

Unbiased ‘walk-on-spheres’ Monte Carlo methods for the fractional Laplacian

ANDREAS E. KYPRIANOU

Department of Mathematical Sciences, University of Bath, Claverton Down, Bath, BA2 7AY, UK
a.kyprianou@bath.ac.uk

ANA OSOJNIK

Mathematical Institute, University of Oxford, Woodstock Road, Oxford OX2 6GG, UK
ana.osojnik@maths.ox.ac.uk

AND

TONY SHARDLOW*

Department of Mathematical Sciences, University of Bath, Claverton Down, Bath, BA2 7AY, UK

*Corresponding author: t.shardlow@bath.ac.uk

[Received on 10 September 2016; revised on 24 June 2017]

We consider Monte Carlo methods for simulating solutions to the analogue of the Dirichlet boundary value problem in which the Laplacian is replaced by the fractional Laplacian and boundary conditions are replaced by conditions on the exterior of the domain. Specifically, we consider the analogue of the so-called ‘walk-on-spheres’ algorithm. In the diffusive setting, this entails sampling the path of Brownian motion as it uniformly exits a sequence of spheres maximally inscribed in the domain. As this algorithm would otherwise never end, it is truncated when the ‘walk-on-spheres’ comes within $\varepsilon > 0$ of the boundary. In the setting of the fractional Laplacian, the role of Brownian motion is replaced by an isotropic α -stable process with $\alpha \in (0, 2)$. A significant difference to the Brownian setting is that the stable processes will exit spheres by a jump rather than hitting their boundary. This difference ensures that disconnected domains may be considered and that, unlike the diffusive setting, the algorithm ends after an almost surely finite number of steps.

Keywords: Walk-on-spheres; fractional Laplacian; Lévy processes; Monte Carlo methods

1. Introduction

We start by recalling the classical Dirichlet problem in d -dimensions and re-examining a, now, classical Monte Carlo algorithm that is used to numerically simulate its solution. Suppose that D is a domain in \mathbb{R}^d , $d \geq 2$, with sufficiently smooth boundary. We are interested in finding $u: D \rightarrow \mathbb{R}$ such that

$$\begin{aligned} \Delta u(x) &= 0, & x \in D, \\ u(x) &= g(x), & x \in \partial D, \end{aligned} \tag{1.1}$$

where g is a given continuous function on the boundary. Feynman–Kac representation tells us that, for example, if $u \in C^2(\bar{D})$ is a solution to (1.1), then

$$u(x) = \mathbb{E}_x[g(W_{\tau_D})], \quad x \in D, \tag{1.2}$$

where $\tau_D := \inf\{t > 0 : W_t \notin D\}$ and $W := (W_t, t \geq 0)$ is the standard d -dimensional Brownian motion with probabilities $(\mathbb{P}_x, x \in \mathbb{R}^d)$.

The representation (1.2) suggests that solutions to (1.1) can be generated numerically via straightforward Monte Carlo simulations of the path of W until first exit from D . That is to say, if $(W_t^i, t \leq \tau_D^i)$, $i \in \mathbb{N}$ are *i.i.d.* copies of $(W_t, t \leq \tau_D)$ issued from $x \in D$, then, by the strong law of large numbers,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n g(W_{\tau_D^i}^i) = u(x), \quad \text{almost surely (a.s.).} \quad (1.3)$$

For practical purposes, because it is impossible to take the limit, one truncates the series of estimates for large n , and the central limit theorem gives $\mathcal{O}(1/n)$ upper bounds on the variance of the n -term sum, which serves as a numerical error estimate.

Although forming the fundamental basis of most Monte Carlo methods for diffusive Dirichlet-type problems, (1.3) is an inefficient numerical approach. Least of all, this is because the Monte Carlo simulation of $u(x)$ is independent for each $x \in D$. Moreover, it is unclear how exactly to simulate the path of a Brownian motion on its first exit from D , that is to say, the quantity W_{τ_D} . This is because of the fractal properties of Brownian motion, making its path difficult to simulate. This introduces additional numerical errors over and above that of Monte Carlo simulation.

A method proposed by Muller (1956), for the case that D is convex, subsamples special points along the path of the Brownian motion to the boundary of the domain D . The method does not require a complete simulation of its path and takes advantage of the distributional symmetry of the Brownian motion. To describe the so-called ‘walk-on-spheres’, we need to first introduce some notation. We may thus set $\rho_0 = x$ for $x \in D$ and define r_1 to be the radius of the largest sphere inscribed in D that is centred at x . This sphere we will call $S_1 = \{y \in \mathbb{R}^d : |y - \rho_0| = r_1\}$. To avoid special cases, we, henceforth, assume that the surface area of $S_1 \cap \partial D$ is zero (this excludes, for example, the case that $x = 0$ and D is a sphere centred at the origin).

Now set $\rho_1 \in D$ to be a point uniformly distributed on S_1 and note that, given the assumption in the previous sentence, $\mathbb{P}_x(\rho_1 \in \partial D) = 0$. Construct the remainder of the sequence $(\rho_n, n \geq 1)$ inductively. Given ρ_{n-1} , we define the radius, r_n , of the largest sphere inscribed in D that is centred at ρ_{n-1} . Calling this sphere S_n , we have that $S_n = \{y \in \mathbb{R}^d : |y - \rho_{n-1}| = r_n\}$. We now select ρ_n to be a point that is uniformly positioned on S_n . Once again, we note that if $\rho_{n-1} \in D$ almost surely, then the uniform distribution of both ρ_{n-1} and ρ_n ensures that $\mathbb{P}(\rho_n \in \partial D) = 0$. Consequently, the sequence ρ_n continues for all $n \geq 1$. In the case that ρ_n approaches the boundary, the sequence of spheres S_n become arbitrarily small in size.

Thanks to the strong Markov property and the stationary and independent increments of Brownian motion, it is straightforward to prove the following result.

LEMMA 1.1 Fix $x \in D$ and define $\rho'_1 = W_{\tau_{S'_1}}$, where $\tau_{S'_1} = \inf\{t > 0 : W_t \in S'_1\}$ and S'_1 is the largest sphere, centred at x , inscribed in D . For $n \geq 2$, given $\rho'_{n-1} \in D$, let $\rho'_n = W_{\tau_{S'_n}}$, where $\tau_{S'_n} = \inf\{t > 0 : W_t \in S'_n\}$ and S'_n is the largest sphere, centred at ρ'_{n-1} . Then, the sequences $(\rho_n, n \geq 0)$ and $(\rho'_n, n \geq 0)$ have the same law.

As an immediate consequence, $\lim_{n \rightarrow \infty} \rho_n$ almost surely exists and, moreover, it is equal in distribution to W_{τ_D} . The sequence $\rho := (\rho_n, n \geq 0)$ may now replace the role of $(W_t, t \leq \tau_D)$ in (1.2), and hence in (1.3), albeit that one must stop the sequence ρ at some finite N . By picking a threshold $\varepsilon > 0$, we can choose $N(\varepsilon)$ as a cut-off for the sequence ρ , such that $N(\varepsilon) = \min\{n \geq 0 : \inf_{z \in \partial D} |\rho_n - z| \leq \varepsilon\}$. Intuitively, one is inwardly ‘thickening’ the boundary ∂D with an ‘ ε -skin’ and stopping once the walk-on-spheres hits the ε -skin. As the sequence ρ is random, $N(\varepsilon)$ is also random. Starting with Theorem 6.6 of

Muller (1956) and the classical computations in Mutoo (1959), it is known that $\mathbb{E}_x[N(\varepsilon)] = \mathcal{O}(|\log \varepsilon|)$. To be more precise, we have the following result.

THEOREM 1.2 Suppose that D is a convex domain. There exist constants $c_1, c_2 > 0$, such that $\mathbb{E}_x[N(\varepsilon)] \leq c_1 |\log \varepsilon| + c_2$, $\varepsilon \in (0, 1)$.

The Monte Carlo simulation (1.3) can now be replaced by the one based on simulating the quantity $g(\rho_{N(\varepsilon)})$, $\rho_0 = x \in D$, which, in turn, is justified by the strong law of large numbers:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n g(\rho_{N^i(\varepsilon)}) = \mathbb{E}_x [g(\rho_{N(\varepsilon)})] \approx \mathbb{E}_x [g(W_{\tau_D})] = u(x), \quad \text{a.s.}, \quad (1.4)$$

where $\varepsilon > 0$ is some threshold and $(\rho_n^i, n \leq N^i(\varepsilon)), i \geq 0$ are *i.i.d.* copies of the walk-on-spheres process stopped at a distance ε or smaller from ∂D . Formally speaking, a convention is required to evaluate g just inside the boundary ∂D in (1.4). In many cases, g can be evaluated without introducing any additional bias (Given *et al.*, 1997; Hwang *et al.*, 2001).

The Laplacian serves as the infinitesimal generator of Brownian motion, in the sense that, for appropriately smooth functions $\phi: \mathbb{R}^d \rightarrow \mathbb{R}$,

$$\lim_{t \rightarrow 0} \frac{\mathbb{E}_x[\phi(W_t)] - \phi(x)}{t} = \frac{1}{2} \Delta \phi(x), \quad x \in \mathbb{R}^d. \quad (1.5)$$

Intuitively speaking, this explains an underlying connection between the Dirichlet problem (1.1) and the Feynman–Kac representation of the solution (1.2). In this article, we consider the analogue of (1.1) when the operator $\Delta/2$ is replaced by the fractional Laplacian $-(-\Delta)^{\alpha/2}$ for $\alpha \in (0, 2)$. In this case, the fractional Laplacian corresponds, in the same sense as (1.5), to an isotropic stable Lévy process with index α . This is a strong Markov process with stationary and independent increments, say $X = (X_t, t \geq 0)$ with probabilities $(\mathbb{P}_x, x \in \mathbb{R}^d)$, whose semigroup is represented by the Fourier transform

$$\mathbb{E}_0 [e^{i\langle \theta, X_t \rangle}] = e^{-|\theta|^\alpha t}, \quad \theta \in \mathbb{R}^d, t \geq 0,$$

where $\langle \cdot, \cdot \rangle$ represents the usual Euclidian inner product. Stable processes enjoy an isotropy in the following sense: if U is any orthogonal matrix in $\mathbb{R}^{d \times d}$, then $(UX_t, t \geq 0)$ under \mathbb{P}_0 has the same law as (X, \mathbb{P}_0) . Moreover, we have the following important scaling property: for all $c > 0$,

$$((cX_{c^{-\alpha}t}, t \geq 0), \mathbb{P}_0) \text{ is equal in law to } ((X_t, t \geq 0), \mathbb{P}_0). \quad (1.6)$$

In dimension two or greater, the operator $-(-\Delta)^{\alpha/2}$ can be expressed in the form

$$-(-\Delta)^{\alpha/2} u(x) = -\frac{2^\alpha \Gamma((d+\alpha)/2)}{\pi^{d/2} \Gamma(-\alpha/2)} \lim_{\varepsilon \downarrow 0} \int_{\mathbb{R}^d \setminus B(0, \varepsilon)} \frac{[u(y) - u(x)]}{|y-x|^{d+\alpha}} dy, \quad x \in \mathbb{R}^d,$$

where $B(0, \varepsilon) = \{x \in \mathbb{R}^d : |x| < \varepsilon\}$ and u is smooth enough for the limit to make sense.

Noting that $-(-\Delta)^{\alpha/2}$ is no longer a local operator, the analogous formulation of (1.1) needs a little more care. In particular, the boundary condition on the domain D is no longer stated on ∂D , but must now be stated on the complement of D , written D^c . To avoid pathological cases, we must assume throughout

that D^c has positive d -dimensional Lebesgue measure. The Dirichlet problem for $(-\Delta)^{\alpha/2}$ requires one to find $u: D \rightarrow \mathbb{R}$, such that

$$\begin{aligned} -(-\Delta)^{\alpha/2}u(x) &= 0, & x \in D, \\ u(x) &= g(x), & x \in D^c, \end{aligned} \tag{1.7}$$

where g is a suitably regular function. The fractional Dirichlet problem and variants thereof appear in many applications, in particular in physical settings where anomalous dynamics occur and where the spread of mass grows faster than linearly in time. Examples include turbulent fluids, contaminant transport in fractured rocks, chaotic dynamics and disordered quantum ensembles (see [Shlesinger *et al.*, 1995](#); [Klages *et al.*, 2008](#); [Klafter *et al.*, 2011](#)). The numerical analysis of (1.7) is no less deserving than in the diffusive setting.

