

The complexity of simulating Brownian Motion

Ilia Binder*

Mark Braverman†

Abstract

We analyze the complexity of the Walk on Spheres algorithm for simulating Brownian Motion in a domain $\Omega \subset \mathbb{R}^d$. The algorithm, which was first proposed in the 1950s, produces samples from the hitting probability distribution of the Brownian Motion process on $\partial\Omega$ within an error of ε . The algorithm is used as a building block for solving a variety of differential equations, including the Dirichlet Problem.

The WoS algorithm simulates a BM starting at a point $X_0 = x$ in a given bounded domain Ω until it gets ε -close to the boundary $\partial\Omega$. At every step, the algorithm measures the distance d_k from its current position X_k to $\partial\Omega$ and jumps a distance of $d_k/2$ in a uniformly random direction from X_k to obtain X_{k+1} . The algorithm terminates when it reaches X_n that is ε -close to $\partial\Omega$.

It is not hard to see that the algorithm requires at least $\Omega(\log 1/\varepsilon)$ steps to converge. Only partial results with respect to the upper bound existed. In 1959 M. Motoo established an $O(\log 1/\varepsilon)$ bound on the running time for convex domains. The results were later generalized for a wider, but still very restricted, class of planar and 3-dimensional domains by G.A. Mikhailov (1979). In our earlier work (2007), we established an upper bound of $O(\log^2 1/\varepsilon)$ on the rate of convergence of WoS for arbitrary planar domains.

In this paper we introduce energy functions using Newton potentials to obtain very general upper bounds on the convergence of the algorithm. Special instances of the upper bounds yield the following results for bounded domains Ω :

- if Ω is a planar domain with connected exterior, the WoS converges in $O(\log 1/\varepsilon)$ steps;
- if Ω is a domain in \mathbb{R}^3 with connected exterior, the WoS converges in $O(\log^2 1/\varepsilon)$ steps;
- for $d > 2$, if Ω is a domain in \mathbb{R}^d , the WoS converges in $O((1/\varepsilon)^{2-4/d})$ steps;

- for $d > 3$, if Ω is a domain in \mathbb{R}^d with connected exterior, the WoS converges in $O((1/\varepsilon)^{2-4/(d-1)})$ steps;
- for any d , if Ω is a domain in \mathbb{R}^d bounded by a smooth surface $\partial\Omega$, the WoS converges in $O(\log 1/\varepsilon)$ steps.

We also demonstrate that the bounds are tight, i.e. we construct a domain from each class for which the upper bound is exact. Our results give the optimal upper bound of $O(\log 1/\varepsilon)$ in many cases for which only a bound polynomial in $1/\varepsilon$ was previously known.

1 Introduction

Brownian Motion (BM) is the most important model of randomized motion in \mathbb{R}^d . It is the simplest (but, in a sense, generic) example of a continuous diffusion process. BM has found an astonishing number of application to diverse areas of Mathematics and Science, including Biomathematics, Finance, Partial Differential Equations, and Statistical Physics [5, 8, 10, 15, 16].

Because of the ubiquity of BM, its effective simulation provides a way to efficiently solve a variety of problems, such as computation of Conformal Maps, Tomography, and Stochastic PDEs. One of the main ways in which simulations of BM are used is to study its first hitting probabilities with respect to some stopping conditions.

A particularly interesting stopping condition is of hitting the boundary of some topologically open bounded connected domain Ω . For any starting point x the *harmonic measure* h_x on $\partial\Omega$ is given by the hitting point distribution. In many of the BM's applications, it is enough to obtain information about the harmonic measure, more specifically, to efficiently sample from it.

One of the immediate applications of the ability to sample from harmonic measures is solving the Dirichlet problem in \mathbb{R}^d . The Dirichlet problem on a domain $\Omega \subset \mathbb{R}^d$ with continuous boundary condition $f : \partial\Omega \rightarrow \mathbb{R}$ is the problem of finding a C^2 -smooth function $u : \Omega \rightarrow \mathbb{R}$ continuous up to the boundary $\partial\Omega$ satisfying

$$(1.1) \quad \begin{cases} \Delta u(x) = 0 & x \in \Omega \\ u(x) = f(x) & x \in \partial\Omega \end{cases}$$

In other words, finding a harmonic function u subject to the boundary conditions f . By the celebrated

*Department of Mathematics, University of Toronto. Partially supported by NSERC Discovery grant 5810-2004-298433

†Microsoft Research. This research was partially conducted during the period the author was a PhD student at the University of Toronto and partially while the author was employed by the Clay Mathematics Institute as a Liffott Fellow.

Kakutani’s Theorem [7, 6], the value of u at $x \in \Omega$ is exactly the expected value of f with respect to the harmonic measure h_x on $\partial\Omega$: $u(x) = \mathbf{E}_{h_x(z)}[f(z)]$.

In the present paper we study the amount of time it takes to sample from the harmonic measure with precision ε using the Walk on Spheres algorithm – the simplest and most commonly used algorithm for sampling from the harmonic measure.

1.1 The Walk on Spheres algorithm The Walk on Spheres (WoS) algorithm was first proposed in 1956 by M. Muller in [14]. In his paper, the method was applied to the solution of various boundary problems for the Laplace operator, including the Dirichlet problem. Logarithmic running time of the process X_t was established for convex domains by M. Motoo in [13] and was later generalized for a wider, but still very restricted, class of planar and 3-dimensional domains by G.A. Mikhailov in [11]. See also [4] and [12] for additional historical background and the use of the algorithm for solving other types of boundary value problems.

In our earlier work [1], we established polylogarithmic, but not logarithmic, upper bounds on the rate of convergence of WoS for planar domains, and for a restricted class of higher-dimensional domains. Unfortunately, the techniques of [1] do not generalize well to general domains in higher dimensions.

Our present results subsume all prior work on the rate of convergence of the WoS. We introduce an easily verified geometric condition on the domain which provides tight bounds for the rates of convergence. In particular, the optimal logarithmic rate of convergence is established for all domains whose boundary is a smooth hypersurface.

Let us now define the WoS. We would like to simulate a BM in a given bounded domain Ω until it gets ε -close to the boundary $\partial\Omega$. Of course one could simulate it using jumps of size δ in a random direction on each step, but this would require $O(1/\delta^2)$ steps. Since we must take $\delta = O(\varepsilon)$, this would also mean that the process may take $O(1/\varepsilon^2)$ steps to converge.

