insmatheco.2008.06.001. URL http://www.sciencedirect.com/science/article/B6V8N-4SSG4PR-1/2/333e62e7c38787453b5bfc5392891dd4
8. Ditlevsen, S., Lansky, P.: Estimation of the input parameters in the Ornstein-Uhlenbeck neuronal model. Phys. Rev. E **71**(1), 011,907 (2005). doi:10.1103/PhysRevE.71.011907
9. Doob, J.L.: Heuristic approach to the Kolmogorov-Smirnov theorems. Ann. Math. Stat. **20**, 393–403 (1949)
10. Gaines, J.G., Lyons, T.J.: Variable step size control in the numerical solution of stochastic differential equations. SIAM J. Appl. Math. **57**(5), 1455–1484 (1997). doi:10.1137/S0036139995286515
11. Gillespie, D.T.: Exact numerical simulation of the Ornstein-Uhlenbeck process and its integral. Phys. Rev. E **54**(2), 2084–2091 (1996). doi:10.1103/PhysRevE.54.2084
12. Giraudo, M.T., Sacerdote, L., Zucca, C.: A Monte Carlo method for the simulation of first passage times of diffusion processes. Methodol. Comput. Appl. Probab. **3**(2), 215–231 (2001)
13. Gobet, E.: Weak approximation of killed diffusion using Euler schemes. Stoch. Process. Appl. **87**(2), 167–197 (2000). doi:10.1016/S0304-4149(99)00109-X. URL http://www.sciencedirect.com/science/article/B6V1B-405KF32-1/2/3d79c4e26426d724ed2ea670b406b3a6
14. Gobet, E., Menozzi, S.: Exact approximation rate of killed hypoelliptic diffusions using the discrete Euler scheme. Stoch. Process. Appl. **112**(2), 201–223 (2004) doi:10.1016/j.spa.2004.03.002. URL http://www.sciencedirect.com/science/article/B6V1B-4C4VXT3-1/2/a1783ee30dc66ded7af8f604e017bc36
15. Hida, T.: Canonical representations of Gaussian processes and their applications. Mem. Coll. Sci. Univ. Kyoto Ser. A. Math. **33**, 109–155 (1960/1961)
16. Honeycutt, R.L.: Stochastic Runge-Kutta algorithms. i. white noise. Phys. Rev. A **45**(2), 600–603 (1992). doi:10.1103/PhysRevA.45.600
17. Kac, M., Siegert, A.J.F.: An explicit representation of a stationary Gaussian process. Ann. Math. Stat. **18**, 438–442 (1947)
18. van Kampen, N.G.: Stochastic processes in physics and chemistry. Lecture Notes in Mathematics, vol. 888. North-Holland, Amsterdam (1981)
19. Karatzas, I., Shreve, S.E.: Brownian motion and stochastic calculus, second edn. Graduate Texts in Mathematics, vol. 113. Springer, New York (1991)
20. Lansky, P., Ditlevsen, S.: A review of the methods for signal estimation in stochastic diffusion leaky integrate-and-fire neuronal models. Biol. Cybern. **99**(4–5), 253–262 (2008). doi:10.1007/s00422-008-0237-x
21. Lánský, P.L., Lánská, V.: Diffusion approximation of the neuronal model with synaptic reversal potentials. Biol. Cybern. **56**(1), 19–26 (1987). doi:10.1007/BF00333064
22. Lehmann, A.: Smoothness of first passage time distributions and a new integral equation for the first passage time density of continuous Markov processes. Adv. Appl. Probab. **34**(4), 869–887 (2002)
23. Lo, C.F.: Exact solutions of the Fokker-Planck equations with moving boundaries. Ann. Phys. **319**(2), 326–332 (2005)

24. Lo, C.F., Hui, C.H.: Computing the first passage time density of a time-dependent Ornstein-Uhlenbeck process to a moving boundary. Appl. Math. Lett. **19**(12), 1399–1405 (2006)
25. Metwally, S.A.K., Atiya, A.F.: Using Brownian bridge for fast simulation of jump-diffusion processes and barrier options. J. Deriv. **10**, 43–54 (2002)
26. Mota-Furtado, F., O'Mahony, P.F.: Exact propagator for generalized Ornstein-Uhlenbeck processes. Phys. Rev. E **75**(4), 041102 (2007). doi:10.1103/PhysRevE.75.041102. URL http://link.aps.org/abstract/PRE/v75/e041102
27. Platen, E.: An introduction to numerical methods for stochastic differential equations. Research Paper Series 6, Quantitative Finance Research Centre, University of Technology, Sydney (1999). URL http://ideas.repec.org/p/uts/rpaper/6.html
28. Pötzelberger, K., Wang, L.: Boundary crossing probability for Brownian motion. J. Appl. Probab. **38**(1), 152–164 (2001)
29. Renault, O., Scaillet, O., Leblanc, B.: A correction note on the first passage time of an Ornstein-Uhlenbeck process to a boundary. Finance Stoch. **4**(1), 109–111 (2000). URL http://ideas.repec.org/a/spr/finsto/v4y2000i1p109-111.html
30. Ricciardi, L.M., Sacerdote, L.: The Ornstein-Uhlenbeck process as a model for neuronal activity. Biol. Cybern. **35**(1), 1–9 (1979). doi:10.1007/BF01845839
31. Ricciardi, L.M., Sato, S.: First-passage-time density and moments of the Ornstein-Uhlenbeck process. J. Appl. Probab. **25**(1), 43–57 (1988)
32. Risken, H.: The Fokker-Planck equation, second edn. Springer Series in Synergetics, vol. 18. Springer, Berlin (1989). Methods of solution and applications
33. Römisch, W., Winkler, R.: Stepsize control for mean-square numerical methods for stochastic differential equations with small noise. SIAM J. Sci. Comput. **28**(2), 604–625 (2006). doi:10.1137/030601429
34. Siegert, A.J.F.: On the first passage time probability problem. Phys. Rev. (2) **81**, 617–623 (1951)
35. Taillefumier, T.: A discrete construction for Gaussian Markov processes. ArXiv e-prints (2008)
36. Taillefumier, T., Magnasco M.: A Haar-like construction for the Ornstein Uhlenbeck process. J. Stat. Phys. **132**(2), 397–415 (2008)
37. Uhlenbeck, G.E., Ornstein, L.S.: On the theory of the Brownian motion. Phys. Rev. **36**(5), 823–841 (1930). doi:10.1103/PhysRev.36.823
38. Wang, L., Pötzelberger, K.: Boundary crossing probability for Brownian motion and general boundaries. J. Appl. Probab. **34**(1), 54–65 (1997)
39. Wang, L., Pötzelberger, K.: Crossing probabilities for diffusion processes with piecewise continuous boundaries. Methodol. Comput. Appl. Probab. **9**(1), 21–40 (2007)
40. Zhang, D., Melnik, R.V.N.: First passage time for multivariate jump-diffusion processes in finance and other areas of applications. Appl. Stoch. Models Bus. Ind. **25**(5), 565–582 (2009). doi:10.1002/asmb.v25:5