Just as with the classical Dirichlet setting, the solution to (1.7) has a Feynman–Kac representation, expressed as an expectation at first exit from D of the associated stable process. The theorem below is proved in this article in a probabilistic way. Similar statements and proofs we found in the existing literature take a more analytical perspective. See for example, the review in [Bucur \(2016\)](#), as well as the monographs [Bliedtner & Hansen \(1986\)](#), [Bucur & Valdinoci \(2016\)](#) and [Bogdan *et al.* \(2009\)](#); the articles [Bogdan & Byczkowski \(1999\)](#), [Ros-Oton & Serra \(2014\)](#) and [Ros-Oton \(2016\)](#) and references therein.

We say a real-valued function ϕ on a Borel set $S \subset \mathbb{R}^d$ belongs to $L^1_\alpha(S)$ if it is a measurable function that satisfies

$$\int_S \frac{|\phi(x)|}{1 + |x|^{\alpha+d}} dx < \infty. \tag{1.8}$$

THEOREM 1.3 For dimension $d \geq 2$, suppose that D is a bounded domain in \mathbb{R}^d and that g is a continuous function in $L^1_\alpha(D^c)$. Then, there exists a unique continuous solution to (1.7) in $L^1_\alpha(\mathbb{R}^d)$, which is given by

$$u(x) = \mathbb{E}_x[g(X_{\sigma_D})], \quad x \in D,$$

where $X = (X_t, t \geq 0)$ is an isotropic stable Lévy process with index α and $\sigma_D = \inf\{t > 0 : X_t \notin D\}$.

The case that D is a ball can be found, for example, in Theorem 2.10 of [Bucur \(2016\)](#). We exclude the case $d = 1$, because convex domains are intervals for which exact solutions are known; see again [Bucur \(2016\)](#) or the forthcoming Theorem 3.1 lifted from [Blumenthal *et al.* \(1961\)](#). Theorem 1.3 follows in fact as a corollary of a more general result stated later in Theorem 6.1, which is proved in the appendix.

In this article, our objective is to demonstrate that the walk-on-spheres method may also be extended to the setting of the Dirichlet problem with fractional Laplacian. In particular, we will show that, thanks to various distributional and path properties of stable processes, notably spatial homogeneity, isotropy, self-similarity and that it exits D by a jump, simulations can be made unbiased, without the need to truncate the algorithm at an ε tolerance. Although there exist many methods for numerically examining the fractional Dirichlet problem (1.7), which mostly appeal to classical methodology for diffusive operators (see, for example, [Zoia *et al.*, 2007](#); [Huang & Oberman, 2014](#); [Dybiec & Szczepaniec, 2015](#); [Szczepaniec & Dybiec, 2015](#); [Acosta *et al.*, 2016](#); [D’Elia & Gunzburger, 2016](#); [Nochetto *et al.*, 2016](#)), to name a few, but not all of the existing literature, we believe that no other work appeals to the walk-on-spheres algorithm in this context.

The remainder of this article is structured as follows. In the next section, we give a brief historical review of Theorem 1.2 and its proofs as well as providing a new, short proof. In Section 3, we show how an old result of Blumenthal *et al.* (1961) can be used to give an exact simulation of the paths of stable processes. In Section 4, we introduce the walk-on-spheres algorithm for the fractional Laplacian Dirichlet problem. We start with domains D that are convex, but not necessarily bounded. Our main result shows that the walk-on-spheres algorithm ends in an a.s. finite number of steps (without the need for approximation), which can be stochastically bounded by a geometric distribution. Moreover, the parameter of this distribution does not depend on the starting point of the walk-on-spheres algorithm. Section 5 looks at the extensions to nonconvex domains. In Section 6, we consider a fractional Poisson equation, where an inhomogeneous term is introduced on the right-hand side of the fractional Laplacian Dirichlet problem (1.7). Appealing to the related results concerning the resolvent of stable processes until first exit from the unit ball, we are able to develop the walk-on-spheres algorithm further. Finally, in Section 7, we discuss some numerical experiments to illustrate the methods developed as well as their implementation.

2. The classical setting

As promised above, we give a brief historical review of the classical walk-on-spheres algorithm and, below, for completeness, we provide a proof of Theorem 1.2, which, to the authors’ knowledge, is new. The walk-on-spheres algorithm was first derived by Muller (1956). In Theorem 6.1 of his article, Muller claims that one can compare $\mathbb{E}_x[N(\varepsilon)]$ with the mean number of steps of a walk-on-spheres process that is stopped when it reaches an ε -skin of the tangent hyperplane that passes through a point on ∂D that is closest to x . Although the claim is correct (indeed the proof that we give for our main result Theorem 4.1 below provides the basis for an alternative justification of this fact), it is not entirely clear from Muller’s reasoning. Motoo (1959) uses Muller’s comparison of the mean number of steps to prove Theorem 1.2. He considers the total expected occupation of an appropriately time changed version of Brownian motion when crossing each sphere of the walk until touching the aforementioned ε -skin of the tangent hyperplane. Using the self-similarity of Brownian motion, Motoo argues that the time change during passage to the boundary of each sphere is such that the expected occupation across each step is uniformly bounded below. It follows that the sum of these weighted expected occupations can be bounded below by $\mathbb{E}_x[N(\varepsilon)]$. On the other hand, the aforesaid sum can also be bounded above by the total expected time-changed occupation until exiting the half-space (as defined by the tangent plane), which can be computed explicitly, thereby providing the $|\log \varepsilon|$ comparison.

Following the foundational work of Muller and Motoo, there have been many reproofs and generalizations of the original algorithm to different processes and domain types. Notable in this respect are the works of Mikhailov (1979) and Binder & Braverman (2012), who consider nonconvex domains, and Sabelfeld (1991), who appeals to renewal theory to analyse the growth in ε of the mean number of steps to completion of the walk-on-spheres algorithm. His method also allows for variants of the algorithm in which the sphere sizes do not need to be optimally inscribed in D . Later, Sabelfeld & Talay (1995) gave an elementary proof of the $|\log \varepsilon|$ bound. Mascagni and colleagues have extensively developed the walk-on-spheres algorithm in applications (see, for example, Given *et al.*, 2001, 2002; Hwang & Mascagni, 2001; Hwang *et al.*, 2003; Mackoy *et al.*, 2013).

Proof Theorem 1.2. We break the proof into two parts. In the first part, we analyse the walk-on-spheres process over one step, by considering the distance of the next point in the algorithm from the orthogonal tangent hyperplane of the first point. (Note the existence of a tangent hyperplane requires convexity of




FIG. 1. Geometric setting of the proof.

the domain.) In the second part of the proof, we use this analysis to build a supermartingale, from which the desired result follows via optional stopping.

For the first part of the proof, we start by introducing notation. For any $x = (x_1, \dots, x_d) \in \mathbb{R}^d$, such that $x_1 > 0$, let us write $V(x) = \{(z_1, \dots, z_d) \in \mathbb{R}^d : z_1 > 0\}$ for the open half-space containing x and denote its boundary $\partial V(x) = \{(z_1, \dots, z_d) \in \mathbb{R}^d : z_1 = 0\}$. Suppose that we choose our coordinate system, so that $x \in D$ is such that $\rho_0 = x = (x_1, 0, \dots, 0)$ and $\partial V(\rho_0)$ is a tangent hyperplane to both D and S_1 . This assumption comes at no cost, thanks to the isotropy and the spatial homogeneity of Brownian motion. Let us define ζ_0 , the orthogonal distance of ρ_0 from $\partial V(\rho_0)$. With the assumed choice of coordinate system, write $\zeta_0 := r_1 = x_1 = |x| = |\rho_0|$ and define

$$\zeta_1 = \min \left\{ \varepsilon, \inf_{z \in \partial V(\rho_0)} |\rho_1 - z| \right\};$$

that is, the minimum of ε and the orthogonal distance of ρ_1 from $\partial V(\rho_0)$. Next, define θ_1 , the angle that subtends at ρ_0 between ρ_1 and the origin $(0, \dots, 0)$ and recall that symmetry implies that θ_1 is uniformly distributed on $[0, 2\pi]$. Simple geometric considerations tell us that

$$\zeta_1 = x_1 - r_1 \sin \left(\frac{\pi}{2} - \theta_1 \right) = \zeta_0 - \zeta_0 \sin \left(\frac{\pi}{2} - \theta_1 \right) = \zeta_0 (1 - \cos(\theta_1)). \quad (2.1)$$

This provides an implicit expression for θ_1 in terms of the orthogonal distance ρ_0 from the nearest tangent hyperplane. See Fig. 1.

Assuming that $\zeta_0 > \varepsilon$, thanks to the isotropic symmetry, the walk-on-sphere algorithm will end at the first step if θ_1 lies in a certain critical interval dictated by the choice of skin thickness ε . We can compute

this critical (and obviously) symmetric interval as a function of ζ_0 , say $(-\theta^*(\zeta_0), \theta^*(\zeta_0))$, where

$$\theta^*(\zeta_0) = \arccos\left(\frac{\zeta_0 - \varepsilon}{\zeta_0}\right). \quad (2.2)$$

A quantity that will be of interest to us to complete the proof is the expectation $\mathbb{E}_x[\sqrt{\zeta_1}] = \mathbb{E}_{\rho_0}[\sqrt{\zeta_1}]$. To this end, we compute

$$\begin{aligned} \mathbb{E}_{\rho_0}[\sqrt{\zeta_1}] &\leq \sqrt{\varepsilon} \mathbb{P}_{\rho_0}(\theta_1 \in (-\theta^*(\zeta_0), \theta^*(\zeta_0))) + \mathbb{E}_{\rho_0}[\mathbf{1}_{(\theta_1 \notin (-\theta^*(\zeta_0), \theta^*(\zeta_0)))} \sqrt{\zeta_1}] \\ &= \sqrt{\varepsilon} \frac{\theta^*(\zeta_0)}{\pi} + \frac{1}{\pi} \int_{\theta^*(\zeta_0)}^{\pi} \sqrt{\zeta_0(1 - \cos(u))} du \\ &=: \Lambda(\varepsilon/\zeta_0) \sqrt{\zeta_0}, \end{aligned} \quad (2.3)$$

where $\mathbf{1}_S$ denotes the indicator function on the set S . Using the primitive $\int \sqrt{1 - \cos(u)} du = -2\sqrt{1 - \cos(u)} \cot(u/2)$, we have

$$\Lambda(u) = \sqrt{u} \frac{\arccos(1 - u)}{\pi} + \frac{2}{\pi} \sqrt{u} \cot\left(\frac{\arccos(1 - u)}{2}\right).$$

One easily verifies that there is a constant $\lambda \in (0, 1)$, such that $\sup_{u \in [0,1]} \Lambda(u) < \lambda$.

Next, we move to the second part of the proof. At each step of the walk-on-spheres, we can construct the quantities ζ_{n+1} , the orthogonal distance of ρ_{n+1} to the tangential hyperplane that passes through the closest point on ∂D to ρ_n ; and θ_n , the angle that is subtended at ρ_n between the aforesaid point and ρ_{n+1} . Note that ε is an absorbing state for the sequence $(\zeta_n, n \geq 0)$ in the sense that, if $\zeta_n = \varepsilon$, then $\zeta_{n+k} = \varepsilon$ for all $k \geq 0$. We may thus write $N(\varepsilon) \leq N'(\varepsilon) := \min\{n \geq 0 : \zeta_n = \varepsilon\}$.

By the strong Markov property and the spatial homogeneity of Brownian motion, given the analysis leading to (2.3), we have, on $\{n < N(\varepsilon)\}$,

$$\mathbb{E}\left[\sqrt{\zeta_{(n+1) \wedge N(\varepsilon)}} \mid \zeta_0, \dots, \zeta_n\right] = \mathbb{E}\left[\sqrt{\zeta_{(n+1) \wedge N(\varepsilon)}} \mid \zeta_n\right] \leq \Lambda(\varepsilon/\zeta_n) \sqrt{\zeta_n} < \lambda \sqrt{\zeta_n}.$$

As a consequence, the process $(\lambda^{-(n \wedge N(\varepsilon))} \sqrt{\zeta_{n \wedge N(\varepsilon)}}, n \geq 0)$ is a supermartingale. The optional sampling theorem and Jensen’s inequality give us

$$\varepsilon \lambda^{-\mathbb{E}_x[N'(\varepsilon)]} \geq \mathbb{E}_x[\lambda^{-N'(\varepsilon)} \varepsilon] \leq \sqrt{r_1}, \quad x \in D.$$

The result now follows by taking logarithms. □

3. Exact simulation of stable paths

The key ingredient to the walk-on-spheres in the Brownian setting is the knowledge that spheres are exited continuously and uniformly on the boundary of spheres. In the stable setting, the inclusion of path discontinuities means that the process will exit a sphere by a jump. The analogous key observation that

makes our analysis possible is the following result, which gives the distribution of a stable process issued from the origin, when it first exits a unit sphere.

