

A Hit-and-Run Approach for Generating Scale Invariant Small World Networks *

Archis Ghatge
Industrial Engineering Program
The University of Washington
Seattle, Washington 98195
(archis@u.washington.edu)

Robert L. Smith
Department of Industrial and Operations Engineering
The University of Michigan
Ann Arbor, Michigan 48105

Abstract

Hit-and-Run is a well-known class of Markov chain algorithms for sampling from essentially arbitrary distributions over bounded regions of the Euclidean space. We present a class of Small World network models constructed using Hit-and-Run in a Euclidean ball. We prove that there is a unique scale invariant model in this class that admits efficient search by a decentralized algorithm. This research links two seemingly unrelated areas: Markov chain sampling techniques and scale invariant Small World networks, and may have interesting implications for stochastic search methods for continuous optimization.

Keywords: Small World Networks, Scale Invariant Networks, Markov Chain Sampling, Continuous Optimization.

1 Introduction

Hit-and-Run is a Markov chain technique originally proposed by Smith [22] for generating approximately uniformly distributed points on bounded, open subsets S of R^n , and later extended by Belisle et al, [4] and Romeijn and Smith [21] to arbitrary distributions. Its simplest version makes a one-step transition from a point $x \in S$ to another point $y \in S$ by generating a direction vector uniformly distributed on the surface of a d -dimensional unit hypersphere around x , followed by generating a point y uniformly distributed on the line segments created by the intersection of the direction vector and S (this is accomplished by employing a one-dimensional rejection method on the line segment intersected by an enclosing box for the open region). The sequence of iterates generated in this intuitive manner was shown to converge in total variation to the uniform

*This work has been funded in part by an NSF collaborative research grant number DMI-0244291

distribution on S [22]. This version of Hit-and-Run was shown to be the fastest known method for generating an asymptotically uniform point from a convex body in R^n [12] assuming that the initial distribution of the Hit-and-Run Markov chain was not far from uniform, i.e., a ‘warm start’. This assumption was later relaxed [15] making Hit-and-Run the only known random walk that converges efficiently to a uniform distribution starting from any point inside a convex body. Moreover, it was also shown [13] to be very efficient for sampling from log-concave distributions over convex bodies. The Hit-and-Run sampler has found many applications including identifying non-redundant constraints [5], global optimization [20, 21, 26], convex optimization [6, 7], and computing the volume of convex bodies [14]. In this paper, we employ Hit-and-Run to accomplish a novel task; that of growing Small World networks that admit efficient search by a decentralized algorithm.

The *Small World* phenomenon gained fame due to the pioneering work of Stanley Milgram and his team in the 1960’s [16, 23]. The phenomenon was observed through a series of social experiments. Milgram’s team realized that we are all linked by short chains of acquaintances. The social experiment randomly identified a pair of people in the United States unknown to each other; a source and a target. Basic information such as address and occupation of the target was provided to the source. A letter was handed to the source and the person was told that the letter should be passed on to a person she knows on a first name basis. The aim of the experiment was to deliver the letter to the target in as few steps as possible. Each intermediate person who received the letter also followed the same rule: the letter should be passed on to one of your acquaintances. The chain was continued until the target received the letter. Over a large number of trials, the average number of steps in a successful chain was found to be between five and six. This observation was popularized as *Six Degrees of Separation*.