The idea of the WoS algorithm is very simple: we do not care about the path the BM takes, but only about the point at which it hits the boundary. Thus if we are currently at a point $X_n \in \Omega$ and we know that

$$d(X_n) := d(X_n, \partial\Omega) \geq r,$$

i.e. that X_n is at least r -away from the boundary, then we can just jump $r/2$ units in a random direction from X_n to a point X_{n+1} . To justify the jump we observe that a BM hitting the boundary would have to cross

the sphere

$$S_n = \{x : |x - X_n| = r/2\}$$

at some point, and the first crossing location X_{n+1} is distributed uniformly on the sphere. There is nothing special about a jump of $d(X_n)/2$ and it can be replaced with any $\beta d(X_n)$ where $0 < \beta < 1$.

Let $\{\gamma_n\}$ be a sequence of i.i.d. random variables each being a vector uniformly distributed on the unit sphere in \mathbb{R}^d . We could take, for example, $\gamma_n = \Gamma_n^d / |\Gamma_n^d|$, where Γ_n^d is a normally distributed d -dimensional Gaussian variable. Then, schematically, the Walk on Spheres algorithm can be presented as follows:

```

WalkOnSpheres( $X_0, \varepsilon$ )
 $n := 0$ ;
while  $d(X_n) = d(X_n, \partial\Omega) > \varepsilon$  do
    compute  $r_n$ : a multiplicative estimate on  $d(X_n)$ 
    such that  $\beta \cdot d(X_n) < r_n < d(X_n)$ ;
     $X_{n+1} := X_n + (r_n/2) \cdot \gamma_n$ ;
     $n := n + 1$ ;
endwhile
return  $X_n$ 

```

Thus at each step of the algorithm we jump at least $\beta/2$ and at most $1/2$ -fraction of the distance to the boundary in a random direction. An example of running the WoS algorithm in 2-d is illustrated on Figure 1.

As mentioned earlier, it is clear that the algorithm is correct. Moreover, it is not hard to see that it converges in $O(1/\varepsilon^2)$ steps. However, in many situations, this rate of convergence is unsatisfactory. In particular, if we wanted to get 2^{-n} -close to the boundary, it would take us a number of steps *exponential* in n . As it turns out, in many natural situations, the rate of convergence is *polynomial* or even *linear* in n (i.e. logarithmic in $1/\varepsilon$). The object of the paper is to prove that this is the case, and give precise condition on when the faster convergence occurs.

While an actual implementation of the WoS would involve round-off errors introduced through an imperfect simulation, we will ignore those to simplify the presentation as they do not affect any of the main results. Thus the problem becomes purely that of analyzing the family of stochastic processes $\{X_t\}$ and their convergence speed to $\partial\Omega$.

Providing the domain Ω to the algorithm It is worth noting that the algorithm needs access to the input domain Ω in a very weak sense. We need an oracle

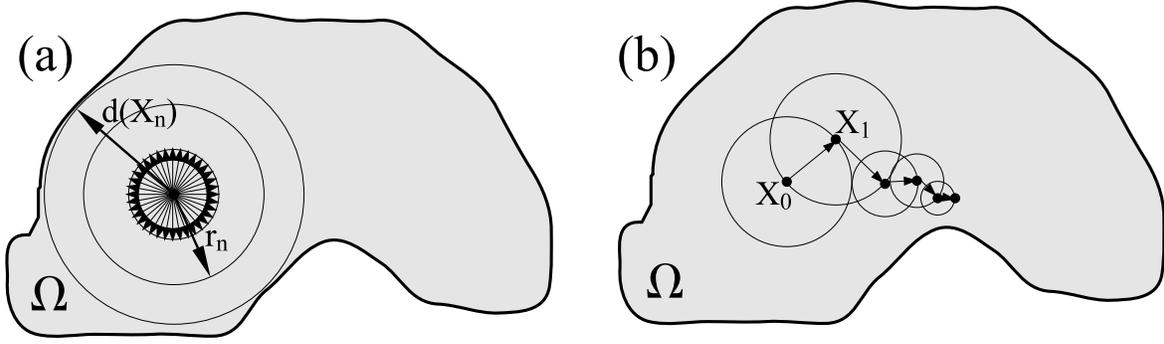


Figure 1: An illustration of the WoS algorithm for $d = 2$: one step jump (a), and a possible run of the algorithm for several steps (b)

$\text{dist}_\Omega(x)$ that satisfy the following:

$$(1.2) \quad \text{dist}_\Omega(x) \in \begin{cases} (\beta d(x), d(x)) & \text{if } x \in \Omega, d(x) > \beta\varepsilon \\ [0, \beta\varepsilon) & \text{if } x \in \Omega, d(x) \leq \beta\varepsilon \\ 0 & \text{if } x \notin \Omega \end{cases}$$

for some $0 < \beta < 1$. Note that dist_Ω would also allow us to decide both the size of the jump on step n and whether X_n is sufficiently close to $\partial\Omega$ for the algorithm to terminate.

If Ω is given to the algorithm as a union of squares on a ε -fine grid, then dist_Ω can be computed in time $\text{poly}(1/\varepsilon)$. In many applications, however, this function can be computed in time $\text{poly}(\log 1/\varepsilon)$, because we only need to estimate the distance within a *multiplicative* error of β . The precise condition for this is that the complement set Ω^c is poly-time computable as a subset of \mathbb{R}^d in the sense of Computable Analysis. See for example [2, 17, 3] for more details on poly-time computability of real sets. The vast majority of domains in applications satisfy this condition.

Thus, in cases when the domain Ω is sufficiently nice for Ω^c to be poly-time computable, the rate of convergence of the WoS becomes the crucial component in the running time of its execution. In particular, depending on whether the rate of convergence is $\text{poly}(1/\varepsilon)$ or $\text{poly}(\log 1/\varepsilon)$ it could take time that is either exponential or polynomial in n to sample points that are 2^{-n} -away from $\partial\Omega$.

1.2 Results To state the results precisely we will need a somewhat technical notion of α -thickness. We will see that this condition is naturally satisfied for many interesting classes of domains. Intuitively, a d -dimensional domain is α -thick, if any neighborhood of any point x on the boundary $\partial\Omega$ contains a $(d - \alpha)$ -dimensional set $A_x \ni x$ in the complement of Ω . Thus

the “nicer” $\partial\Omega$ is the smaller is α . As we will see, smaller α indeed translates directly to faster convergence of the WoS. The formal definition of α -thickness involves measures:

DEFINITION 1. A domain $\Omega \subset \mathbb{R}^d$ is said to be α -thick $0 \leq \alpha \leq d$ if there exists a constant $c > 0$ such that for every $x \in \partial\Omega$ there is a Borel measure μ_x which satisfies the following conditions:

1. $\text{supp}(\mu_x) \cap \Omega = \emptyset$, or equivalently $\mu_x(\Omega) = 0$;
2. for any $y \in \mathbb{R}^d$ and $r > 0$, $\mu_x(B(y, r)) \leq r^{d-\alpha}$;
3. for any $1 > r > 0$, $\mu_x(B(x, r)) \geq c \cdot r^{d-\alpha}$.