THEOREM 3.1 (Blumenthal *et al.*, 1961) Suppose that $B(0, 1)$ is a unit ball centred at the origin and write $\sigma_{B(0,1)} = \inf\{t > 0 : X_t \notin B(0, 1)\}$. Then,

$$\mathbb{P}_0(X_{\sigma_{B(0,1)}} \in dy) = \pi^{-(d/2+1)} \Gamma(d/2) \sin(\pi\alpha/2) |1 - |y|^2|^{-\alpha/2} |y|^{-d} dy, \quad |y| > 1.$$

This result provides a method for constructing precise sample paths of stable processes in phase space (i.e., exploring sample paths as ordered subsets of \mathbb{R}^d rather than as functions $[0, \infty] \rightarrow \mathbb{R}^d$). Choose a tolerance ϵ and initial point $X_0 = x$. Denote by E_1 a sampling from the distribution given in Theorem 3.1. This gives the exit from a ball of radius one when X is issued from the origin. By the scaling property (1.6) and the stationary and independent increments, $x + \epsilon E_1$ is distributed as the exit position from a ball of radius ϵ centred at x when the process is issued from x . Hence, we define $X_1 = x + \epsilon E_1$ and then, inductively for $n \geq 1$, generate X_{n+1} as the exit point of the ball centred on X_n with radius ϵ by noting this is equal in distribution to $X_n + \epsilon E_{n+1}$, where E_{n+1} is an *i.i.d.* copy of E_1 . It is important to remark for later that the value of ϵ in this algorithm does not need to be fixed and may vary with each step. Note, however, the method does not generate the corresponding time to exit from each ball. Therefore, the sample paths that are produced, while being exact in the distribution of points that the stable process will pass through, cannot be represented graphically in time, as there is only an equal mean duration to exiting each sphere. If the tolerance ϵ is altered on each step, then even this mean duration feature is lost. The method is used to generate Fig. 2.

On account of classical Feynman–Kac representation, simulation of solutions to parabolic and elliptic equations involving the fractional Laplacian and, more generally, the infinitesimal generator of a Lévy process are synonymous with the simulation of the paths of the associated stochastic process. On account of the fact that such equations occur naturally in mathematical finance in connection with (exotic) option pricing, there are already many numerical and stochastic methods in existence for the general Lévy setting. The reader is referred, for example, to the books Cont & Tankov (2004), Boyarchenko & Levendorskiĭ (2002) and the references therein. Other sources offering simulation techniques can be found (e.g., Janicki & Weron, 1994; Asmussen & Rosiński, 2001; Cohen & Rosinski, 2007; Cohen *et al.*, 2010). Similarly to works in mathematical finance, they are mostly focused on the approximation of the stable process (and indeed the general Lévy process) by a compound Poisson process or a power series representation of the path, with a diffusive component to mimic the effect of small jumps. To our knowledge, however, the walk-on-spheres approach to path simulation has not been used in the context of simulating stable processes to date, nor, as alluded to above, to the end of numerically solving Dirichlet-type problems for the fractional Laplacian.

4. Walk-on-spheres for the fractional Laplacian

We start by describing the walk-on-spheres for the fractional Laplacian Dirichlet problem (1.7) on a convex domain D . The domain D may be unbounded, as long as D^c has nonzero measure (even though Theorem 1.3 requires boundedness). Fix $x \in D$. The walk-on-spheres $(\rho_n, n \geq 0)$, with $\rho_0 = x$ is defined in a similar way to the Brownian setting in the sense that, given ρ_{n-1} , the distribution of ρ_n is selected according to an independent copy of $X_{\sigma_{B_n}}$ under $\mathbb{P}_{\rho_{n-1}}$, where $B_n = \{x \in \mathbb{R}^d : |x - \rho_{n-1}| < r_n\}$ and $\sigma_{B_n} = \inf\{t > 0 : X_t \notin B_n\}$. The algorithm comes to an end at the random index $N = \min\{n \geq 0 : \rho_n \notin D\}$, again using the standard understanding that $\min \emptyset := \infty$. See, for example, the depiction in Fig. 3.

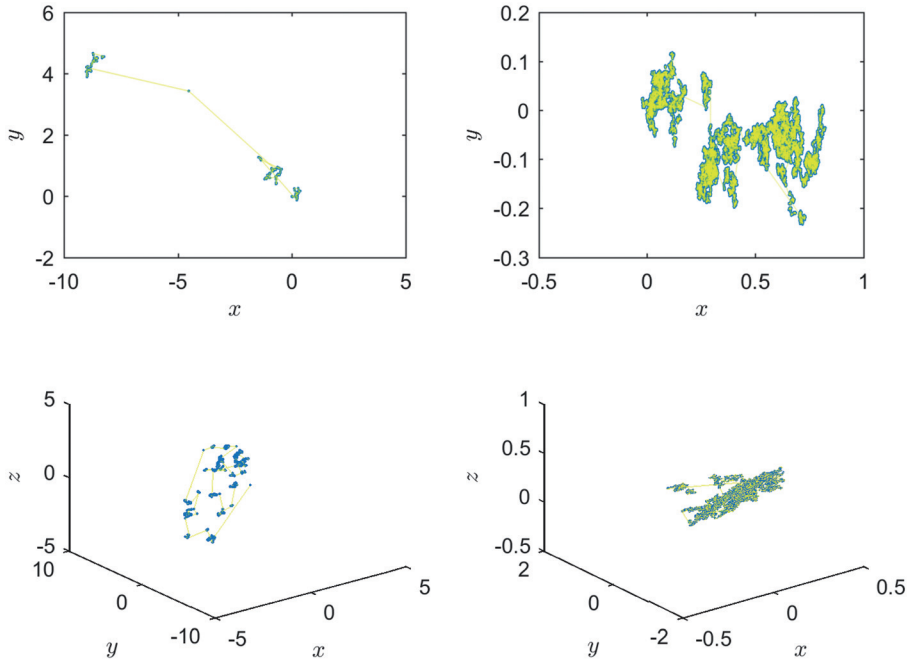


FIG. 2. Example sample paths for the α -stable Levy process generated using the exit distribution in Theorem 3.1 for spheres of radius 10^{-6} . Rows show sample paths in two and three dimensions for $\alpha = 0.9$ (left) and $\alpha = 1.8$ (right). The yellow lines indicate jumps of the process and blue dots show where the process has been.




FIG. 3. Steps of the walks-on-sphere algorithm until exiting the convex domain D in the stable setting. In this realization, $N = 3$.

Even though the domain D may be unbounded, our main result predicts that, irrespective of the point of issue of the algorithm, there will always be at most a geometrically distributed number of steps (whose parameter also does not depend on the point of issue) before the algorithm ends.

THEOREM 4.1 Suppose that D is a convex domain. For all $x \in D$, there exists a constant $p = p(\alpha, d) > 0$ (independent of x and D) and a real-valued random variable Γ , such that $N \leq \Gamma$ a.s., where

$$\mathbb{P}(\Gamma = k) = (1 - p)^{k-1} p, \quad k \in \mathbb{N}.$$

There are a number of remarks that we can make from the conclusion above.

1. Although Γ has the same distribution for each $x \in D$, it is not the same random variable for each $x \in D$. As we shall see in the proof of the above theorem, the inequality $N \leq \Gamma$ is derived by comparing each step of the walk-on-spheres algorithm with a sequence of Bernoulli random variables. This sequence of Bernoulli random variables are defined up to null sets, which may be different under each \mathbb{P}_x . Therefore, although the distribution of Γ does not depend on x , its null sets do.
2. The stochastic domination in Theorem 4.1 is much stronger than the usual comparison of the mean number of steps. Indeed, although it immediately implies that $\mathbb{E}_x[N] = 1/p$, we can also deduce that there is an exponentially decaying tail in the distribution of the number of steps. Specifically, for any $x \in D$,

$$\mathbb{P}(N > n) \leq \mathbb{P}(\Gamma > n) = (1 - p)^n, \quad n \in \mathbb{N}.$$

3. The randomness in the geometric random variables Γ is heavily correlated to N . The fact that each of the Γ are geometrically distributed has the advantage that

$$\sup_{x \in D} \mathbb{E}_x[N] \leq \sup_{x \in D} \mathbb{E}_x[\Gamma] = \frac{1}{p}.$$

However, it is less clear as to what kind of distributional properties can be said of the random variable $\sup_{x \in D} \Gamma$, which a.s. upper bounds $\sup_{x \in D} N$.

Finally, it is worth stating formally that the walk-on-spheres algorithm is unbiased and therefore, providing $\mathbb{E}_x[g(X_{\tau_D})] < \infty$, the strong law of large numbers applies and a straightforward Monte Carlo simulation of the solution to (1.7) is possible. Moreover, providing $\mathbb{E}_x[g(X_{\tau_D})^2] < \infty$, the central limit theorem offers the rate of convergence.

COROLLARY 4.2 When D is bounded and convex and g is continuous and in $L^1_\alpha(D^c)$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n g(\rho_N^i) = \mathbb{E}_x[g(\rho_N)] = \mathbb{E}_x[g(X_{\tau_D})] = u(x), \quad (4.1)$$

a.s., where $(\rho_n^i, n \leq N^i)$, $i \geq 1$ are *i.i.d.* copies of the walk-on-spheres with $\rho_0^i = x \in D$, $i \geq 1$ and $u(x)$ is the solution to (1.7). Moreover, when

$$\int_{D^c} \frac{g(x)^2}{1 + |x|^{\alpha+d}} dx < \infty, \quad (4.2)$$

then $\text{Var}(g(\rho_N)) < \infty$ and, in the sense of weak convergence,

$$\lim_{n \rightarrow \infty} n^{1/2} \left(\frac{1}{n} \sum_{i=1}^n g(\rho_{N^i}) - u(x) \right) = \text{Normal}(0, \text{Var}(g(\rho_N))).$$

Proof. The first part is a straightforward consequence of the earlier mentioned strong law of large numbers and the fact that Theorem 1.3 ensures that $\mathbb{E}_x[g(\rho_N)] = \mathbb{E}_x[g(X_{\tau_D})] < \infty$. For the second part, we need to show that (4.2) implies $\mathbb{E}_x[g(\rho_N)^2] = \mathbb{E}_x[g(X_{\tau_D})^2] < \infty$. However, if we consider the computation in (A.2) of the appendix, which shows that $\mathbb{E}_x[g(X_{\tau_D})] < \infty$ when g is continuous and in $L^1_\alpha(D^c)$, then it is easy to see that the same statement holds replacing g by g^2 . Under finiteness of the second moment, the central limit theorem completes the proof. \square

We now return to the proof of Theorem 4.1. Our approach is to break it into several parts. For convenience, we shall henceforth write $X^{(x)} = (X^{(x)}(t) : t \geq 0)$ to indicate the dependency of X on its initial position $X_0 = x$ (equivalent to writing (X, \mathbb{P}_x)). For any $x = (x_1, \dots, x_d) \in \mathbb{R}^d$, such that $x_1 > 0$, we have $V(x) = \{(z_1, \dots, z_d) \in \mathbb{R}^d : z_1 > 0\}$ for the open half-space containing x and denote its boundary $\partial V(x) = \{(z_1, \dots, z_d) \in \mathbb{R}^d : z_1 = 0\}$. For any Borel set $A \subset \mathbb{R}^d$, we write $\sigma_A = \inf\{t > 0 : X_t \notin A\}$. We will typically use in place of A the set $V(x)$ as well as $B(x, 1) = \{z \in \mathbb{R}^d : |z - x| < 1\}$, the unit ball centred at $x \in \mathbb{R}^d$. Finally write $\mathbf{i} = (1, 0, \dots, 0) \in \mathbb{R}^d$.

LEMMA 4.3 Without loss of generality (by appealing to the spatial homogeneity of X , which allows us to appropriately choose our coordinate system) suppose that $x = |x| \mathbf{i} \in D$ is such that $\partial V(x)$ is a tangent hyperplane to both D and B_1 . Then, $X_{\sigma_{B_1}}^{(x)}$ is equal in distribution to $|x| X_{\sigma_{B(\mathbf{i}, 1)}}^{(\mathbf{i})}$ and $X_{\sigma_{V(x)}}^{(x)}$ is equal in distribution to $|x| X_{\sigma_{V(\mathbf{i})}}^{(\mathbf{i})}$.