In recent years, researchers have extensively studied modeling and algorithmic aspects of networks to characterize the Small World phenomenon mathematically [8, 9, 17, 24, 25]. The reader is referred to Albert and Barabasi [1], Newman [18], and Kleinberg [11] for comprehensive surveys of literature in this area. For example, it was shown in [24] that it is possible to construct networks in which any two nodes are connected by short chains. This can be done by superimposing a small number of uniformly random ‘long-range connections’ on a lattice where each node originally has a few ‘local contacts’. This approach resolved the issue of existence of short chains. However, the most striking component of Milgram’s experiment is not that people are connected by short chains, but that they are able to discover these chains by employing a simple letter forwarding procedure. Existence of short chains in a network where the long-range contacts are chosen uniformly randomly *does not* imply that nodes acting only on local information will be able to *find* them [8]. In his seminal work, Kleinberg [9] showed that it is indeed possible to construct networks in which nodes can accomplish this task using a decentralized algorithm, i.e., an algorithm which uses only *local information* at the nodes it visits. This approach goes beyond mere existence of short chains in that it attempts to actually find them, and hence, is algorithmic. In Kleinberg’s lattice network, the probability of a long range connection between two nodes was allowed to decrease with a power of the distance between them. In particular, he proved that there is a *unique* power law exponent (equal to the lattice dimension) for which a decentralized algorithm can find short chains. Kleinberg argued that this is possible because of a structural property of the power law network with exponent equal to the lattice dimension... “it is the unique exponent at which a node’s long range contacts are nearly uniformly distributed over all distance scales..”, in other words, the network is *scale invariant*. It was evident from Kleinberg’s work that employing some notion of ‘proximity’ is essential to building network models that admit efficient search by a decentralized algorithm (also see Kleinberg [11] and Watts et al [25]). We generalize these ideas to Markov chain

based network models in Euclidean spaces.

More specifically, we introduce a class of network models within a unit ball (*the world*) in \mathfrak{R}^2 . As we shall see, a point u in a typical network from this family is connected to a finite number of points on the boundary of an ‘extremely small’ ball of radius $1 \gg \epsilon > 0$ around that point. These are u ’s local contacts. In addition, u also has one long-range contact in *the world* outside the small ball. This contact is generated by a one-step transition of a Markov chain from a Hit-and-Run family parameterized by a non-negative, real number a . When a is zero, the probability density function of the Markov transition kernel does not take geographic proximity into account. On the other hand, as the value of a increases, it is more likely that u ’s long-range contact will be near u . The balance of geographic cues provided by a and the size of the world is delicate; we will show that $a = 1$ is the *only* value of a for which the network model is scale invariant. As a result, for this value of a , *there exists* a greedy decentralized algorithm that finds a path of expected length $O(\text{poly}(\log(2/\epsilon)))$ connecting a ‘source’ point s in the network to a point in the network that is within a distance $\epsilon > 0$ of any other ‘target’ point t in the world. On the other hand, for $0 \leq a < 1$, *every* decentralized algorithm finds it hard to close-in on the target once it reaches near the target. Analogously, when $a > 1$, *every* decentralized algorithm faces a bottleneck when it is far away from the target. Mathematically, *any* decentralized algorithm expends $\Omega(\text{poly}(2/\epsilon))$ expected number of steps to find a path when $a \neq 1$.

2 A Family of Network Models

In this section, we propose our basic model and make preliminary observations that will be revisited repeatedly in Sections 3 and 4. We pose our central question toward the end of this section and answer it through three theorems. Detailed proofs of these theorems are provided in Section 4.

Our model has three components : the world in which we build our network, and local as well as long-range contacts of a point in the network. These are illustrated in Figure 2 and explained in detail here.

The World: Let B denote the closed unit ball centered at the origin $(0,0)$ in \mathfrak{R}^2 , i.e., $B = \{x \in \mathfrak{R}^2 : \|x\| \leq 1\}$, where $\|\cdot\|$ is the Euclidean norm. This is our *world*. Starting at any arbitrary point in this world, the network is built by adding points one-by-one. We describe this process by considering a generic point u in the network.

Local Contacts: Let $\epsilon > 0$ be a constant such that $\epsilon \ll 1$, $B_\epsilon(u)$ the closed ball of radius ϵ around u , and $\partial B_\epsilon(u)$ its boundary $\{x \in \mathfrak{R}^2 : \|x - u\| = \epsilon\}$. Let $D_\epsilon(u) = \partial B_\epsilon(u) \cap B$. Let $\mathcal{P}_{\pi/4}(D_\epsilon(u))$ be a partition of $D_\epsilon(u)$ into a finite number of arcs where the central angle of every arc (at u) is at most $\pi/4$. Point u has exactly one local contact in every arc in $\mathcal{P}_{\pi/4}(D_\epsilon(u))$.