We call the constant c the *thickness* of the domain Ω . It is not hard to see that the property of α -thickness is monotone: an α -thick domain is α' -thick for $\alpha < \alpha' \leq d$. The class of α -thick domains is very rich, even for specific constant values of α . In particular we have the following interesting special cases.

CLAIM 1.

1. All d -dimensional domains are d -thick;
2. all bounded d -dimensional domains Ω such that the complement Ω^c is connected are $d - 1$ -thick.
3. all convex domains are 0-thick;
4. all domains Ω that are bounded by a smooth hypersurface $\partial\Omega$ are 0-thick.

Proof. The first statement simply follows by placing a δ -measure $\mu_x(\{x\}) = 1$ at x .

To prove the second statement, consider a measurable function $f_x : [0, \infty) \rightarrow \Omega^c$ such that $|x - f_x(r)| =$

r. Existence of such a function follows from the connectedness of $\Omega^c \ni x$ by a standard topological argument. Define the measure μ_x as

$$\mu_x(B) = \frac{1}{2}m_1(f_x^{-1}(B)),$$

where m_1 is the standard Lebesgue measure on $[0, \infty)$. Let $y \in \mathbb{R}^d$ with $|y - x| = a$ and $r > 0$, then

$$f_x^{-1}(B(y, r)) \subset [a - r, a + r],$$

and hence $\mu_x(B(y, r)) \leq r$. On the other hand, for each $r > 0$,

$$\mu_x(B(x, r)) = \frac{1}{2}m_1(f_x^{-1}(B(x, r))) = \frac{1}{2}m_1([0, r]) = \frac{1}{2}r.$$

To prove the third and the fourth claim, we just take μ_x to be a scaled d -dimensional Lebesgue measure restricted to the complement of Ω .

We are now ready to state the main theorem.

THEOREM 1.1. *Let Ω be an open, bounded α -thick domain in \mathbb{R}^d . Then the expected rate of convergence of the WoS from any $x \in \Omega$ until termination at distance $< \varepsilon$ to the boundary is given by the following table:*

	Rate of convergence
$\alpha < 2$	$O(\log 1/\varepsilon)$
$\alpha = 2$	$O(\log^2 1/\varepsilon)$
$\alpha > 2$	$O((1/\varepsilon)^{2-4/\alpha})$

The $O(\cdot)$ in the expressions above depends on the dimension d , on α , on the thickness constant c from Definition 1 and on $\beta > 0$ from the definition of the WoS. It does not depend directly on Ω .

Moreover, the rates of convergence above are tight. That is, for each α there is a family of α -thick domains Ω_n^α with some thickness c , such that the rate of convergence with $\varepsilon = 1/n$ on Ω_n^α is asymptotically given by the formulas in (1.3).

The rate of convergence cannot be better than $O(\log 1/\varepsilon)$ since at each step of the WoS, the distance of X_t to the boundary $\partial\Omega$ decreases by at most a multiplicative constant. An intuitive explanation to the phase transition phenomenon occurring at $\alpha = 2$, is that a BM in \mathbb{R}^d almost surely “misses” sets of co-dimension > 2 , while hitting sets of co-dimension ≤ 2 with positive probability.

It is worth noting that the main result in [1] is the special case $\alpha = d = 2$ of the theorem.

The following corollaries are implied directly by the main theorem and Claim 1.

COROLLARY 1.1.

1. For any planar domain, the WoS converges in $O(\log^2 1/\varepsilon)$ steps;
2. for any planar domain with connected exterior, the WoS converges in $O(\log 1/\varepsilon)$ steps;
3. for any $d \geq 3$, the WoS converges in $O((1/\varepsilon)^{2-4/d})$ steps;
4. for any 3-dimensional domain with connected exterior, the WoS converges in $O(\log^2 1/\varepsilon)$ steps;
5. for any $d \geq 4$, and for any d -dimensional domain with connected exterior, the WoS converges in $O((1/\varepsilon)^{2-4/(d-1)})$ steps;
6. for any domain bounded by a smooth hypersurface, the WoS converges in $O(\log 1/\varepsilon)$ steps.

2 Upper bounds: energy function

2.1 Energy Function of optimal growth The heart of the proof of the upper bounds in Theorem 1.1 is the construction of a subharmonic function with optimal growth at the boundary, the *Energy Function* U on Ω . We will construct $U(x)$ so that it is “small” in the interior of Ω , and grows to ∞ as x approaches the boundary $\partial\Omega$. The α -thickness of the domain allows us to establish that the value of $U(X_t)$ grows in expectation as the WoS progresses. Thus after a certain number of steps $U(X_t)$ will be large in expectation which would imply that X_t is close to $\partial\Omega$ with high probability.

The construction of the function is based on the notion of a *Riesz potential*. For a finite Borel measure μ on \mathbb{R}^d , and $\alpha < d$, the α -Riesz potential of the measure μ is defined by

$$U_\alpha^\mu(x) = \frac{1}{d - \alpha} \int \frac{d\mu(z)}{\|z - x\|^{d-\alpha}}.$$

For $\alpha = d$, the d -Riesz potential is defined by

$$U_\alpha^\mu(x) = \int \log \frac{1}{\|z - x\|} d\mu(z).$$

The value $U_\alpha^\mu(x) = \infty$ is allowed when the integral diverges.

An important special case is the case of $\alpha = 2$, the so-called *Newton potential*. We will denote U_2^μ simply by U^μ . In this case the expression under the integral is harmonic in \mathbb{R}^d . It is well known (e.g. see [9]) that the function U^μ is superharmonic on \mathbb{R}^d , and harmonic outside of $\text{supp } \mu$.

More generally, for $0 < \alpha < 2$, the function U_α^μ is subharmonic outside of $\text{supp } \mu$.