Proof. The scaling property of X ensures that we can write

$$X_s^{(x)} = |x| \hat{X}_{|x|^{-\alpha}s}^{(\mathbf{i})}, \quad s \geq 0, \quad (4.3)$$

where $\hat{X}^{(x)}$ is equal in law to $X^{(x)}$. Note that

$$\begin{aligned} \sigma_{B_1} &= \inf \{t > 0 : X^{(x)}(t) \notin B(x, |x|)\} \\ &= |x|^\alpha \inf \left\{ |x|^{-\alpha} t > 0 : |x| \hat{X}^{(\mathbf{i})}(|x|^{-\alpha} t) \notin B(x, |x|) \right\} \\ &= |x|^\alpha \inf \left\{ u > 0 : \hat{X}^{(\mathbf{i})}(u) \notin B(\mathbf{i}, 1) \right\} \\ &=: |x|^\alpha \hat{\sigma}_{B(\mathbf{i}, 1)}. \end{aligned} \quad (4.4)$$

It follows that

$$X_{\sigma_{B_1}}^{(x)} = |x| \hat{X}_{|x|^{-\alpha} |x|^\alpha \hat{\sigma}_{B(\mathbf{i}, 1)}}^{(\mathbf{i})} \stackrel{d}{=} |x| X_{\sigma_{B(\mathbf{i}, 1)}}^{(\mathbf{i})}, \quad (4.5)$$

as required. The proof of the second claim follows the same steps and is omitted for the sake of brevity. \square

An important consequence of the previous result is the comparison between the first exit from the largest sphere in D centred at x and the first exit from the tangent hyperplane to the latter sphere. Recall that $B_n = \{z \in \mathbb{R}^d : |z - \rho_{n-1}| < r_n\}$ denotes the n th sphere.

COROLLARY 4.4 Suppose that $x \in D$ is such that $\partial V(x)$ is a tangent hyperplane to both D and B_1 . Define under \mathbb{P}_x the indicator random variables

$$I_D = \mathbf{1}_{\{X_{\sigma_{B_1}} \notin D\}} \quad \text{and} \quad I_V = \mathbf{1}_{\{X_{\sigma_{B_1}} \notin V(x)\}}.$$

Then, $\mathbb{P}_x(I_D \geq I_V) = 1$ and, independently of $x \in D$, $\mathbb{P}_x(I_V = 1) = p(\alpha, d)$, where

$$\begin{aligned} p(\alpha, d) &:= \mathbb{P}_i(X_{\sigma_{B(i,1)}} \notin V(\mathbf{i})) \\ &= \frac{\Gamma(d/2)}{\pi^{(d+2)/2}} \sin(\pi\alpha/2) \int_{|x| < 1} |1 - |x|^2|^{-\alpha/2} |x|^{-d} dx, \end{aligned}$$

which is a number in $(0, 1)$.

Proof. The inequality follows from the inclusion $D \subset V(x)$. The formula for $p(\alpha, d)$ uses the coordinate system and scaling property of stable processes in Lemma 4.3 as well as the identity for the first exit from a sphere given by Theorem 3.1. \square

We are now ready to prove our main result.

Proof of Theorem 4.1. Suppose we condition on the previous positions of the walk-on-spheres, $\rho_0, \dots, \rho_{k-1}$ as well as on the event $\{N > k - 1\}$. Thanks to the stationary and independent increments as well as isotropy in the law of a stable process, we can always choose a coordinate system, or equivalently reorient D in such a way that $\rho_k = |\rho_k|\mathbf{i}$. This has the implication that, with the aforesaid conditioning, the random variable $\mathbf{1}_{\{N=k\}}$ is independent of $\rho_0, \dots, \rho_{k-1}$ and equal in law to $I_D(\rho_{k-1})$, where we have abused our original notation to indicate the initial position of X in the definition of I_D . Similarly, with the same abuse of notation, the event $I_V(\rho_{k-1})$ is independent of $\rho_0, \dots, \rho_{k-1}$ and equal in law to a Bernoulli random variable with a probability of success $p = p(\alpha, d)$. In particular, the sequence $I_V(\rho_k)$, $k \geq 0$ is a sequence of Bernoulli trials. That is to say, if we define

$$\Gamma = \min\{k \geq 1 : I_V(\rho_k) = 1\},$$

then it is geometrically distributed with parameter p . Thanks to Corollary 4.4, we also have that $\mathbb{P}_x(I_D \geq I_V)|_{x=\rho_k} = 1$, $k < N$, that is to say, $\{I_V(\rho_k) = 1\}$ a.s. implies $\{I_D(\rho_k) = 1\}$, for $k < N$, and hence

$$\min\{n \geq 1 : I_D(\rho_k) = 1\} \leq \min\{n \geq 1 : I_V(\rho_k) = 1\}$$

a.s. In other words, we have $N \leq \Gamma$, a.s., as required. \square




FIG. 4. A domain that satisfies the regularized uniform exterior-cone condition.

5. Nonconvex domains

The key element in the proof of Theorem 4.1 is the comparison of the event that the next step of the walk-on-spheres exits the domain D with the event that the next step of the walk-on-spheres exits a larger, more regular domain. More precisely, the aforesaid regular domain is taken to be the half-space that contains D with boundary hyperplane that is tangent to both the current maximal sphere and D . It is the use of a half-space that allows us to work with unbounded domains, but that forces the assumption that D is convex. With a little more care, we can remove the need for convexity without disturbing the main idea of the proof. However, this will come at the cost of insisting that D is bounded. It does, however, open the possibility that D is not a connected domain. We give two results in this respect.

For the first one, we introduce the following definition, which has previously been used in the potential analysis of stable processes; (see, for example, [Chen & Song, 1998](#)).

DEFINITION 5.1 A domain D in \mathbb{R}^d is said to satisfy the *uniform exterior-cone condition* (UECC), if there exist constants $\eta > 0$, $r > 0$ and a cone

$$\text{Cone}(\eta) = \{x = (x_1, \dots, x_d) \in \mathbb{R}^d : |x| < \eta x_1\},$$

such that, for every $z \in \partial D$, there is a cone C_z with vertex z , isometric to $\text{Cone}(\eta)$ satisfying $C_z \cap B(z, r) \subset D^c$.

It is well known that, for example, bounded $C^{1,1}$ domains satisfy (UECC). We need a slightly more restrictive class of domains than those respecting UECC. See Fig. 4.

DEFINITION 5.2 We say that D satisfies the *regularized uniform exterior-cone condition* (RUECC), if it is UECC, and the following additional condition holds: for each $x \in D$, suppose that $\partial(x)$ is the closest point on the boundary of D to x . Then, the isometric cone that qualifies D as UECC can be placed with its vertex at $\partial(x)$ and symmetrically oriented around the line that passes through x and $\partial(x)$.

THEOREM 5.3 Suppose that D is open and bounded (but not necessarily connected) and satisfies RUECC. Then, for each $x \in D$, there exists a random variable \hat{T} such that $N \leq \hat{T}$ a.s. and

$$\mathbb{P}(\hat{T} = k) = (1 - \hat{q})^{k-1} \hat{q}, \quad k \in \mathbb{N},$$

for some $\hat{q} = \hat{q}(\alpha, D)$.

Proof. Reviewing the proof of Theorem 4.1, we note that it suffices to prove that, in the context of Corollary 4.4, for each $x \in D$, there exists a Bernoulli random variable \hat{J}_x with parameter \hat{q} (independent of x), such that $\mathbb{P}_x(I_D \geq \hat{J}_x) = 1$. To this end, we recall that, without loss of generality, we may choose our coordinate system such that $x = |x|\mathbf{i} \in D$ is such that $\partial(x) = 0$. The assumption that D is bounded implies that there exists an η such that $|x| \leq \eta$. From the definition of RUECC, we know that there exists an $r > 0$ and a cone, C_0 , with vertex at 0, the closest point on ∂D to x , which is symmetrically oriented around the line passing through x and 0, such that $C_{0,r} := C_0 \cap B(0, r) \subset D^c$. We have

$$\begin{aligned} \mathbb{P}_x(X_{\sigma_{B_1}} \in C_{0,r}) &= \mathbb{P}_{\mathbf{i}}(X_{\sigma_{B(\mathbf{i},1)}} \in C_{0,r/|x|}) \\ &\geq \mathbb{P}_{\mathbf{i}}(X_{\sigma_{B(\mathbf{i},1)}} \in C_{0,r/\eta}) \\ &= \frac{\Gamma(d/2)}{\pi^{(d+2)/2}} \sin(\pi\alpha/2) \int_{C_{-\mathbf{i}(r/\eta)}} |1 - |y|^2|^{-\alpha/2} |y|^{-d} dy \\ &=: \hat{q}, \end{aligned}$$

where $C_{z,u} := [C_0 \cap B(0, u)] - \{z\}$, for $z \in \mathbb{R}^d$ and $u > 0$. Note that \hat{q} is necessarily strictly positive. Taking account of scaling, we have \mathbb{P}_x , a.s. that

$$I_D \geq \mathbf{1}_{\{|x|^{-1}X_{\sigma_{B_1}}^{(x)} \in C_{0,r/|x|}\}} \geq \mathbf{1}_{\{|x|^{-1}X_{\sigma_{B_1}}^{(x)} \in C_{0,r/\eta}\}} =: \hat{J},$$

where \hat{J} is a Bernoulli random variable with parameter \hat{q} . Stochastic dominance, $N \leq \hat{J}$ a.s., follows by the same line of reasoning as in the proof of Theorem 4.1. \square

For the second result, we completely relax the geometrical requirements on D at the expense of efficiency. With an abuse of our earlier notation, we introduce

$$N(\varepsilon) = \min \left\{ n \geq 0 : \rho_n \notin D \text{ or } \inf_{z \in \partial D} |\rho_n - z| < \varepsilon \right\}.$$

Intuitively, $N(\varepsilon)$ is the step that exits the inner ε -thickened boundary of D .

THEOREM 5.4 Suppose that D is open and bounded (but not necessarily connected). Then, for all $x \in D$, there exists a constant $q_\varepsilon = q_\varepsilon(\alpha, D) > 0$ (independent of x) and a random variable Γ^ε such that $N \leq \Gamma^\varepsilon$ a.s., where

$$\mathbb{P}_x(\Gamma^\varepsilon = k) = (1 - q_\varepsilon)^{k-1} q_\varepsilon, \quad k \in \mathbb{N}.$$

Moreover, $q_\varepsilon = \mathcal{O}(\varepsilon^\alpha)$ as $\varepsilon \downarrow 0$. In particular

$$\mathbb{E}_x[N(\varepsilon)] = \mathcal{O}(\varepsilon^{-\alpha}), \quad \text{as } \varepsilon \downarrow 0. \quad (5.1)$$

Proof. Define

$$\delta := \inf \{ r > 0 : D \subset B(x, r) \text{ for all } x \in D \},$$

so that any sphere of radius δ centred at $x \in D$ contains D . Once again, we recall that, without loss of generality, we may choose our coordinate system such that $x = |x| \mathbf{i} \in D$ is such that $\partial V(x)$ is a tangent hyperplane to B_1 and such that $0 \in \partial B_1 \cap \partial V(x) \cap \partial D$. Then, taking account of scaling and that, for all $x \in D$ such that $\inf_{z \in \partial D} |x - z| \geq \varepsilon$, with the particular choice of coordinates described above, $\delta/|x| \leq \delta/\varepsilon$, we have

$$\mathbf{1}_{\{N(\varepsilon)=1\}} \geq \mathbf{1}_{\{X_{\sigma_{B_1}}^{(x)} \notin B(x, \delta)\}} = \mathbf{1}_{\{|x|^{-1} X_{\sigma_{B_1}}^{(x)} \notin B(\mathbf{i}, \delta/|x|)\}} \geq \mathbf{1}_{\{|x|^{-1} X_{\sigma_{B_1}}^{(x)} \notin B(\mathbf{i}, \delta/\varepsilon)\}}.$$

Recall, however, from (4.5) that $|x|^{-1} X_{\sigma_{B_1}}^{(x)} \stackrel{d}{=} X_{\sigma_{B(\mathbf{i}, 1)}}^{(\mathbf{i})}$. It therefore follows that, \mathbb{P}_x , a.s.,

$$\mathbf{1}_{\{N(\varepsilon)=1\}} \geq \mathbf{1}_{\{X_{\sigma_{B(\mathbf{i}, 1)}}^{(\mathbf{i})} \notin B(\mathbf{i}, \delta/\varepsilon)\}} =: J^\varepsilon,$$

where J^ε is a Bernoulli random variable with parameter

$$q_\varepsilon(\alpha, D) = \mathbb{P}_\mathbf{i}(X_{\sigma_{B(\mathbf{i}, 1)}} \notin B(\mathbf{i}, \delta/\varepsilon)) = \frac{\Gamma(d/2)}{\pi^{(d+2)/2}} \sin(\pi\alpha/2) \int_{|y| \geq \delta/\varepsilon} |1 - |y|^2|^{-\alpha/2} |y|^{-d} dy.$$

Reverting to generalized spherical polar coordinates, in particular recalling that the Jacobian with respect to Cartesian coordinates is no larger than $|x|^{d-1}$ (see [Blumenson \(1960\)](#)), we can estimate

$$q_\varepsilon(\alpha, D) \leq \frac{\Gamma(d/2)}{\pi^{(d+2)/2}} \sin(\pi\alpha/2) \int_{\delta/\varepsilon}^{\infty} r^{-(\alpha+1)} dr = \mathcal{O}(\varepsilon^\alpha).$$

Reviewing the line of reasoning in the proof of [Theorem 4.1](#), we see that this comparison of events on the first step can be repeated at each surviving step of the algorithm to deduce the claimed result. \square

The $\mathcal{O}(\varepsilon^{-\alpha})$ bound in (5.1) can be compared with the bounds achieved by [Binder & Braverman \(2012\)](#) for the classical walk-on-spheres with Brownian motion for domains with more general geometries than convex. The worst case in [Binder & Braverman \(2012\)](#) is $\mathcal{O}(\varepsilon^{2-4/a})$ for a parameter $a > 0$ (describing the domain’s thickness or fractal boundary). Notably in the limit $\alpha \rightarrow 2$ (X converges to Brownian motion) and $a \rightarrow \infty$ (the domain loses regularity), the two agree with an $\mathcal{O}(\varepsilon^{-2})$ bound.