Long-Range Contacts: The long-range contact of u is generated by employing a one-step transition of the Hit-and-Run Markov chain parameterized by a non-negative real number a on state space $\{B \setminus B_\epsilon(u)\}$ as follows:

1. Choose a unit direction vector ξ emanating from u and uniformly distributed on the boundary $\partial B_1(u)$ of the unit ball centered at u . Let θ be the angle between ξ and $-u$ (see Figure 1).
2. Let $l_u(\theta)$ be the set of real numbers that characterize the line segment of $\{B \setminus B_\epsilon(u)\}$ passing through u parallel to direction ξ , i.e., $l_u(\theta) = \{\lambda \in \mathfrak{R} : u + \lambda\xi \in \{B \setminus B_\epsilon(u)\}\}$. Independently

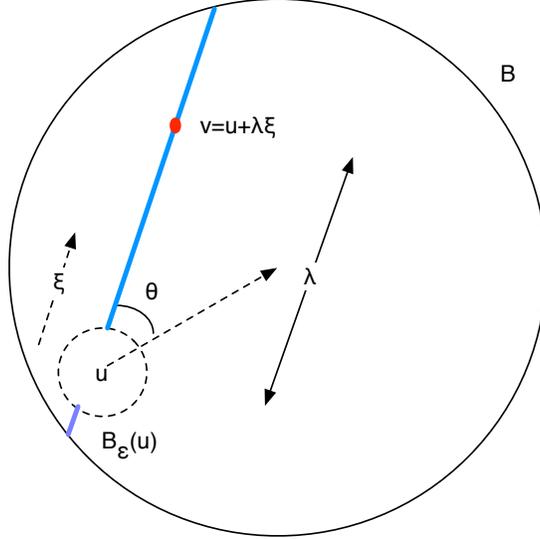


Figure 1: Directions ξ and $-u$ are shown with dotted arrows and θ is the angle between them. A chord of $B \setminus B_\epsilon(u)$ parallel to direction ξ is shown with a thick blue line. v is a point on this chord at distance λ from u .

choose a random number λ from $l_u(\theta)$ according to a probability distribution with density

$$\frac{\frac{1}{|\lambda|^a}}{\int_{\mu \in l_u(\theta)} \frac{1}{|\mu|^a} d|\mu|} \text{ when } l_u(\theta) \neq \emptyset, \quad a \geq 0. \quad (1)$$

Place the long-range contact v of u at $u + \lambda\xi$.

Remark 2.1. Notice that our local contact model essentially ‘discretizes’ $D_\epsilon(u)$. The upper bound on the central angle controls fineness of the partition. The bound of $\pi/4$ is not necessary, and can be increased in some cases (see Figure 2 (a) for example). However, it is sufficient to ensure that for any point $t \in B \setminus B_\epsilon(u)$, u will have at least one local contact w such that t is closer to w than u (see Lemma A.1 in Appendix A).

Remark 2.2. Observe that the integral in the denominator of Equation (1) is well-defined since $|\mu| > \epsilon$ for all $\mu \in l_u(\theta)$. We will use $C_{u,\epsilon}(\theta)$ to denote this integral and write the above density as $\frac{1}{C_{u,\epsilon}(\theta)|\lambda|^a}$. $C_{u,\epsilon}(\theta)$ is evaluated in Appendix A. Moreover, since $l_u(\theta)$ is one-dimensional, sampling according to density function (1) can be easily done by employing the standard left-continuous inverse method (see [4]).

We now compute the probability density for the distribution of long-range contacts. For $u \in B$ and $v \in \{B \setminus B_\epsilon(u)\}$, let $f(u, v)$ denote the probability density for the distribution of u ’s long-range contact. We have

Proposition 2.3. (Power Law Probability Density) The probability density function for the distribution of u ’s long-range contact is given by

$$f(u, v) = \frac{1}{\pi} \frac{1}{C_{u,\epsilon}(\theta_{uv}) \|u - v\|^{a+1}} \quad (2)$$

where θ_{uv} is the angle between $(v - u)$ and $-u$.