The following important technical identity, which easily follows from Fubini's Theorem and substitution, relates the local behavior of the measure μ and the growth of its potential U_α^μ . For $\alpha < d$, we have

$$(2.4) \quad U_\alpha^\mu(y) = \frac{1}{d-\alpha} \int_0^\infty \mu(B(y, t^{-1/(d-\alpha)})) dt = \int_0^\infty \frac{\mu(B(y, r))}{r^{d-\alpha+1}} dr,$$

and for $\alpha = d$,

$$(2.5) \quad U_\alpha^\mu(y) = \int_{-\infty}^\infty \mu(B(y, e^{-t})) dt = \int_0^\infty \frac{\mu(B(y, r))}{r} dr$$

Let us now fix an α -thick domain $\Omega \subset B(0, 1) \subset \mathbb{R}^d$. Let us consider the set \mathcal{M} of all Borel measures μ supported inside $\overline{B(0, 2)}$ and outside of Ω (i.e. $\mu(\Omega) = 0$), satisfying the following condition:

$$(2.6) \quad \text{for any } y \in \mathbb{R}^d \text{ and } r > 0, \mu(B(y, r)) \leq r^{d-\alpha}$$

Let us now introduce the *Energy Function* $U(y)$. Recall that $U^\mu(y) := U_2^\mu(y)$.

$$(2.7) \quad U(y) := \begin{cases} \sup_{\mu \in \mathcal{M}} U_\alpha^\mu(y), & \text{when } \alpha \leq 2 \\ \sup_{\mu \in \mathcal{M}} U^\mu(y), & \text{when } \alpha \geq 2. \end{cases}$$

Let us summarize the properties of $U(y)$ in the following claim, proved in the full version of the paper. The proof makes use of identities (2.4) and (2.5). Recall that $d(y) = \text{dist}(y, \partial\Omega)$.

CLAIM 2. *Let Ω be an α -thick domain. Then*

1. $U(y)$ is subharmonic in Ω .
2. For $\alpha \leq 2$, $U(y) \leq \log \frac{2}{d(y)}$ for all $y \in \Omega$.
3. For $\alpha > 2$, $U(y) \leq \frac{1}{\alpha-2} d(y)^{2-\alpha}$ for all $y \in \Omega$.

Let X_t be the WoS process initiated at some point $X_0 = y \in \Omega$. Let us define a new process $U_t = U(X_t)$, the value of the energy function at the t -th step of the process. Note that because U is subharmonic, U_t is a submartingale, that is $\mathbf{E}[U_{t+1}|U_t] \geq U_t$.

For the rest of the section let $n = 1/\varepsilon$. Claim 2 immediately implies that a large value of U_t will guarantee the closeness to the boundary. More specifically,

CLAIM 3. *For $\alpha \leq 2$, if $U_t > \log 2n$ then $d(X_t) < 1/n$. For $\alpha > 2$, $U_t > (\alpha-2)n^{\alpha-2}$ implies $d(X_t) < 1/n$.*

The proof of Theorem 1.1 relies on finer lower bounds on the function U , which would guarantee the optimal rate of boundary convergence. These bounds depend heavily on the value of α . We outline the ideas of the proof in the next sections, with the full details deferred to the full version of the paper.

2.2 Logarithmic convergence: the case $\alpha < 2$.

At the heart of the proof for this case lies the following strong estimate on the behavior of the Riesz potentials near the boundary.

LEMMA 2.1. *For any $\alpha < 2$ and $c > 0$, there exists a constant δ , such that the following holds.*

Let Ω be an α -thick domain in \mathbb{R}^d with thickness c . Let $y \in \Omega$ and $x \in \partial\Omega$ be the closest point to y . Let $\mu \in \mathcal{M}$ (recall that \mathcal{M} is the class of measures defined in the previous section).

Then either

$$(2.8) \quad U(z) > U_\alpha^\mu(z) + 1 \text{ whenever } \delta/4 \cdot d(y) < |z-x| < \delta \cdot d(y).$$

or

$$(2.9) \quad \mu(B(y, 2d(y))) \geq \delta d(y)^{d-\alpha}$$

The lemma is established in the full version of the paper. Note that after $k = O(|\log \delta|)$ steps of the WoS process,

$$(2.10) \quad \delta/4 \cdot d(X_t) < |X_{t+k} - x| < \delta \cdot d(X_t)$$

with a certain probability p ,

where x is the point of $\partial\Omega$ that is closest to X_t , and $p > 0$ depends only on β and the dimension d .

Let us fix X_t and take the measure $\mu \in \mathcal{M}$ maximizing the value of $U_\alpha^\mu(X_t)$. Such a measure exists by a compactness argument. By the preceding observation, in the first case in Lemma 2.1, the subharmonicity of U implies that the expectation of U_{t+k} , conditioned on U_t , will increase by some definite constant.

On the other hand, using the α -thickness of Ω , one can see that the Laplacian of U_α^μ is large near the point X_t in the second case of Lemma 2.1. Thus, since large Laplacian leads to a fast build-up of mean values, we have the above-mentioned increase by a constant after the first step. Thus we arrive to the following estimate, which shows that U_t grows at least linearly in expectation.

LEMMA 2.2. *There are constants L and k , depending only on c , β , and α , such that*

$$\mathbf{E}[(U_{t+k} - U_t)|U_t] > L.$$

A detailed proof of the lemma can be found in the full version of the paper.

Lemma 2.2 implies that $\mathbf{E}[U_t] > tL/k + U_0$. Since $d(X_t) \geq (1 - \beta)^t d(X_0)$, Claim 3 implies that $U_t \leq U_0 + t|\log(1 - \beta)| + \log 2$. This implies that $U_t > U_0 + tL/2k$ with probability at least P , where P depends only on β . This, together with Claim 3 implies the necessary upper bound in the case $\alpha < 2$.

2.3 Polylogarithmic convergence: the case $\alpha = 2$. In the case $\alpha = 2$ the steady linear growth of U_t given by Lemma 2.2 no longer holds. In fact, the only thing that generally holds in this case is the submartingale property $\mathbf{E}[U_{t+1}|U_t] \geq U_t$. We are able to overcome this difficulty, by showing that the submartingale process $\{U_t\}$ has a deviation bounded from below by a constant at every step. To this end it suffices to show that U_t can grow by some η with a non-negligible probability. We use the following estimate on the energy function (established in the full version of the paper).

LEMMA 2.3. *There exists a constant δ , dependent only on the thickness c , such that the following holds. Let Ω be a 2-thick domain. Let $y \in \Omega$ and $x \in \partial\Omega$ be the closest point to y . Then*

$$(2.11) \quad U(z) > U(y) + 1 \text{ whenever } |z - x| < \delta \cdot d(y).$$

Since the function U is subharmonic, observation (2.10) implies the following estimate (we defer to the full version of the paper for a detailed proof).

LEMMA 2.4. *Let Ω be a 2-thick domain in \mathbb{R}^d . There are constants k and L , depending only on the thickness c , the jump ratio β , and the dimension d , such that*

$$\mathbf{E}[(U_{t+k} - U_t)^2 | U_t] > L.$$

The derivation of the upper bound on the rate of convergence from the Lemma is pretty standard, and is established in the full version of the paper.

2.4 Polynomial convergence: the case $\alpha > 2$. For the case $\alpha > 2$, the required converse to Claim 3 is very simple.