6. Fractional Poisson problem

We are now interested in using the walk-on-spheres process to find the solution to the inhomogeneous version of (1.7), namely

$$\begin{aligned} -(-\Delta)^{\alpha/2} u(x) &= -f(x), & x \in D, \\ u(x) &= g(x), & x \in D^c, \end{aligned} \tag{6.1}$$

for suitably regular functions $f: D \rightarrow \mathbb{R}$ and $g: D^c \rightarrow \mathbb{R}$. We want to identify a Feynman–Kac representation for solutions to (6.1) for suitable assumptions on g, f and D . Throughout this section, we adopt the setting of the following theorem.

THEOREM 6.1 Let $d \geq 2$ and assume that D is a bounded domain in \mathbb{R}^d . Suppose that g is a continuous function that belongs to $L^1_\alpha(D^c)$. Moreover, suppose that f is a function in $C^{\alpha+\varepsilon}(\overline{D})$ for some $\varepsilon > 0$. Then, there exists a unique continuous solution to (6.1) in $L^1_\alpha(\mathbb{R}^d)$, which is given by

$$u(x) = \mathbb{E}_x[g(X_{\sigma_D})] + \mathbb{E}_x \left[\int_0^{\sigma_D} f(X_s) ds \right], \quad x \in D, \quad (6.2)$$

where $\sigma_D = \inf\{t > 0: X_t \notin D\}$.

The combinations of Theorem 2.10 and 3.2 in [Bucur \(2016\)](#) treat the case that D is a ball. In the more general setting, among others, [Bogdan & Byczkowski \(1999\)](#), [Ros-Oton & Serra \(2014\)](#) and [Ros-Oton \(2016\)](#) (see also citations therein) offer results in this direction, albeit from a more analytical perspective. We give a new probabilistic proof of Theorem 6.1 in the appendix using a method that combines the idea of walk-on-spheres with the version of Theorem 6.1 when D is a ball. It is for this reason that the (otherwise unclear) need for the assumption that $f \in C^{\alpha+\varepsilon}(\overline{D})$ enters. Note in particular that Theorem 1.3 follows as a corollary.

We can develop the expression in (6.2) in terms of the walk-on-spheres $(\rho_n, n \leq N)$, providing the basis for a Monte Carlo simulation. What will work to our advantage here is another explicit identity that appears in [Blumenthal *et al.* \(1961\)](#). Define

$$V_r(x, dy) := \int_0^\infty \mathbb{P}_x(X_t \in dy, t < \sigma_{B(x,r)}) dt, \quad x \in \mathbb{R}^d, |y| < 1, r > 0.$$

THEOREM 6.2 ([Blumenthal *et al.*, 1961](#)) The expected occupation measure of the stable process prior to exiting a unit ball centred at the origin is given, for $|y| < 1$, by

$$V_1(0, dy) = 2^{-\alpha} \pi^{-d/2} \frac{\Gamma(d/2)}{\Gamma(\alpha/2)^2} |y|^{\alpha-d} \left(\int_0^{|y|^{-2}-1} (u+1)^{-d/2} u^{\alpha/2-1} du \right) dy. \quad (6.3)$$

Although the above identity is presented in a probabilistic context, it has a much older history in the analysis literature. Known as Boggio's formula, the original derivation in the setting of potential theory dates back to [Boggio \(1905\)](#). See the discussion in [Delaurentis & Romero \(1990\)](#) and [Bucur \(2016\)](#).

In the next result, we will write as a slight abuse of notation $V_r(x, f(\cdot)) = \int_{|y-x|<r} f(y) V_r(x, dy)$ for bounded measurable f .

LEMMA 6.3 For $x \in D$, $g \in L^1_\alpha(D^c)$ and $f \in C^{\alpha+\varepsilon}(\overline{D})$, we have the representation

$$u(x) = \mathbb{E}_x[g(\rho_N)] + \mathbb{E}_x \left[\sum_{n=0}^{N-1} r_n^\alpha V_1(0, f(\rho_n + r_n \cdot)) \right].$$

Proof. Given the walk-on-spheres $(\rho_n, n \leq N)$ with $\rho_0 = x \in D$, define σ_n jointly with ρ_n so that, given ρ_{n-1} , (ρ_n, σ_n) is equal in law to $(X_{\sigma_{B_n}}, \sigma_{B_n})$ under $\mathbb{P}_{\rho_{n-1}}$. We can now represent the second expectation on the right-hand side of (6.2) in the form

$$\mathbb{E}_x \left[\sum_{n \geq 0} \mathbf{1}_{\{\rho_n \in D\}} \int_0^{\sigma_{n+1}} f(\rho_n + X_s^{(n+1)}) ds \right], \quad x \in D, \quad (6.4)$$

where $X^{(n)}$ are independent copies of (X, \mathbb{P}_0) . Applying Fubini's theorem, then conditioning each expectation on $\mathcal{F}_n := \sigma(\rho_k : k \leq n)$ followed by Fubini's theorem again, we have

$$\begin{aligned} \mathbb{E}_x \left[\int_0^{\sigma_D} f(X_s) ds \right] &= \sum_{n \geq 0} \mathbb{E}_x \left[\mathbf{1}_{\{\rho_n \in D\}} \mathbb{E}_y \left[\int_0^{\sigma_{B(y,r)}} f(X_s) ds \right] \Big|_{y=\rho_n, r=r_n} \right] \\ &= \sum_{n \geq 0} \mathbb{E}_x \left[\mathbf{1}_{\{\rho_n \in D\}} V_{r_n}(\rho_n, f(\cdot)) \right] \\ &= \mathbb{E}_x \left[\sum_{n=0}^{N-1} V_{r_n}(\rho_n, f(\cdot)) \right]. \end{aligned}$$

The proof is completed once we show that $V_r(x, g) = r^\alpha V_1(0, f(x + r \cdot))$, for $r > 0$, $x \in \mathbb{R}^d$ and bounded measurable f . To this end, we appeal to spatial homogeneity and the, now, familiar computations using the scaling property of stable processes:

$$\begin{aligned} V_r(x, f(\cdot)) &= \mathbb{E}_x \left[\int_0^{\sigma_{B(x,r)}} f(X_t) dt \right] \\ &= \mathbb{E}_0 \left[\int_0^{\sigma_{B(0,r)}} f(x + X_t) dt \right] \\ &= \mathbb{E}_0 \left[\int_0^{\sigma_{B(0,1)}} r^\alpha f(x + r X_s) ds \right] \\ &= \int_{|y| < 1} r^\alpha f(x + r y) V_1(0, dy) \\ &= r^\alpha V_1(0, f(x + r \cdot)). \end{aligned} \tag{6.5}$$

The proof is now complete. □

Lemma 6.3 now informs a Monte Carlo procedure based on simulating the quantity

$$\chi := g(\rho_N) + \sum_{n=0}^{N-1} r_n^\alpha V_1(0, f(\rho_n + r_n \cdot)), \quad x \in D,$$

which is again justified by an obvious strong law of large numbers and the central limit theorem in the spirit of Corollary 4.2.

COROLLARY 6.4 When D is bounded and convex, g is continuous and in $L_\alpha^1(D^c)$ and f is a function in $C^{\alpha+\varepsilon}(\bar{D})$ for some $\varepsilon > 0$, then

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \chi^i = \mathbb{E}_x \left[g(\rho_N) + \sum_{n=0}^{N-1} r_n^\alpha V_1(0, f(\rho_n + r_n \cdot)) \right] = u(x), \tag{6.6}$$

a.s., where χ^i , $i \geq 1$ are *i.i.d.* copies of χ and $u(x)$ is the solution to (6.1). Moreover, when

$$\int_{D^c} \frac{g(x)^2}{1 + |x|^{\alpha+d}} dx < \infty. \quad (6.7)$$

Then, $\text{Var}(\chi) < \infty$ and, in the sense of weak convergence,

$$\lim_{n \rightarrow \infty} n^{1/2} \left(\frac{1}{n} \sum_{i=1}^n \chi^i - u(x) \right) = \text{Normal}(0, \text{Var}(\chi)).$$

Proof. Theorem 6.1 and Lemma 6.3 ensure that the strong law of large numbers may be invoked. For the central limit theorem, we need $\mathbb{E}_x[\chi^2] < \infty$. Taking account of the fact that χ is the sum of two terms, the Cauchy–Schwarz inequality ensures that $\mathbb{E}_x[\chi^2]$ is finite if $\mathbb{E}_x[g(\rho_N)^2]$ and $\mathbb{E}_x \left[\left(\sum_{n=0}^{N-1} r_n^\alpha V_1(0, f(\rho_n + r_n \cdot)) \right)^2 \right]$ are finite. Recall that $\mathbb{E}_x[g(\rho_N)^2] = \mathbb{E}_x[g(X_{\sigma_D})^2]$ and, from Corollary 4.2, that (6.7) is sufficient to ensure that this expectation is bounded.

Now note that, on account of the fact that f is bounded, there exists a constant $\kappa \in (0, 1)$, such that, for each $n \leq N$, appealing to (6.5), we have $r_n^\alpha V_1(0, f(\rho_n + r_n \cdot)) \leq \kappa \sigma_n$, where σ_n is the time it takes for the walk-on-spheres to exit the n th sphere. Thus, $\sum_{n=0}^{N-1} r_n^\alpha V_1(0, f(\rho_n + r_n \cdot)) \leq \kappa \sum_{n=0}^{N-1} \sigma_n = \kappa \sigma_D$. We thus have that

$$\mathbb{E}_x \left[\left(\sum_{n=0}^{N-1} r_n^\alpha V_1(0, f(\rho_n + r_n \cdot)) \right)^2 \right] \leq \kappa^2 \mathbb{E}_x[\sigma_D^2].$$

However, the latter expectation can be bounded by $\mathbb{E}_x[\sigma_{B^*}^2]$, where $B^* = B(x, R)$ for some suitably large R , such that D is compactly embedded in B^* . Moreover, appealing to Gettoor (1961), we know that $\mathbb{E}_x[\sigma_{B^*}^2]$ is bounded. \square

7. Numerical experiments

In the following section, all of the routines associated with the simulations are publicly available at the following repository: https://bitbucket.org/wos_paper/wos_repo.