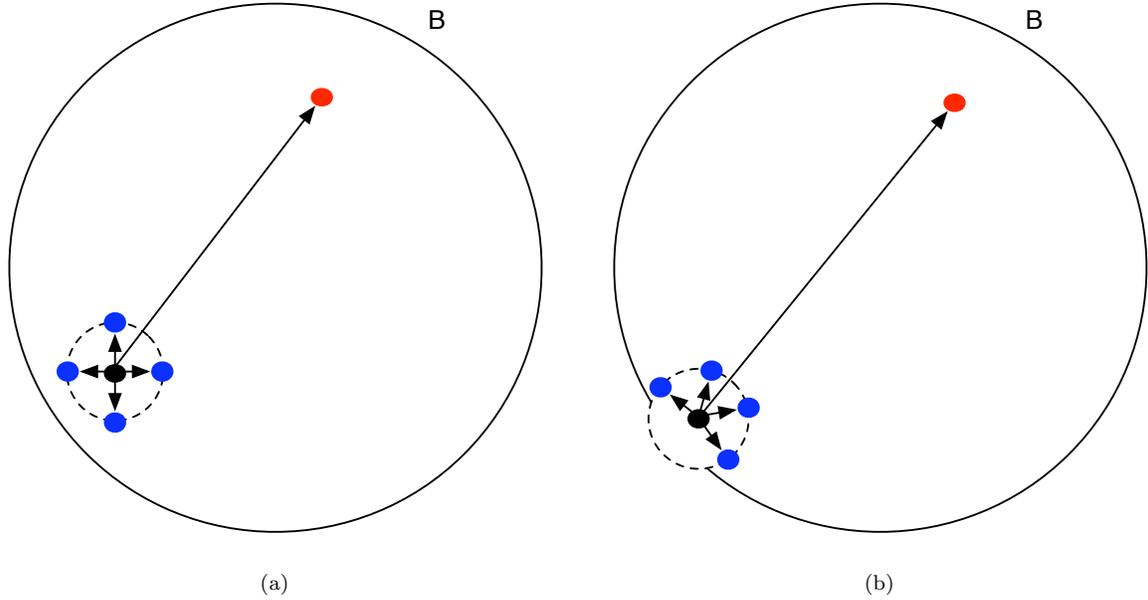


Figure 2: (a) Point u shown as a black dot is ϵ -deep inside B , i.e., $B_\epsilon(u) \subset B$. The four blue dots are u 's local contacts on $\partial B_\epsilon(u)$ shown as a dotted circle. Point v shown as a red dot is u 's long-range contact. (b) Point u is not deep inside B .

Proof. Let S_w be an infinitesimally small square with side w oriented along $(v - u)$. $f(u, v) = \lim_{A(S_w) \rightarrow 0} \frac{P(u, S_w)}{A(S_w)}$, where $A(S_w)$ is the area of S_w and $P(u, \cdot)$ is the probability that u 's long range contact is in S_w . Notice that $\|u - v\|$ is the same for all points in S_w when w is infinitesimally small. In order for u 's long-range contact to be in S_w , the unit direction ξ defined in step 1 of our long range contact model must pass through S_w and the number λ sampled in step 2 must be such that $u + \lambda\xi \in S_w$. Since ξ is chosen uniformly on $\partial B_1(u)$, the probability of the former is $2\frac{w}{2\pi\|u-v\|}$ whereas the latter is $\frac{w}{C_\epsilon(\theta_{uv})\|u-v\|^a}$. We have,

$$f(u, v) = \lim_{w \rightarrow 0} \frac{2\frac{w}{2\pi\|u-v\|} \frac{w}{C_\epsilon(\theta_{uv})\|u-v\|^a}}{w^2} = \frac{1}{\pi C_\epsilon(\theta_{uv})\|u-v\|^{a+1}}.$$

□

We refer to the probability density in Equation (2) as **the power law density**.

2.1 Milgram's Experiment: Efficient Message Delivery

Following Kleinberg, we put our network model into the context of Milgram's social experiments. Suppose a network with a countably infinite number of points is built by employing the above procedure starting at an arbitrary point $s \in B$. Let t be any arbitrary point in B other than s . Our objective is to deliver a message from s to a point in our network within a distance ϵ of t using a decentralized algorithm, that is, a procedure in which each message holder

1. knows the location of t .
2. knows its local and long-range contacts.