LEMMA 2.5. *For $\alpha > 2$, and an α -thick domain Ω in \mathbb{R}^d with the thickness c ,*

$$U(y) \geq K \cdot d(y)^{2-\alpha}$$

for all $y \in \Omega$. Here the constant $K = K(c, \alpha)$ depends only on c and α .

Proof. Let x be the closest to y point at $\partial\Omega$, and let $\mu = \mu_x$ be the corresponding measure from the definition of the α -thick domains. Then, by the identity

(2.4) and since $B(x, r) \subset B(y, r + d(y))$,

$$(2.12) \quad U(y) \geq U_2^\mu(y) = \int_{d(y)}^2 \frac{\mu(B(y, r))}{r^{d-1}} dr \geq c2^{\alpha-d} \int_{2d(y)}^2 t^{1-\alpha} \geq K \cdot d(y)^{2-\alpha}.$$

The idea of the proof of Theorem 1.1 in this case is now as follows. When the WoS is far from the boundary $\partial\Omega$ it makes fairly big steps and when it is close it makes small steps. There are not too many big steps because the number of big steps of length $> \varepsilon$ confined to $B(0, 1)$ is bounded by $O(1/\varepsilon^2)$. On the other hand, there are not too many small steps, because a small step means that the WoS is very close to $\partial\Omega$, and should converge before an opportunity to make many more steps. More precisely, the number of “big” jumps is bounded by the following claim (we defer to the full version of the paper for a proof).

CLAIM 4. *Let $N(\varepsilon, T)$ be the number of the jumps in the WoS process before the time t which are bigger than ε , i.e.*

$$N(\varepsilon, T) = \#\{t \leq T \mid |X_t - X_{t-1}| \geq \varepsilon\}.$$

Then

$$\mathbf{P} \left[N(\varepsilon, T) > \frac{4}{\varepsilon^2} \right] < 1/4.$$

To bound the number of small jumps, we denote by $R_0 \subset \Omega$ the $1/n$ -neighborhood of $\partial\Omega$, and more generally, by

$$R_k := \{x \in \Omega : 2^{k-1}/n < d(x, \partial\Omega) \leq 2^k/n\}$$

(see Figure 2(b)).

Using Lemma 2.5, we can see that

CLAIM 5. *Denote by v_k the number of visits of X_t to R_k before the time T when X_t first hits the $1/n$ -neighborhood of the boundary $\partial\Omega$,*

$$v_k = \#\{t < T : X_t \in R_k\}.$$

Then

$$\mathbf{P}[v_k > C_2 \cdot 2^{k(\alpha-2)} M] < 1/4^M,$$

for some constant $C_2 = C_2(c, d, \alpha, \beta)$ and for any $M > 1$.

As shown in the full version of the paper, if we select k_0 such that $2^{k_0} \approx n^{2/\alpha}$ in Claim 5, sum up the values of v_k for $k \leq k_0$, and then let $\varepsilon = 2^{k_0}/n$ in Claim 4, we obtain the required upper bounds on the rate of convergence for the case $\alpha > 2$.

3 Lower bounds: examples

In this section we construct examples of α -thick domains for which the bounds in Theorem 1.1 are tight. The main idea of the construction is as follows. We take a domain A in \mathbb{R}^d , such as the unit ball or a cylinder. We remove a “thin” subset of points C from A to obtain $\Omega = A \setminus C$. The set C can be thought of as the subset of the grid $(\gamma\mathbb{Z})^d$, for some small $\gamma > 0$. The set C will be chosen so that it “separates” the origin from the boundary of A . We set $n = 1/\varepsilon$. We choose γ so that the probability of the WoS originated at 0 hitting a $1/n$ -neighborhood of C before hitting the boundary of A is $< 1/2$ (this means that C is “thin”). Hence, with high probability, the WoS will reach ∂A before terminating. However, in this case the WoS will have to “pass through” the set C , where its step magnitudes are bounded by γ . This will, in turn, yield an $\Omega(1/\gamma^2)$ bound on the convergence time. The analysis is more intricate in the case when $\alpha = 2$. In the case when $\alpha > 2$ is not an integer, a slight modification to this construction is needed, as will be described below.

3.1 Proof of the lower bound in the case $\alpha > 2$

In this section we will give an example of a “thin” α -thick domain Ω_α for which the WoS will likely take $\Omega(n^{2-4/\alpha})$ steps to converge within $\varepsilon = 1/n$ from the boundary $\partial\Omega_\alpha$. The domain Ω_α will reside in \mathbb{R}^d , where $d = \lceil \alpha \rceil \geq 3$. It is easy to see that the examples in higher dimensions $d' > d$ can be constructed from Ω_α by simply multiplying Ω_α by $[-1, 1]^{d'-d}$.

The set

$$\Omega_\alpha := \left(B(0, 1)_{d-1} \times [-1, 1] \right) \setminus S$$

is comprised of a d -dimensional cylinder with a set of points S removed. Here $B(0, 1)_{d-1}$ denotes the unit ball in \mathbb{R}^{d-1} . We take A to be the “middle 1/3” shell of the d -dimensional cylinder:

$$A = \{z \in \mathbb{R}^{d-1} : 1/3 < |z| < 2/3\} \times \{x \in [-1, 1] : 1/3 < |x| < 2/3\}.$$

Let $0 < \gamma \ll 1$ be the grid size that will be selected later. We consider the set A_γ of gridpoints in A .

$$A_\gamma = (\gamma\mathbb{Z})^d \cap A.$$

Let $0 \leq \eta := d - \alpha < 1$. Denote by C_η the η -dimensional Cantor set in the interval $[0, 1]$. It is obtained by removing the middle λ -fraction of the interval, then removing the middle λ -fraction of each subinterval etc. For the set C_η to be η -dimensional, we choose λ so that

$$\eta = \frac{\log 2}{\log 2 - \log(1 - \lambda)}.$$

In the special case when $\eta = 0$, we set $C_0 = \{0\}$. We now define the set S :

$$S := A_\gamma + \{0\} \times \gamma C_\eta.$$

In other words, S is obtained by attaching a γ -scaled copy of C_η to each gridpoint of A_γ . This completes the definition of the set $\Omega_\alpha = (B(0, 1)_{d-1} \times [-1, 1]) \setminus S$. The intuition that each point in $\partial\Omega_\alpha$ has an η -dimensional set in Ω_α^c attached to it is captured by the following claim.