For the Monte Carlo procedure, independent copies of the walk-on-spheres $(\rho_n, n \leq N)$ need to be simulated whereby, by the Markov property, every new point in the sequence can be expressed as $\rho_{n+1} = \rho_n + X'_{\sigma'_{B(0, r_n)}}$, where X' is an independent version of X and

$$\sigma'_{B(0, r_n)} = \inf\{t > 0: X'_t \notin B(0, r_n)\}.$$

In other words, ρ_{n+1} is an exit point from a ball $B(0, r_n)$ under \mathbb{P}_0 translated by ρ_n . A consequence of Lemma 3.1 is that the exit distribution of X'_t from $B(0, r_n)$, $r_n > 0$, can be, via a change of variable $y = \tilde{y}/r_n$, written as

$$\mathbb{P}_0(X_{\sigma_{B(0, r_n)}} \in d\tilde{y}) = \pi^{-(d/2+1)} \Gamma(d/2) \sin(\pi\alpha/2) |r_n^2 - |\tilde{y}|^2|^{-\alpha/2} |\tilde{y}|^{-d} r_n^\alpha d\tilde{y}, \quad |\tilde{y}| > r_n. \quad (7.1)$$

For $d = 2$, it is more convenient to work with polar coordinates (r, θ) to separate variables in (7.1). Indeed, recalling that $d\tilde{y} = r dr d\theta$, we have

$$\mathbb{P}_0(X_{\sigma_{B(0,r_n)}} \in d\tilde{y}) = \frac{2}{\pi} \sin(\pi\alpha/2) (r^2 - r_n^2)^{-\alpha/2} r_n^\alpha \frac{dr}{r} \times \frac{d\theta}{2\pi}, \quad r > r_n. \quad (7.2)$$

From (7.2), we see that the angle θ is sampled uniformly on $[0, 2\pi]$, whereas we can sample the radius r via the inverse transform sampling method. To this end, noting that $\sin(\pi\alpha/2)B(\alpha/2, 1 - \alpha/2) = \pi$, the first factor on the right-hand side of (7.2) is the density of a distribution with cumulative distribution function F . The inverse of F can be identified as follows: For $x \in [0, 1]$,

$$F^{-1}(x) = r_n (I^{-1}(1 - x; \alpha/2, 1 - \alpha/2))^{-1/2},$$

where $I^{-1}(x; z, w)$ is the inverse of the incomplete beta function

$$I(x; z, w) := \frac{1}{B(z, w)} \int_0^x u^{z-1} (1-u)^{w-1} du, \quad x \in [0, 1],$$

and $B(z, w) := \int_0^1 u^{z-1} (1-u)^{w-1} du$ is the beta function.

The homogeneous part of the solution to (6.1) is somewhat easier to compute than the inhomogeneous part, which additionally involves numerical computation of the integral $r_n^\alpha V_1(0, f(\rho_n + r_n \cdot))$ in (6.6). To develop this expression, we use the substitution $u = (1-t)/t$ for the integral in (6.3) and hence, when $d = 2$, for $|y| < 1$,

$$V_1(0, dy) = c_{2,\alpha} B(1 - \alpha/2, \alpha/2) |y|^{\alpha-2} (1 - I(|y|^2; 1 - \alpha/2, \alpha/2)), \quad (7.3)$$

with $c_{2,\alpha} = 2^{-\alpha} \pi^{-1} \Gamma(\alpha/2)^{-2}$. Moreover, by converting to polar coordinates (r, θ) , the simulated quantity at step n becomes

$$\begin{aligned} & r_n^\alpha V_1(0, f(\rho_n + r_n \cdot)) \\ &= r_n^\alpha c_{2,\alpha} B(1 - \alpha/2, \alpha/2) \int_{|y|<1} f(\rho_n + r_n y) |y|^{\alpha-2} (1 - I(|y|^2; 1 - \alpha/2, \alpha/2)) dy \\ &= r_n^\alpha c_{2,\alpha} B(1 - \alpha/2, \alpha/2) 2\pi \alpha^{-1} \int_0^1 \int_{-\pi}^\pi f(\rho_n + r_n r (\cos \theta, \sin \theta)) \\ &\quad \times (1 - I(r^2; 1 - \alpha/2, \alpha/2)) \frac{d\theta}{2\pi} \times \alpha r^{\alpha-1} dr. \end{aligned}$$

We used the Monte Carlo approach for evaluating this integral. Consider independent random variables $\Theta \sim U(-\pi, \pi)$ and $R = X^{1/\alpha}$, such that $X \sim U(0, 1)$. Then, R has the probability density function $f_R(r) = \alpha r^{\alpha-1}$ and we want to evaluate

$$r_n^\alpha V_1(0, f(\rho_n + r_n \cdot)) = a_{2,\alpha} r_n^\alpha \mathbb{E} \left[(1 - I(R^2; 1 - \alpha/2, \alpha/2)) f(\rho_n + r_n R (\cos \Theta, \sin \Theta)) \right], \quad (7.4)$$

with $a_{2,\alpha} = \alpha^{-1} 2^{-\alpha+1} \Gamma(\alpha/2)^{-2} B(1 - \alpha/2, \alpha/2)$. We simulate $n_{R,\Theta}$ samples of pairs (R, Θ) and compute the sample mean of the quantity in (7.4). The quantity is evaluated more efficiently by writing

$$f(\rho_n + r_n R (\cos \Theta, \sin \Theta)) = f(\rho_n) + [f(\rho_n + r_n R (\cos \Theta, \sin \Theta)) - f(\rho_n)]. \quad (7.5)$$




FIG. 5. Example simulation for (1.7) with exterior data $g(x) = G(x, y)$ with $y = (2, 0)$ on the domain given by the unit ball centred at the origin, based on 10^6 samples. The left-hand plot shows the relative error and the right-hand plot shows the sample variance. The sample variance is larger for small α as the process stops further away from the boundary and can see the singularity at $(2, 0)$ in the exterior data. Accordingly, the relative error is higher as we are using a fixed number of samples.

This gives two terms: one can be evaluated directly (by storing $\mathbb{E}[1 - I(R^2; 1 - \alpha/2, \alpha/2)]$) and the second can be evaluated using a Monte Carlo method, but with smaller variance (as the quantity in square brackets in (7.5) is $\mathcal{O}(r_n)$). It is worth noting that a similar mixed approach using the trapezoidal rule over θ and randomizing r as earlier for evaluating the left-hand side of (7.4) was also tested. However, the results showed that the pure Monte Carlo approach is, in comparison with the mixed one, superior with regard to accuracy and computational cost. With this view, we decided to focus on the first one.

Accuracy of this algorithm and its feasibility of implementation were checked with model solutions to problems of the type (6.1) and they are presented below in order.

7.1 Free-space Green's function

The free-space Green's function for the fractional Laplacian $(-\Delta)^{\alpha/2}$ is

$$G(x, y) = c_{d,\alpha} \frac{1}{|x - y|^{\alpha-d}},$$

for a constant $c_{d,\alpha}$ for $d > 1$ and $\alpha \in (0, 2)$ (see Bucur, 2016). If the point y is chosen outside a domain D , then we can construct G as an exact solution to the homogeneous version of the fractional Dirichlet problem in (1.7); that is, $u(x) = G(x, y)$ for $x \in D$ and $g(x) = G(x, y)$ for $x \notin D$. Figure 5 shows the results of applying the walk-on-spheres algorithm to evaluate $u(0.6, 0.6)$ with 10^6 samples, where D is a unit ball in \mathbb{R}^2 centred at the origin and $y = (2, 0)$. We observe that the samples $g(\rho_N)$ have larger variance when α is small and a larger error results from the same number of samples.

7.2 Gaussian data

For the Poisson problem (6.1), we take D to be the unit ball in \mathbb{R}^2 , exterior data

$$g(x) = \exp(-|x - y|^2), \quad x \in D^c,$$

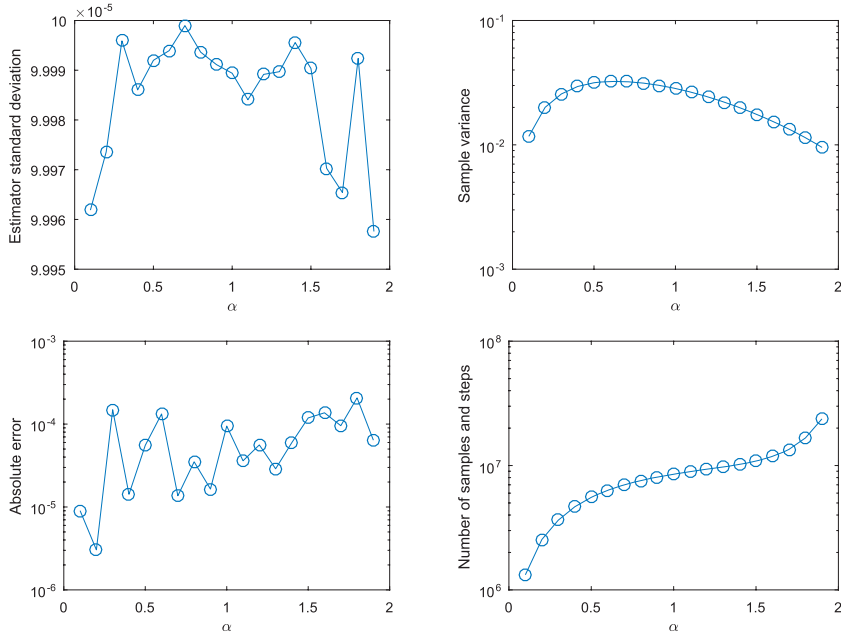


FIG. 6. Example simulation with the walk-on-spheres algorithm for (1.7) based on desired tolerance of 10^{-4} . From top left to bottom right, we see the standard deviation of the estimator, the sample variance, the absolute error (using a quadrature approximation for (7.6) for the reference value) and the amount work (number of samples \times mean number of steps).

for a given $y \in \mathbb{R}^2$ and zero source term $f = 0$. We can represent the solution to (1.7) in D by

$$u(x) = \pi^{-2} \sin(\pi \alpha / 2) \int_{D^c} \left(\frac{1 - |x|^2}{|y|^2 - 1} \right)^{\alpha/2} \frac{1}{|y - x|^2} \exp(-|x - y|^2) dy, \quad x \in D. \quad (7.6)$$

This integral can be computed numerically via a quadrature approximation. Here, instead of a fixed number of samples, the number of samples is taken adaptively based on a tolerance ε for the computed sample standard deviation. Figure 6 shows the results with $y = (2, 0)$ and tolerance $\varepsilon = 10^{-4}$ for the evaluation of $u(0.6, 0.6)$ as previously. The estimator standard deviation and absolute error exhibit no obvious trend, whereas the sample variance peaks at about $\alpha = 0.6$. Also, at this value, the largest number of samples is needed to satisfy the tolerance. Despite the sample variance decreasing after $\alpha = 0.6$, there is an increasing trend in the amount of work required. This implies that the increase in the number of steps with α (see Fig. 10) dominates, and therefore a solution point of accuracy 10^{-4} is computationally more costly for larger values of α .

7.3 Nonconstant source term

Suppose that, again in the context of (6.1), we again take D to be equal to the unit ball and the source term equal to

$$f(x) = 2^\alpha \Gamma(2 + \alpha/2) \Gamma(1 + \alpha/2) (1 - (1 + \alpha/2) \|x\|^2), \quad x \in D,$$

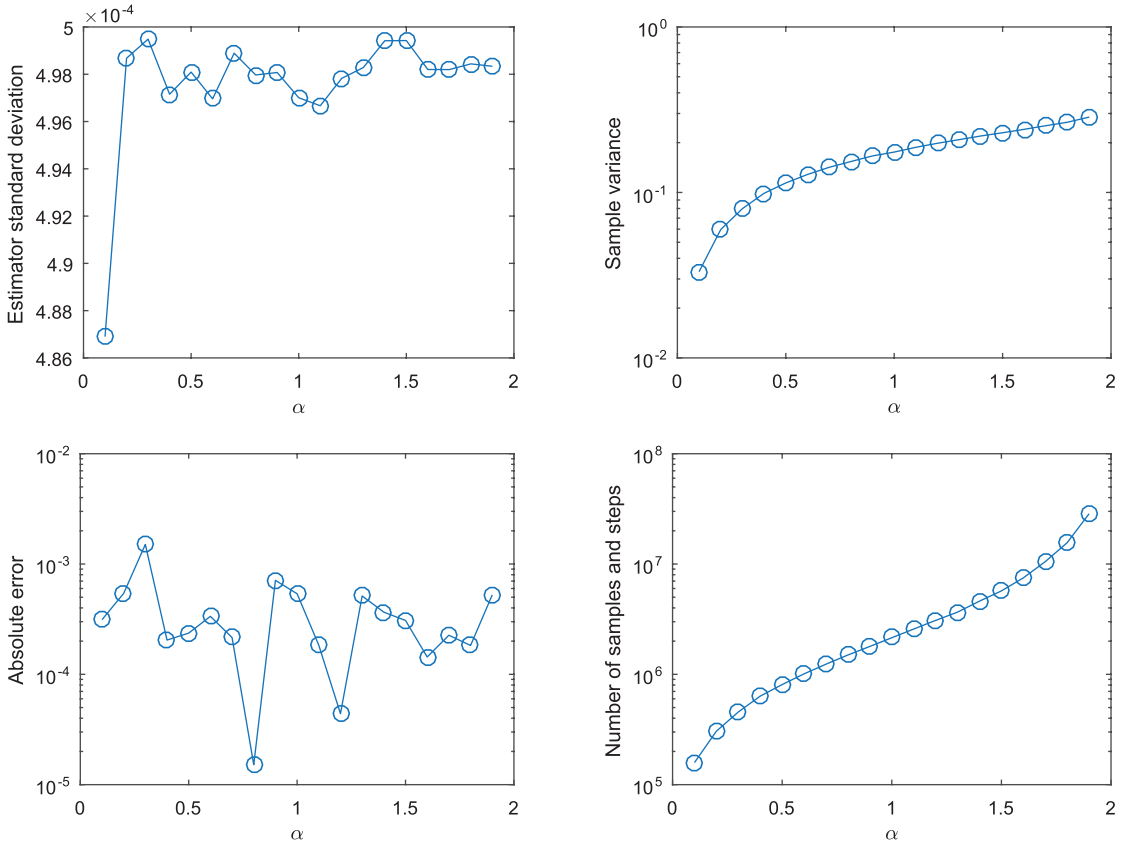


FIG. 7. Example simulation with the walk-on-spheres algorithm for (6.1) based on desired tolerance of 10^{-3} and $n_{R,\Theta} = 1000$. From top left to bottom right, we see the standard deviation of the estimator, the sample variance, the absolute error and the amount work (number of samples \times mean number of steps).

and zero exterior data $g = 0$. This has the exact solution $u(x) = \max\{0, 1 - \|x\|^2\}^{1+\alpha/2}$ (cf. Dyda, 2012). The behaviour of the algorithm is shown in Fig. 7. As expected, we again observe no obvious trend in estimator standard deviation and absolute error. The sample variance of sums of Monte Carlo-generated integrals increases with α , as does the number of samples accordingly. Work required grows with α as in Fig. 6, but with a slightly steeper trend. Notice that accuracy of 10^{-4} for the inhomogeneous part of the solution would demand a lot more work than the homogeneous part in Fig. 6.