Most importantly, a message holder *does not* have any information about the local and long-range contacts of other points in the network, and hence, must forward the message to one of *its own* contacts. We will say that the message has been ϵ -delivered when it enters an ϵ ball around t . We consider the following question.

Question: For what values of parameter a (if any) does there exist a decentralized algorithm that ϵ -delivers the message in $O(\text{poly}(\log(2/\epsilon)))$ expected number of steps? We will call any such algorithm *efficient*.

Remark 2.4. *Although it is beneficial to think of an underlying countably infinite network in B that is built prior to the implementation of a decentralized algorithm to relate our work to Milgram's experiments, this is not required for the analysis. One can assume that the network is 'built on the fly', i.e., starting with s , a point in B forms its contacts **when and only when** the message is delivered to it. We will employ this framework for our analysis.*

Theorems 2.5, 2.6 and 2.7 answer the above question.

Theorem 2.5. *When $a = 1$, there exists a decentralized algorithm (in particular, the greedy algorithm G^* defined in Section 4) whose expected ϵ -delivery time for any (s, t) pair in $B \times B$ is at most $b \left(\log \frac{2}{\epsilon}\right)^2$, where b is a constant that does not depend on ϵ .*

Theorem 2.6. *Let $0 \leq a < 1$. Then there exist (s, t) pairs in $B \times B$ such that the expected ϵ -delivery time of **any** decentralized algorithm is at least $c_a \left(\frac{2}{\epsilon}\right)^{(1-a)/2}$, where c_a is a constant that depends only on a .*

Theorem 2.7. *Let $a > 1$. Then there exist (s, t) pairs in $B \times B$ such that the expected ϵ -delivery time of **any** decentralized algorithm is at least $d_a \left(\frac{2}{\epsilon}\right)^{(a-1)/a}$, where d_a is a constant that depends only on a .*

Remark 2.8. *Note here that $2/\epsilon$ is a measure of distance in our world as the diameter of the unit ball is 2. When the radius of our world is R , this distance measure is replaced by $2R/\epsilon$ in the above theorems whereas the constants remain independent of both R and ϵ .*

Proofs of these results are presented in Section 4. However, we first investigate two properties of the power law density in Section 3 below.

3 Scale Invariance and Steady Improvement

This section is devoted to establishing two important properties of the power law density: **scale invariance** and **steady improvement**. We begin by introducing the notion of scale invariance in B . See Kleinberg [8] for an informal description of a similar idea in the discrete lattice context.

Definition 3.1. *Let $0 < \beta < 1$ be any positive fraction and u be at the origin. For integers $j \geq 1$ such that $\beta^{j+1} > \epsilon$, define $P_j^\beta(u)$ to be the probability that the distance of u 's long-range contact from u is between β^{j+1} and β^j . We will say that the probability density function (2) is **scale invariant** for generating long-range contacts from the origin if for any β , $P_j^\beta(u)$ does not depend on j .*

Theorem 3.2. (Scale Invariance) *The power law density function is scale invariant for generating long-range contacts from the origin if and only if $a = 1$.*

Proof. We use the spherical coordinate system (r_{uw}, θ_{uw}) centered at u . In particular, for any point $w \in \mathfrak{R}^2$, $r_{uw} = \|u - w\|$, and θ_{uw} is the angle between $(w - u)$ and $-u$. Note that the infinitesimal area element in this system is given by $(r_{uw} dr_{uw} d\theta_{uw})$.

Proof of scale invariance of $a = 1$:

Note that since u is at the origin, it is ϵ -deep, and moreover, $\|u\| = 0$. From Lemma A.2 for $a = 1$, $C_{u,\epsilon}(\theta) = 2 \log\left(\frac{1}{\epsilon}\right)$ for any angle θ . Then for any $0 < \beta < 1$ and integer $j \geq 1$ with $\beta^{j+1} > \epsilon$,

$$P_j^\beta(u) = \int_0^{2\pi} \int_{\beta^{j+1}}^{\beta^j} \frac{1}{\pi} \frac{1}{2 \log \frac{1}{\epsilon}} \frac{1}{r_{uv}^2} (r_{uv} dr_{uv} d\theta_{uv}) = \frac{1}{2\pi} \frac{1}{\log \frac{1}{\epsilon}} \int_0^{2\pi} \log\left(\frac{\beta^j}{\beta^{j+1}}\right) (d\theta_{uv}) = \frac{\log(1/\beta)}{\log \frac{1}{\epsilon}},$$

which does not depend on j .