CLAIM 6. *There is a universal constant $c \geq 1/16$ such that for every γ , the set Ω_α is α -thick with the thickness c .*

Proof. The statement is trivial in the case $\alpha = d$ (and $\beta = 0$). In all other cases, let x be a point in $\partial\Omega_\alpha$. We will construct μ_x as in Definition 1. The construction is obvious if x is in the boundary of the unit ball, hence we only need to concentrate on the case when $x \in S$. The measure μ_x is supported on the fractal line S_x in S that passes through x :

$$S_x := \{z \in S : \text{the first } d-1 \text{ coordinates of } z \text{ are identical to those of } x\}.$$

S_x is composed of $\Theta(1/\gamma)$ small copies of the Cantor set C_β . The metric μ_x will assign weight w_0 to the two copies that are closest to x , weight w_1 to the next two closest copies of C_β etc. Within each copy, μ_x will be the β -dimensional Hausdorff measure on C_β scaled to have a total weight of w_i for the appropriate w_i .

We choose $w_0 = w_1 = \gamma^\eta/4$, $w_i = (i^\eta - (i-1)^\eta)\gamma^\eta/4$, so that for $k \geq 1$, $\sum_{i=0}^k 2w_i = i^\eta\gamma^\eta/2 + \gamma^\eta/2 \leq (i\gamma)^\eta$. It is not hard to see that the resulting measure μ_x satisfies the conditions of Definition 1 with $c = 1/16$.

The following two claims assert that for an appropriately chosen γ , the WoS originated at the origin $0 \in \mathbb{R}^d$ and terminated at the $1/n$ neighborhood of $\partial\Omega_\alpha$ is likely to hit the boundary of the external cylinder (as opposed to the neighborhood of S), and is likely to spend $\Omega(n^{2-4/\alpha})$ steps getting there.

CLAIM 7. *If $\gamma > 8n^{2/\alpha-1}$ then a WoS originated at 0 and terminated at the $1/n$ -neighborhood of the boundary $\partial\Omega_\alpha$ will hit the boundary of the cylinder $B(0, 1)_{d-1} \times [-1, 1]$ with probability at least $3/4$.*

Proof. It is not hard to see that we can choose a finite subset P of points in S such that $|P| < 2\gamma^{-\alpha} \cdot n^\beta$, and for every x such that $d(x, S) < 1/n$ there is a $p \in P$ such that $|x - p| < 2/n$. Consider the harmonic function

$$(3.13) \quad \Phi(x) := \sum_{y \in P} \frac{1}{|x - y|^{d-2}} > 0.$$

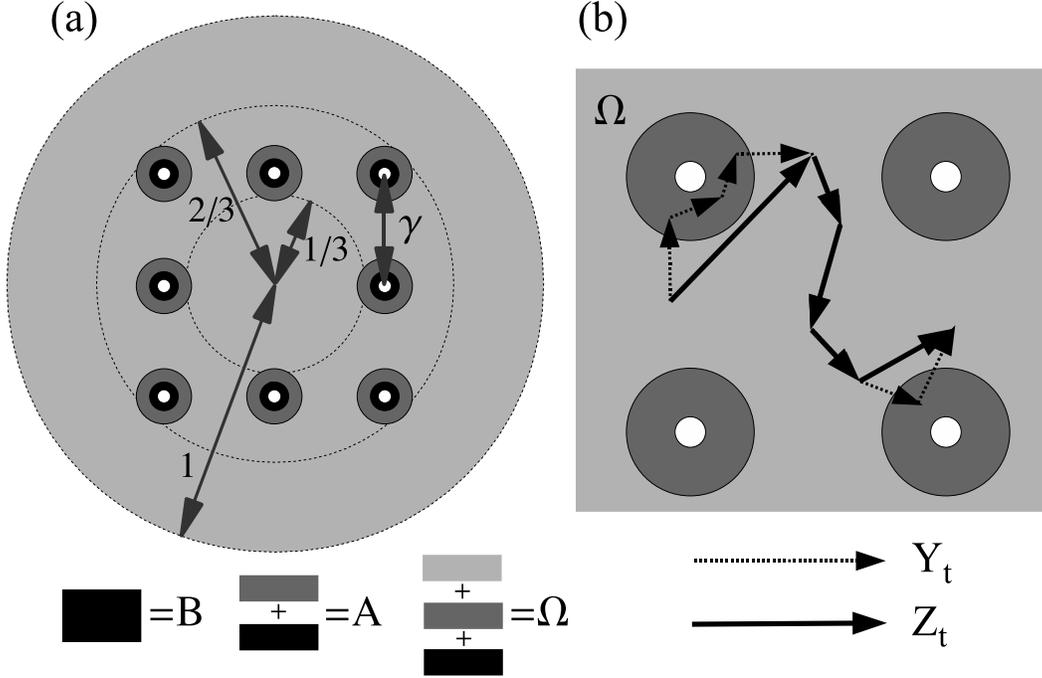


Figure 2: An illustration the sets Ω , A and B (a), and a possible sequence of jumps in the processes $\{Y_t\}$ and $\{Z_t\}$ (b)

Since the function Φ is harmonic, its application to the WoS process X_t gives a martingale. Hence if T is the stopping time of the process,

$$\mathbf{E}[\Phi(X_T)] = \Phi(X_0) = \Phi(0) < 3^{d-2} \cdot |P| < 6\gamma^{-\alpha} \cdot n^\eta.$$

On the other hand, if $d(X_T, S) < 1/n$, then there is a $y \in P$ with $|X_T - y| < 2/n$, and

$$\Phi(X_T) \geq 1/|y - X_T|^{d-2} > (n/2)^{d-2}.$$

Hence the probability of X_T being near S is bounded by

$$\frac{\mathbf{E}[\Phi(X_T)]}{(n/2)^{d-2}} < \frac{6\gamma^{-\alpha} \cdot n^\eta}{(n/2)^{d-2}} < \frac{2^{d+1}\gamma^{-\alpha}}{n^{\alpha-2}} < \frac{8^\alpha n^2}{4(\gamma n)^\alpha} < 1/4.$$

The last inequality follows from the condition on γ .

CLAIM 8. *There is a universal constant $\delta > 0$ such that for γ as above, with probability at least $1/2$ the WoS takes at least $\delta(1/\gamma)^2$ steps to reach the boundary of the cylinder $B(0, 1)_{d-1} \times [-1, 1]$.*

The proof is done analogously to the proof of Claim 11 below.

Hence the expected number of steps is at least

$$\frac{\delta}{2} \cdot \left(\frac{1}{8n^{2/\alpha-1}} \right)^2 = \Omega(n^{2-4/\alpha}),$$

which completes the proof of the lower bound for Theorem 1.1 in the case when $\alpha > 2$.