7.4 Distribution of the number of steps in convex and nonconvex domains

In previous sections, a large focus was put on deriving upper bounds and limiting distributions for N . Here, we provide numerical support for these theoretical results. The walk-on-spheres algorithm was simulated inside a unit-ball domain centred at the origin as well as inside a domain of a hundred touching unit balls centred at points (i, j) , $i, j = -10, \dots, 10$, the so-called ‘Swiss cheese’ domain as shown in Figure 8. The first represents a convex domain, whereas the latter a nonconvex one. The algorithm was started at a point $x = (\sqrt{0.29}, -\sqrt{0.7})$ which lies very close to the boundary in both domains. This





FIG. 8. ‘Swiss cheese’ domain (interior of balls).

FIG. 9. Histogram of the proportion of runs of the walk-on-spheres algorithm with $\alpha = 1$, for which $N > n$ for the unit-ball domain (left) and the Swiss cheese domain (right). The dotted curve shows the tail of $\text{Geom}(p(1, 2))$, this is $(1 - p(1, 2))^n$, as in Remark 2 of Theorem 4.1.

point was chosen, as numerical simulations in a unit-ball domain revealed that the mean number of steps decreases with increasing distance from the boundary of the starting point. Theorem 4.1 states that N is stochastically dominated by a geometric distribution with parameter $p(\alpha, d)$. In two dimensions, we are able to numerically compute $p(\alpha, 2)$, because it is the solution to (1.7) with $D = B(0, 1)$, $g(x) = \mathbf{1}_{\{x_1 < -1\}}(x)$ and zero source term $f = 0$ as deduced from Corollary 4.4. We computed values of $p(\alpha, 2)$ for different α to accuracy 10^{-4} .

The left-hand histogram in Fig. 9 confirms the stochastic dominance of Γ and an exponentially decaying tail as stated in Remark 2 of Theorem 4.1. However, the right-hand histogram shows that this statement fails in the particular example of the Swiss cheese domain. Moreover, the plot of the mean number of steps against α in Fig. 10 shows the observed value of $\mathbb{E}_x[N]$ is bounded above by $1/p(\alpha, 2)$ for the unit-ball domain. On the other hand, this is not the case for the Swiss cheese domain, where the observed value of $\mathbb{E}_x[N]$ exceeds $1/p(\alpha, 2)$ for α in the range $(0.3, 1.6)$.




FIG. 10. Mean number of steps for the walk-on-spheres algorithm started at $x = (\sqrt{0.29}, -\sqrt{0.7})$ inside the circle domain (left) and inside the Swiss cheese domain (right). The dashed curve on both plots is $1/p(\alpha, 2)$ as in Corollary 4.4.

An explanation for why this is happening might be as follows. At larger values of α , the path of X starts resembling that of a Brownian motion (albeit with a countable infinity of arbitrarily small discontinuities). The process X is started inside a ball in the Swiss cheese. When it exits this ball, its exit position is relatively close to the boundary with high probability. Therefore, the exit point of the ball containing the point of issue is more likely to be in the ‘cheese’ (which would cause an end to the algorithm) and less likely to be inside another vacuous ball. Accordingly, $\mathbb{E}_x[N]$ does not deviate largely from the example of a single ball. However, for small values of α , exit points from the sphere containing the point of issue have a higher probability to be far from the boundary, landing inside another vacuous ball, thereby requiring the algorithm to continue. In that case, the comparison with the case of exiting a single sphere breaks down.

Acknowledgements

We would like to thank Mateusz Kwaśniki for pointing out a number of references to us and Alexander Freudentberg for a close reading of an earlier version of this article.

Funding

Engineering and Physical Sciences Research Council (EPSRC) grant (EP/L002442/1 to A.E.K.).

REFERENCES

- ACOSTA, G., BORTHAGARAY, J. P. BRUNO, O. & MAAS, M. (2016) *Regularity theory and high order numerical methods for the (1d)-fractional Laplacian*. arXiv: 1608.08443.
- ASMUSSEN, S. & ROSIŃSKI, J. (2001) Approximations of small jumps of Lévy processes with a view towards simulation. *J. Appl. Probab.*, **38**, 482–493.
- BINDER, I. & BRAVERMAN, M. (2012) The rate of convergence of the Walk on Spheres Algorithm. *Geom. Funct. Anal.*, **22**, 558–587.
- BLIEDTNER, J. & HANSEN, W. (1986) *Potential Theory. An Analytic and Probabilistic Approach to Balayage*. Berlin: Springer, xiv+435.
- BLUMENSON, L. E. (1960) Classroom notes: a derivation of n -dimensional spherical coordinates. *Amer. Math. Monthly*, **67**, 63–66.

- BLUMENTHAL, R. M., GETOOR, R. K. & RAY, D. B. (1961) On the distribution of first hits for the symmetric stable processes. *Trans. Amer. Math. Soc.*, **99**, 540–554.
- BOGDAN, K. & BYCZKOWSKI, T. (1999) Potential theory for the α -stable Schrödinger operator on bounded Lipschitz domains. *Studia Math.*, **133**, 53–92.
- BOGDAN, K., BYCZKOWSKI, T., KULCZYCKI, T., RYZNAR, M., SONG, R. & VONDRAČEK, Z. (2009) *Potential Analysis of Stable Processes and its Extensions*, vol. 1980. Lecture Notes in Mathematics. (P. Graczyk and A. Stos eds) Berlin: Springer, x+187.
- BOGGIO, T. (1905) Sulle funzioni di green d’ordine m. *Rend. Circ. Matem. Palermo* **20**, pp. 97–135.
- BOYARCHENKO, S. I. & LEVENDORSKIĬ, S. Z. (2002) *Non-Gaussian Merton-Black-Scholes Theory*, vol. 9. Advanced Series on Statistical Science & Applied Probability. World Scientific Publishing Company, xxii+398.
- BUCUR, C. (2016) Some observations on the Green function for the ball in the fractional Laplace framework. *Commun. Pure Appl. Anal.*, **15**, 657–699.
- BUCUR, C. & VALDINOCI, E. (2016) *Nonlocal Diffusion and Applications*, Vol. 20. Lecture Notes of the Unione Matematica Italiana. Berlin: Springer, xii+155.
- CHEN, Z.-Q. & SONG, R. (1998) Estimates on Green functions and Poisson kernels for symmetric stable processes. *Math. Ann.*, **312**, 465–501.
- COHEN, S., MEERSCHAERT, M. M. & ROSINSKI, J. (2010) Modeling and simulation with operator scaling. *Stochastic Process. Appl.*, **120**, 2390–2411.
- COHEN, S. & ROSINSKI, J. (2007) Gaussian approximation of multivariate Lévy processes with applications to simulation of tempered stable processes. *Bernoulli*, **13**, 195–210.
- CONT, R. & TANKOV, P. (2004) *Financial Modelling with Jump Processes*. Financial Mathematics Series. Chapman & Hall/CRC, xvi+535.
- DELAURENTIS, J. M. & ROMERO, L. A. (1990) A Monte Carlo method for Poisson’s equation. *J. Comput. Phys.*, **90**, 123–140.
- D’ELIA, M. & GUNZBURGER, M. (2016) Identification of the diffusion parameter in nonlocal steady diffusion problems. *Appl. Math. Optim.*, **73**, 227–249.
- DYBIEC, B. & SZCZEPANIEC, K. (2015) Escape from hypercube driven by multi-variate α -stable noises: role of independence. *Eur. Phys. J. B*, **88**, 8.
- DYDA, B. (2012) Fractional calculus for power functions and eigenvalues of the fractional Laplacian. *Fractional calculus and applied analysis*, **15**, 536–555.
- GETOOR, R. K. (1961) First passage times for symmetric stable processes in space. *Trans. Amer. Math. Soc.*, **101**, 75–90.
- GIVEN, J. A., HUBBARD, J. B. & DOUGLAS, J. F. (1997) A first-passage algorithm for the hydrodynamic friction and diffusion-limited reaction rate of macromolecules. *J. Chem. Phys.*, **106**, 3761–3771.
- GIVEN, J. A., HWANG, C.-O. & MASCAGNI, M. (2002) First- and last-passage Monte Carlo algorithms for the charge density distribution on a conducting surface. *Phys. Rev. E*, **66**, 056704.
- GIVEN, J. A., MASCAGNI, M. & HWANG, C.-O. (2001) Continuous path Brownian trajectories for diffusion Monte Carlo via first- and last-passage distributions. *Large-Scale Scientific Computing: Third International Conference, LSSC 2001 Sozopol, Bulgaria, June 6-10, 2001 Revised Papers* (S. Margenov, J. Wasniewski and P. Yalamov eds). Berlin: Springer, 46–57.
- HUANG, Y. & OBERMAN, A. (2014) Numerical methods for the fractional Laplacian: a finite difference-quadrature approach. *SIAM J. Numer. Anal.*, **52**, 3056–3084.
- HWANG, C.-O., GIVEN, J. A. & MASCAGNI, M. (2001) The simulation-tabulation method for classical diffusion Monte Carlo. *J. Comput. Phys.*, **174**, 925–946.
- HWANG, C.-O. & MASCAGNI, M. (2001) Efficient modified ‘walk on spheres’ algorithm for the linearized Poisson-Boltzmann equation. *Appl. Phys. Lett.*, **78**, 787–789.
- HWANG, C.-O., MASCAGNI, M. & GIVEN, J. A. (2003) A Feynman-Kac path-integral implementation for Poisson’s equation using an h-conditioned Green’s function. *Math. Comput. Simul.*, **62**, 347–355.
- JANICKI, A. & WERON, A. (1994) *Simulation and chaotic behavior of α -stable stochastic processes*, Vol. 178. Monographs and Textbooks in Pure and Applied Mathematics. Marcel Dekker, xii+355.

- KLAFTER, J., LIM, S. C. & METZLER, R. (2011) *Fractional Dynamics: Recent Advances*. World Scientific Publishing Company.
- KLAGES, R., RADONS, G. & SOKOLOV, I. M. (2008) *Anomalous Transport: Foundations and Applications*. Wiley.
- LANDKOF, N. S. (1972) *Foundations of Modern Potential Theory*. (Translated from the Russian by A. P. Doohovskoy, Die Grundlehren der mathematischen Wissenschaften, Band 180). New York-Heidelberg: Springer, x+424.
- MACKOY, T., HARRIS, R. C., JOHNSON, J., MASCAGNI, M. & FENLEY, M. O. (2013) Numerical optimization of a walk-on-spheres solver for the linear Poisson–Boltzmann equation. *Commun. Comput. Phys.*, **13**, 195–206.
- MIKHAILOV, G. A. (1979) Estimation of the difficulty of simulating the process of ‘random walk on spheres’ for some types of regions. *USSR Comput. Math. Math. Phys.*, **19**, 247–254.
- MOTOO, M. (1959) Some evaluations for continuous Monte Carlo method by using Brownian hitting process. *Ann. Inst. Stat. Math.*, **11**, 49–54.
- MULLER, M. E. (1956) Some Continuous Monte Carlo Methods for the Dirichlet Problem. *Ann. Math. Stat.*, **27**, 569–589.
- NOCHETTO, R. H., OTÁROLA, E. & SALGADO, A. J. (2016) A PDE approach to space-time fractional parabolic problems. *SIAM J. Numer. Anal.*, **54**, 848–873.
- ROS-OTON, X. (2016) Nonlocal elliptic equations in bounded domains: a survey. *Publ. Mat.*, **60**, 3–26.
- ROS-OTON, X. & SERRA, J. (2014) The Dirichlet problem for the fractional Laplacian: regularity up to the boundary. *J. Math. Pures Appl. (9)*, **101**, 275–302.
- SABELFELD, K. K. (1991) *Monte Carlo methods in Boundary Value Problems*. Springer Series in Computational Physics. Berlin: Springer.
- SABELFELD, K. K. & TALAY, D. (1995) Integral formulation of the boundary value problems and the method of random Walk on Spheres. *Monte Carlo Methods Appl.*, **1**, 1–34.
- SHLESINGER, M. F., ZASLAVSKY, G. M. & FRISCH, U. eds. (1995) Lévy Flights and Related Topics in Physics. Lecture Notes in Physics, vol. 450. Berlin: Springer, xvi+347.
- SZCZEPANIEC, K. & DYBIEC, B. (2015) Escape from bounded domains driven by multivariate α -stable noises. *J. Stat. Mech. Theory Exp.*, **6**, 16.
- ZOIA, A., ROSSO, A. & KARDAR, M. (2007) Fractional Laplacian in bounded domains. *Phys. Rev. E (3)*, **76**, 021116, 11.