Proof of lack of scale invariance of $a = 0$:

From Lemma A.2 for $a = 0$, $C_{u,\epsilon}(\theta) = 2(1 - \epsilon)$ for any angle θ when u is at the origin. Then for any $0 < \beta < 1$ and integer $j \geq 1$ with $\beta^{j+1} > \epsilon$,

$$P_j^\beta(u) = \int_0^{2\pi} \int_{\beta^{j+1}}^{\beta^j} \frac{1}{\pi} \frac{1}{2(1 - \epsilon)} \frac{1}{r_{uv}} (r_{uv} dr_{uv} d\theta_{uv}) = \frac{\beta^j(1 - \beta)}{(1 - \epsilon)},$$

which obviously depends on j .

Proof of lack of scale invariance of $0 < a \neq 1$:

From Lemma A.2 for $0 < a \neq 1$, $C_{u,\epsilon}(\theta) = \frac{2}{(1-a)} (1 - \epsilon^{1-a})$ for any angle θ when u is at the origin. Hence,

$$P_j^\beta(u) = \int_0^{2\pi} \int_{\beta^{j+1}}^{\beta^j} \frac{1}{\pi} \frac{1}{\frac{2}{(1-a)} (1 - \epsilon^{1-a})} \frac{1}{r_{uv}^{a+1}} (r_{uv} dr_{uv} d\theta_{uv}) = \frac{\beta^{j(1-a)} \{1 - \beta^{1-a}\}}{(1 - \epsilon^{1-a})},$$

which again depends on j . □

Another interesting property of the parameter $a = 1$ is discussed in the remainder of this section. Informally, we say that the power law density exhibits **steady improvement** if the probability that a point u at distance r from a point t in the world has a long-range contact within distance $r/2$ from t is $\Omega(1/\log(2/\epsilon))$ (independent of r). Making this precise requires some geometric and algebraic work as shown here.

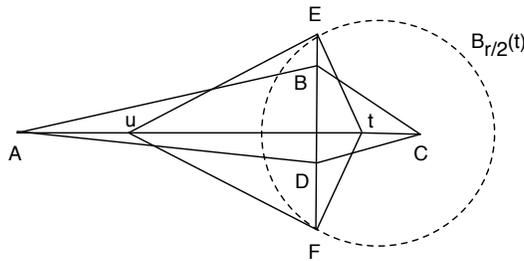


Figure 3: Illustration of the first part of the proof of Theorem 3.3

Theorem 3.3. (Steady Improvement) *Let u, t be any two points in B such that $\|u - t\| = r \geq 2\epsilon$. Let $Q_r(u, t)$ be the probability that u has a long-range contact within distance $r/2$ from t . Then, when $a = 1$, there is a constant c (that is independent of r and ϵ) such that*

$$Q_r(u, t) \geq \frac{c}{\log \frac{2}{\epsilon}}. \quad (3)$$

Proof. The boundary $\partial B_{r/2}(t)$ of disc $B_{r/2}(t)$ of radius $r/2$ centered at t is shown in Figure 3 with a dotted line. $Q_r(u, t)$ is the probability that u 's long-range contact lies in $B_{r/2}(t) \cap B$ (note that $B_{r/2}(t)$ may not entirely lie inside B). Let uE and uF be the two tangents to $\partial B_{r/2}(t)$ from u . We first show by contradiction that at least one of E and F is inside B . Suppose not. Then let B and D be the points of intersection of B with segment EF (these two points exist since E and F are outside B and both u and t are inside B). Let A and C be the end points of the chord of B that passes through u and t . Consider triangles ABC and ADC . Note that $\angle ABC > \pi/2$ and $\angle ADC > \pi/2$. Therefore, the circumcenter of triangle ABC lies outside that triangle and in particular, on the D side of segment AC . Similarly, the circumcenter of triangle ADC lies outside that triangle and in particular, on the B side of segment of AC . However, this is a contradiction since the circumcenters of these two triangles coincide and in particular are at the center of B , i.e., the origin. Thus, at least one of E and F (say E) is in B .