3.2 Proof of the lower bound in the case $\alpha = 2$

We will now give an example of a two dimensional domain Ω such that the expected convergence time of the WoS to a $O(1/n)$ -neighborhood of $\partial\Omega$ is $\Omega(\log^2 n)$. By taking the d -dimensional domain $\Omega_d = \Omega \times \mathbb{R}^{d-2}$ for $d > 2$, we obtain a lower bound of $\Omega(\log^2 n)$ for 2-thick domains in \mathbb{R}^d , proving the lower bound for $\alpha = 2$ in Theorem 1.1.

The domain Ω will consist of the unit disc in \mathbb{R}^2 with $O(\log n)$ holes “poked” out of it in a grid formation. More specifically, let $\gamma = 4/\log^{1/2} n$. We consider the grid $\Gamma = \gamma\mathbb{Z} \times \gamma\mathbb{Z} \subset \mathbb{R}^2$. We take Ω to be the unit disc with points from Γ removed from the “middle third” annulus of the disc. Let $\Omega = B(0, 1) \setminus ((B(0, 2/3) \setminus B(0, 1/3)) \cap \Gamma)$. The set Ω is illustrated on Fig. 2(a).

We will show that a WoS originated at the origin $X_0 = 0$ would require an expected time of $\Omega(\log^2 n)$ to converge. It is immediate to see that the same lower bound holds for any point $X_0 \in B(0, 1/3)$. We first observe the following:

CLAIM 9. *With probability at least $7/8$, a WoS origi-*

ated at $X_0 = 0$ that runs until $d(X_t) < 1/n$ terminates near the unit circle (and not near one of the holes).

Proof. Let $\{a_i\}_{i=1}^k = B(0, 1) \setminus \Omega$ be the set of holes in Ω . Define the harmonic function $\Phi(z) = \sum_{i=1}^k \log(2/|z - a_i|)$.

It is clear that $\Phi(z) > 0$ for all $z \in B(0, 1)$. For any point u in the $1/n$ -neighborhood of any of the holes, $\Phi(u) > \log n$. On the other hand, $\Phi(0) < k \cdot \log 6 < 2/\gamma^2 = (\log n)/8$.

If X_t is the WoS process with $X_0 = 0$ terminated at time T when $d(X_T, \partial\Omega) < 1/n$, then $\Phi(X_t)$ is a martingale. Hence,

$$(\log n)/8 > \Phi(X_0) = \mathbf{E}[\Phi(X_t)] > \mathbf{P}[X_t \text{ near a hole}] \cdot \log n.$$

Hence the probability that the WoS terminates near a hole is less than $1/8$.

For simplicity, we will assume that at every step of the process the WoS jumps exactly half way to the boundary $\partial\Omega$.

To facilitate the analysis we replace the WoS process X_t on Ω with the following process Y_t . It evolves in exactly the same fashion as X_t , except when Y_t is closer than $1/n$ to one of the holes in Ω . In this case, instead of terminating, the process makes a jump of $1/n$ in a direction selected uniformly at random. The process Y_t is guaranteed to terminate near the unit circle. We denote the termination time by T . Further, we set $Y_t = Y_T$ for $t > T$. Note that if the process X_t does not terminate near one of the holes, then the process Y_t coincides with X_t . Claim 9 implies that this happens with probability at least $7/8$:

CLAIM 10. $\mathbf{P}[X_t \text{ does not coincide with } Y_t] < 1/8$.

We define two regions A and B , $B \subset A \subset \Omega$. We take A to be the union of discs with radius $r = \gamma/4$ around the holes in Ω . We take B to be the union of discs with radius $r/2$ around the same holes. The sets Ω , A and B are illustrated on Fig. 2(a).

Let time t_0 be the first time with $|Y_t| > 1/2$. Let t' be the first time afterward with either $|Y_t| > 2/3$ or $|Y_t| < 1/3$. Our goal is to show that with probability at least $3/4$, $|t_0 - t'| = \Omega(\log^2 n)$. We define a subprocess Z_t of Y_t as follows. Let $\{s_i\}_{i=0}^k$ be a subsequence of times s between t_0 and t' such that $Y_s \notin A$. We set $Z_i = Y_{s_i}$. We further define $\Delta_i = Z_i - Z_{i-1}$. An instance of the process Z_i is illustrated on Fig. 2(b). Since Y_t is a martingale, and Z_i is defined by a stopping rule on Y_t , Z_i is also a martingale, and

$$(3.14) \quad \mathbf{E}[\Delta_i \mid \Delta_1, \Delta_2, \dots, \Delta_{i-1}] = 0.$$

In addition, it is not hard to see from the definition of Y_t that $|\Delta_i| < 4/\log^{1/2} n$ for all i . Our first claim is that the number k of steps Z_i is $\Omega(\log n)$.

CLAIM 11. $\mathbf{P}[k < 10^{-4} \log n] < 1/8$.

Proof. Denote $\ell = 10^{-4} \log n$. Then, by (3.14),

$$\begin{aligned} \mathbf{E}[(Z_0 - Z_\ell)^2] &= \mathbf{E}[(\Delta_1 + \Delta_2 + \dots + \Delta_\ell)^2] = \\ &= \sum_{j=1}^{\ell} \mathbf{E}[\Delta_j^2] + \sum_{1 \leq i < j \leq \ell} \mathbf{E}[\Delta_i \cdot \Delta_j] = \\ &= \sum_{j=1}^{\ell} \mathbf{E}[\Delta_j^2] < \ell \cdot 16/\log n < 1/288. \end{aligned}$$

On the other hand, by definition, $|Z_0 - Z_k| > 1/6$, and $(Z_0 - Z_k)^2 > 1/36$. Hence,

$$\mathbf{P}[k \leq \ell] = \mathbf{P}[Z_\ell = Z_k] < (1/288)/(1/36) = 1/8.$$

Thus the number of steps the process Z_t takes is at least $10^{-4} \log n$ w.p. $> 7/8$. The process Y_t consists of the steps of the process Z_t plus, in addition, steps the process takes within the region A . We claim that once the process Y_t enters the region A , it is expected to spend $\Omega(\log n)$ steps there. Moreover, the following holds.

CLAIM 12. *Let $\eta > 2$. Then there is a $\theta > 0$ such that whenever $Y_t \in A$, if $s > t$ is the first time, conditioned on Y_t such that $Y_s \notin A$, then*

$$(3.15) \quad \mathbf{P}[s - t > \theta \log^2 n] > \eta/\log n,$$

for sufficiently large n .