Appendix. Proof of Theorem 6.1

Our proof of Theorem 6.1 uses heavily the joint conclusion of Theorems 2.10 and 3.2 in Bucur (2016), namely that the Theorem 6.1 is true in the case that D is a ball. Our proof is otherwise constructive proving existence and uniqueness separately.

Existence: On account of the fact that D is bounded, we can define a ball of sufficiently large radius $R > 0$, say $B^* = B(X_0, R)$, centred at X_0 , such that D is a subset of B^* and hence $\sigma_D \leq \sigma_{B^*}$ a.s., irrespective of the initial position of X , where $\sigma_{B^*} = \inf\{t > 0: X_t \notin B^*\}$. In particular, thanks to stationary and independent increments, this upper bound for σ_D does not depend on X_0 in law and $\sup_{x \in D} \mathbb{E}_x[\sigma_D] \leq \mathbb{E}_0[\sigma_{B^*}] < \infty$.

Define for convenience $v(x) = g(x)$ for $x \in D^c$ and

$$v(x) = \mathbb{E}_x[g(X_{\sigma_D})] + \mathbb{E}_x \left[\int_0^{\sigma_D} f(X_s) ds \right], \quad x \in D, \quad (\text{A.1})$$

where g and f satisfy the assumptions of the theorem. We want to prove that v is bounded and continuous on \bar{D} . For the boundedness of v , we prove the boundedness of the two expectations in its definition.

First note that, for all $x \in D$,

$$\begin{aligned}
\mathbb{E}_x [|g(X_{\sigma_D})|] &= \mathbb{E}_x [|g(X_{\sigma_D})| \mathbf{1}_{(\sigma_D = \sigma_{B^*})}] + \mathbb{E}_x [|g(X_{\sigma_D})| \mathbf{1}_{(\sigma_D < \sigma_{B^*})}] \\
&\leq \mathbb{E}_x [|g(X_{\sigma_{B^*}})|] + \sup_{x \in B^* \setminus D} |g(x)| \\
&= \mathbb{E}_0 [|g(x + B^* X_{\sigma_{B(0,1)}})|] + \sup_{x \in B^* \setminus D} |g(x)| \\
&= \pi^{-(d/2+1)} \Gamma(d/2) \sin(\pi\alpha/2) \int_{|y|>1} \frac{|g(x + B^*y)|}{|1 - |y|^2|^{\alpha/2} |y|^d} dy + \sup_{x \in B^* \setminus D} |g(x)| \\
&= C \int_{\mathbb{R}^d} \frac{|g(z)|}{1 + |z|^{d+\alpha}} dy + \sup_{x \in B^* \setminus D} |g(x)| < \infty, \tag{A.2}
\end{aligned}$$

for some constant $C \in (0, \infty)$ that does not depend on x (this is ensured, thanks to the boundedness of D). In the inequality, we have used the fact that, on $\{\sigma_D < \sigma_{B^*}\}$, we have $X_{\sigma_D} \in B^* \setminus D$; moreover, that, as a continuous function on \mathbb{R}^d , g is bounded in $B^* \setminus D$. In the second equality, we have used spatial homogeneity and the scaling property of stable processes. In the third equality, we have used Theorem 3.1. The fourth equality follows by changing variables to $z = x + B^*y$ in the integral appropriately estimating the denominator and the assumption that g is continuous and in $L_\alpha^1(D^c)$.

The boundedness of f on D and the uniform finite mean of σ_D ensures that the second expectation in the definition of v is bounded on \bar{D} . We claim that v is continuous in \mathbb{R}^d and belongs to $L_\alpha^1(\mathbb{R}^d)$. Continuity of v follows, thanks to the path regularity of X , the continuity of g , the openness of D and the fact that $\omega \mapsto X_{\sigma_D}(\omega)$ and $\omega \mapsto \int_0^{\sigma_D(\omega)} f(X_s(\omega)) ds$ are continuous in the Skorohod topology (for which it is important that $\omega \mapsto \sigma_D(\omega)$ is finite). Continuity is also a consequence of the classical potential analytic point of view, seeing the identity for v in (A.1) in terms of Riesz potentials (see, for example, the classical texts of [Landkof, 1972](#); [Bliedner & Hansen, 1986](#)).

To check that $v \in L_\alpha^1(\mathbb{R}^d)$, we need some estimates. For $x \in D^c$, $v(x) = g(x)$ and hence, as $g \in L_\alpha^1(D^c)$, it suffices to check that $\int_D |v(x)|/(1 + |x|^{\alpha+d}) dx < \infty$. However, this is trivial on account of the boundedness and continuity of v on \bar{D} .

Now fix $x' \in D$ and let $B(x')$ be the largest ball centred at x' that is contained in D . A simple application of the strong Markov property tells us that

$$\begin{aligned}
v(x) &= \mathbb{E}_x \left[\mathbb{E} \left[g(X_{\sigma_D}) + \int_0^{\sigma_{B(x')}} f(X_s) ds + \int_{\sigma_{B(x')}}^{\sigma_D} f(X_s) ds \middle| \mathcal{F}_{\sigma_{B(x')}} \right] \right] \\
&= \mathbb{E}_x \left[v(X_{\sigma_{B(x')}}) + \int_0^{\sigma_{B(x')}} f(X_s) ds \right], \quad x \in D, \tag{A.3}
\end{aligned}$$

where $(\mathcal{F}_t, t \geq 0)$ is the natural filtration generated by X . Thanks to the fact that Theorem 6.1 is valid on balls, we see immediately that the right-hand side of (A.3) is the unique solution to

$$\begin{aligned}
-(-\Delta u)^{\alpha/2} u(x) &= -f(x), & x \in B(x'), \\
u(x) &= v(x), & x \in B(x')^c. \tag{A.4}
\end{aligned}$$

That is to say, v solves (A.4). Note that it is at this point in the argument that we are using the condition $f \in C^{\alpha+\varepsilon}(\overline{D})$. Since the solution to (A.4) is defined on $B(x')$ and x' is chosen arbitrarily in D , we conclude that v solves

$$\begin{aligned} -(-\Delta u)^{\alpha/2} u(x) &= -f(x), & x \in D, \\ u(x) &= v(x), & x \in D^c. \end{aligned} \tag{A.5}$$

On account of the fact that $\mathbb{P}_x(\sigma_D = 0) = 1$ for all $x \in D^c$, it follows that $v = g$ on D^c and hence (A.5) is identical to (6.1).

Uniqueness: Suppose that \hat{u} solves (6.1), then, in particular, for any $x' \in D$, it must solve

$$\begin{aligned} -(-\Delta u)^{\alpha/2} u(x) &= -f(x), & x \in B(x'), \\ u(x) &= \hat{u}(x), & x \in B(x')^c. \end{aligned}$$

As we know the Feynman–Kac representation of the solution to the above fractional Poisson problem, thanks to Theorem 3.2 in Bucur (2016) for domains that are balls, we are forced to conclude that

$$\hat{u}(x) = \mathbb{E}_x \left[\hat{u}(X_{\sigma_{B(x')}}) + \int_0^{\sigma_{B(x')}} f(X_s) ds \right], \quad x \in B(x'), \quad x' \in D. \tag{A.6}$$

Here again, we are implicitly using that $f \in C^{\alpha+\varepsilon}(\overline{D})$ in the application of Theorem 3.2 of Bucur (2016). Let us now appeal to the same notation we have used for the walk-on-spheres. Specifically, recall the sequential exit times from maximally sized balls σ_{B_k} for the walk-on-spheres that were defined in Section 4. We claim that

$$M_k =: \hat{u}(X_{\sigma_{B_k} \wedge \sigma_D}) + \int_0^{\sigma_{B_k} \wedge \sigma_D} f(X_s) ds, \quad k \geq 0,$$

is a martingale. To see why, note that, by the strong Markov property and then by (A.6),

$$\begin{aligned} \mathbb{E}[M_{k+1} | \mathcal{G}_k] &= \mathbf{1}_{\{k < N\}} \left\{ \mathbb{E}_x \left[\hat{u}(X_{\sigma_{B(x)}}) + \int_0^{\sigma_{B(x)}} f(X_s) ds \right] \Big|_{x=X_{\sigma_{B_k}}} + \int_0^{\sigma_{B_k}} f(X_s) ds \right\} \\ &\quad + \mathbf{1}_{\{k \geq N\}} \left\{ \hat{u}(X_{\sigma_D}) + \int_0^{\sigma_D} f(X_s) ds \right\} \\ &= \mathbf{1}_{\{k < N\}} \left\{ \hat{u}(X_{\sigma_{B_k}}) + \int_0^{\sigma_{B_k}} f(X_s) ds \right\} + \mathbf{1}_{\{k \geq N\}} \left\{ \hat{u}(X_{\sigma_D}) + \int_0^{\sigma_D} f(X_s) ds \right\} \\ &= \hat{u}(X_{\sigma_{B_k} \wedge \sigma_D}) + \int_0^{\sigma_{B_k} \wedge \sigma_D} f(X_s) ds \\ &= M_k, \quad k \geq 1, \end{aligned}$$

where $\mathcal{G}_k = \mathcal{F}_{\sigma_{B_k} \wedge \sigma_D}$, $k \geq 1$. For consistency, we may define $M_0 = \mathbb{E}_x[M_k] = \hat{u}(x)$, thanks to (A.6).

Next, we appeal to the definition of B^* and, in particular, that $\sigma_D \leq \sigma_{B^*}$, as well as the continuity of \hat{u} to deduce that, for all $k \geq 0$,

$$\begin{aligned} & \left| \hat{u}(X_{\sigma_{B_k} \wedge \sigma_D}) + \int_0^{\sigma_{B_k} \wedge \sigma_D} f(X_s) \, ds \right| \\ & \leq \left| \hat{u}(X_{\sigma_{B^*}}) \right| \mathbf{1}_{\{\sigma_{B_k} \wedge \sigma_D = \sigma_{B^*}\}} + \sup_{y \in B^*} \left| \hat{u}(y) \right| \mathbf{1}_{\{\sigma_{B_k} \wedge \sigma_D < \sigma_{B^*}\}} + \sup_{y \in D} |f(y)| \sigma_D \\ & \leq \left| g(X_{\sigma_{B^*}}) \right| + c_1 + c_2 \sigma_{B^*}, \end{aligned}$$

where c_1, c_2 are constants. We know that for each fixed $x \in D$, $\mathbb{E}_x[\sigma_{B^*}] < \infty$ and, moreover, from Theorem 3.1, after scaling (see, for example, (7.1)), $\mathbb{E}_x[|g(X_{\sigma_{B^*}})|] < \infty$ as $g \in L^1_\alpha(D^c)$. Dominated convergence allows us to deduce that $(M_k, k \geq 0)$ is a uniformly integrable martingale, such that, for each fixed $x \in D$,

$$\begin{aligned} \hat{u}(x) &= \lim_{k \rightarrow \infty} \mathbb{E}_x[M_k] \\ &= \mathbb{E}_x[\lim_{k \rightarrow \infty} M_k] \\ &= \mathbb{E}_x \left[\hat{u}(X_{\sigma_D}) + \int_0^{\sigma_D} f(X_s) \, ds \right] \\ &= \mathbb{E}_x \left[g(X_{\sigma_D}) + \int_0^{\sigma_D} f(X_s) \, ds \right], \end{aligned}$$

where in the final equality we have used that $\hat{u} = g$ on D^c . Uniqueness now follows. \square