Now refer to Figure 4. We bound $Q_r(u, t)$ below by the probability that u 's long-range contact is inside band $EGHI$ defined as follows. I is the point of intersection of the arc of radius uE centered at u with segment ut . Note that segment uE has length $\sqrt{3}r/2$ by Pythagoras theorem in triangle uEt ; H is the point of intersection of the arc of radius $3r/4$ centered at u with segment ut whereas G is where this arc intersects $\partial B_{r/2}(t)$. Applying the cosine rule to triangle uGt we get

$$\begin{aligned} (tG)^2 &= (uG)^2 + (ut)^2 - 2(uG)(ut) \cos(\angle tuG), \quad \text{i.e.,} \\ (r/2)^2 &= (3r/4)^2 + (r)^2 - 2(3r/4)(r) \cos(\angle tuG), \quad \text{i.e.,} \\ \angle tuG &= \arccos(7/8). \end{aligned}$$

Finally, Lemmas A.2 and A.3 imply that $C_{u,\epsilon}(\theta) \leq 2 \log(2/\epsilon)$ for any angle θ . Hence,

$$\begin{aligned} Q_r(u, t) &\geq \int_0^{\arccos(7/8)} \int_{3r/4}^{\sqrt{3}r/2} \frac{1}{\pi} \frac{1}{2 \log(2/\epsilon)} \frac{1}{r_{uv}^2} r_{uv} dr_{uv} d\theta_{uv} \geq \int_0^{\pi/12} \int_{3r/4}^{\sqrt{3}r/2} \frac{1}{\pi} \frac{1}{2 \log(2/\epsilon)} \frac{1}{r_{uv}} dr_{uv} d\theta_{uv} \\ &= \frac{\pi/12}{2\pi \log(2/\epsilon)} \log(2/\sqrt{3}) = \frac{\log \frac{2}{\sqrt{3}}}{24 \log \frac{2}{\epsilon}}. \end{aligned}$$

□

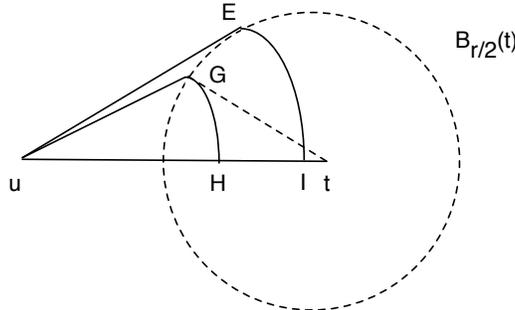


Figure 4: Illustration of the second part of the proof of Theorem 3.3

Theorem 3.4. *The power law density function for generating long-range contacts using one step of Hit-and-Run exhibits steady improvements **only** if $a = 1$.*

Proof. We consider three cases.

Proof for $a = 0$.

It suffices to prove that the density in Equation (2) fails to make significant improvements for a *specific* choice of u and t . Let t be at the center of B , u be ϵ -deep in B . Suppose $r \geq 2\epsilon$ as in Theorem 3.3, and also that $\sqrt{1 - r^2} > 2\epsilon$. That is,