Proof. Denote the hole in Ω that is closest to Y_t by x . Given that $Y_t \in A$, there is some fixed probability $p > 0$ that $Y_{t+1} \in B$. In other words, $|Y_{t+1} - x| < r/2 = \gamma/8$. Consider the harmonic function

$$\Phi(z) = \log(r/|x - z|).$$

Let $t' > t + 1$ be the first time such that either $Y_{t'} \notin A$ (and thus $t' = s$), or $|Y_{t'} - x| < n^{-p/(5\eta)}$. If $Y_{t'} \notin A$, then $\Phi(Y_{t'}) < 0$. In the other case, $\Phi(Y_{t'}) < p \log n/(4\eta)$. Since t' is a stopping time, the optional stopping time theorem applied to the martingale $\Phi(Y_{t+\tau})$ combined with the estimate $\Phi(Y_t) > 1/2$, gives

$$\mathbf{P}[|Y_{t'} - x| < n^{-p/(5\eta)}] > (1/2)/(p \log n/(4\eta)) > 2\eta/(p \log n).$$

To complete the argument, we claim that assuming $|Y_{t'} - x| < n^{-p/(5\eta)}$, it will take the process another

$\Omega(\log^2 n)$ steps to escape A with probability at least $1/2$. We consider the process $\phi_\tau = \Phi(Y_{t'+\tau})$ stopped at time τ_0 when either $Y_{t'+\tau_0}$ escapes A , or gets closer than distance $1/n$ from x . The process ϕ_τ is a martingale. Moreover, it is not hard to see that $|\phi_0 - \phi_{\tau_0}| > p \log n / (6\eta)$, and $|\phi_i - \phi_{i+1}| < 1$ for all i . These two facts imply that

$$\mathbf{E}[\tau_0] > \sum_{i=1}^{\tau_0} \mathbf{E}[(\phi_i - \phi_{i-1})^2] = \mathbf{E}[(\phi_{\tau_0} - \phi_0)^2] > (p \log n / (6\eta))^2 = p^2 \log^2 n / (36\eta^2).$$

Tschebyshev inequality implies that $\theta = p^2 / (72\eta^2)$ satisfies the statement of the claim.

By Claim 11 we know that except with probability $< 1/8$ the walk will contain at least $\Omega(\log n)$ visits to A . It remains to use Claim 12 to show that at least one of these stays must be $\Omega(\log^2 n)$ long. Recall that T is the stopping time of the process Y_T , and k is the number of steps Y_t takes outside of A .

CLAIM 13. *Let $\alpha_1 = 10^{-4}$ from Claim 11. There is a constant $\alpha_2 > 0$ such that*

$$(3.16) \quad \mathbf{P}[k > \alpha_1 \log n \text{ and } T < \alpha_2 \log^2 n] < 1/8.$$

We defer the proof of the lemma to the full version of the paper.

Claims 10, 11 and 13 imply the following.

CLAIM 14. *Let X_t be the WoS process on the set Ω with $X_0 = 0$. Let T' be its termination time. Then*

$$\mathbf{P}[T' > \alpha_2 \log^2 n] > 5/8,$$

where $\alpha_2 > 0$ is the constant from Claim 13. In particular, this implies that $\mathbf{E}[T'] = \Omega(\log^2 n)$.

Proof. We know that $T' > \alpha_2 \log^2 n$ if the following three conditions hold: (C1) the process X_t coincides with the process Y_t ; (C2) the process Y_t makes at least $k > \alpha_1 \log n$ steps outside of A in the $\{z : 1/3 < |z| < 2/3\}$ annulus; and (C3) the stopping time T of Y_t satisfies $T > \alpha_2 \log^2 n$. In fact conditions (C1) and (C3) suffice. We have $\mathbf{P}[\overline{C1}] < 1/8$ by Claim 10, $\mathbf{P}[\overline{C2}] < 1/8$ by Claim 11, and $\mathbf{P}[C2 \cap \overline{C3}] < 1/8$ by Claim 13. Here \overline{C} denotes the complement of an event C . Hence

$$\mathbf{P}[\overline{C1} \cup \overline{C2} \cup \overline{C3}] \leq \mathbf{P}[\overline{C1}] + \mathbf{P}[\overline{C2}] + \mathbf{P}[C2 \cap \overline{C3}] < 3/8,$$

which implies that $\mathbf{P}[T' > \alpha_2 \log^2 n] > 5/8$.

Claim 14 gives the lower bound for Theorem 1.1 in the case $\alpha = 2$.

References

- [1] I. Binder and M. Braverman. Derandomization of Euclidean Random Walks. *RANDOM'07, LNCS*, 4627:353–365, 2007.
- [2] V. Brattka and K. Weihrauch. Computability of subsets of Euclidean space I: Closed and compact subsets. *Theoretical Computer Science*, 219:65–93, 1999.
- [3] M. Braverman and S. Cook. Computing over the reals: Foundations for scientific computing. *Notices of the AMS*, 53(3):318–329, 2006.
- [4] B. S. Elepov, A. A. Kronberg, G. A. Mihailov, and K. K. Sabel'fel'd. *Reshenie kraevykh zadach metodom Monte-Karlo*. “Nauka” Sibirsk. Otdel., Novosibirsk, 1980.
- [5] P. Embrechts, C. Klüppelberg, and T. Mikosch. *Modelling Extremal Events for Insurance and Finance*. Springer, 1997.
- [6] J. B. Garnett and D. E. Marshall. *Harmonic Measure*. Cambridge Univ Press, 2004.
- [7] Shizuo Kakutani. Two-dimensional Brownian motion and harmonic functions. *Proc. Imp. Acad. Tokyo*, 20:706–714, 1944.
- [8] I. Karatzas and S.E. Shreve. *Methods of Mathematical Finance*. Springer, 1998.
- [9] N. S. Landkof. *Foundations of modern potential theory. Translated from the Russian by AP Doohovskoy*, volume 180. 1972.
- [10] R.M. Mazo. *Brownian Motion: Fluctuations, Dynamics, and Applications*. Oxford University Press, 2002.
- [11] G. A. Mihailov. Estimation of the difficulty of simulating the process of “random walk on spheres” for some types of regions. *Zh. Vychisl. Mat. i Mat. Fiz.*, 19(2):510–515, 558–559, 1979.
- [12] G. N. Milstein. *Numerical Integration of Stochastic Differential Equations*. Kluwer Academic Publishers, Dordrecht, 1995.
- [13] Minoru Motoo. Some evaluations for continuous Monte Carlo method by using Brownian hitting process. *Ann. Inst. Statist. Math. Tokyo*, 11:49–54, 1959.
- [14] M. E. Muller. Some continuous Monte Carlo methods for the Dirichlet problem. *Ann. Math. Statist.*, 27:569–589, 1956.
- [15] E. Nelson. *Dynamical Theories of Brownian Motion*. Princeton University Press Princeton, NJ, 1967.
- [16] A.S. Sznitman. *Brownian Motion, Obstacles and Random Media*. Springer, 1998.
- [17] K. Weihrauch. *Computable Analysis*. Springer-Verlag, Berlin, 2000.