$$\sqrt{1 - 4\epsilon^2} > r \geq 2\epsilon. \quad (4)$$

We have,

$$\begin{aligned} Q_r(u, t) &= 2 \int_0^{\pi/6} \int_{r \cos \theta_{uv} - \frac{r}{2} \sqrt{1 - 4 \sin^2 \theta_{uv}}}^{r \cos \theta_{uv} + \frac{r}{2} \sqrt{1 - 4 \sin^2 \theta_{uv}}} \frac{1}{\pi} \frac{1}{2(\sqrt{1 - r^2 \sin^2(\theta_{uv})} - \epsilon)} \frac{1}{r_{uv}} (r_{uv} dr_{uv} d\theta_{uv}) \\ &= \frac{1}{\pi} \int_0^{\pi/6} \frac{r \sqrt{1 - 4 \sin^2 \theta_{uv}}}{(\sqrt{1 - r^2 \sin^2(\theta_{uv})} - \epsilon)} (d\theta_{uv}) \\ &= \frac{1}{2\pi} \int_0^{\pi/6} \frac{\sqrt{1 - 4 \sin^2 \theta_{uv}}}{(\sqrt{(1/r)^2 - \sin^2(\theta_{uv})} - (\epsilon/r))} (d\theta_{uv}) \\ &\leq \frac{1}{\pi} \int_0^{\pi/6} \frac{\sqrt{1 - 4 \sin^2 \theta_{uv}}}{(\frac{1}{2}(\sqrt{(1/r)^2 - \sin^2(\theta_{uv})}))} (d\theta_{uv}) \leq \frac{2}{\pi} \int_0^{\pi/6} \frac{2 \cos(\theta_{uv})}{\sqrt{(1/r)^2 - \sin^2(\theta_{uv})}} (d\theta_{uv}) \\ &= \frac{4}{\pi} \int_0^{r/2} \frac{dy}{\sqrt{1 - y^2}} = \frac{4}{\pi} \arcsin\left(\frac{r}{2}\right). \end{aligned}$$

Now suppose $Q_r(u, t)$ is $\Omega(1/\log(2/\epsilon))$, i.e., there exists a constant $c > 0$ independent of ϵ such that $Q_r(u, t) \geq c/\log(2/\epsilon)$ for all r as in (4). Specifically, $(4/\pi) \arcsin(\epsilon) \geq Q_{2\epsilon}(u, t) \geq c/\log(2/\epsilon)$ implying that $c \leq (4 \log(2/\epsilon)/\pi) \arcsin(\epsilon)$. But the right hand side has limit 0 as $\epsilon \rightarrow 0$. Therefore the inequality contradicts independence of c and ϵ .

Proof for $0 < a < 1$.

We will use the logic in the proof for $a = 0$. However, this time, let u be at the center of B , and t be such that $B_{r/2}(t) \subset B$ where $r \geq 2\epsilon$. Then,

$$Q_r(u, t) = 2 \int_0^{\pi/6} \int_{r \cos \theta_{uv} - \frac{r}{2} \sqrt{1 - 4 \sin^2 \theta_{uv}}}^{r \cos \theta_{uv} + \frac{r}{2} \sqrt{1 - 4 \sin^2 \theta_{uv}}} \frac{1}{\pi} \frac{1}{\frac{2}{(1-a)} \left(1 - \frac{1}{\epsilon^{a-1}}\right)} \frac{1}{r_{uv}^a} (dr_{uv} d\theta_{uv})$$

Let $\theta_{uv}^+ \equiv \cos \theta_{uv} + \frac{1}{2} \sqrt{1 - 4 \sin^2 \theta_{uv}}$ and $\theta_{uv}^- \equiv \cos \theta_{uv} - \frac{1}{2} \sqrt{1 - 4 \sin^2 \theta_{uv}}$ for brevity. Then we have,

$$Q_r(u, t) = \frac{r^{1-a}}{\pi (1 - \epsilon^{1-a})} \int_0^{\pi/6} \{\theta_{uv}^{+1-a} - \theta_{uv}^{-1-a}\} d\theta_{uv} \leq \frac{2}{\pi} r^{1-a} A,$$

where $A = \int_0^{\pi/6} \{\theta_{uv}^{+1-a} - \theta_{uv}^{-1-a}\} d\theta_{uv}$. Thus, as in the $a = 0$ case, suppose $Q_r(u, t)$ is $\Omega(1/\log(2/\epsilon))$, i.e., there exists a constant $c > 0$ independent of ϵ such that $Q_r(u, t) \geq c/\log(2/\epsilon)$ for all r . That is, $\frac{2}{\pi} r^{1-a} A \geq Q_r(u, t) \geq c/\log(2/\epsilon)$ and specifically, $\frac{2^{2-a}}{\pi} \epsilon^{1-a} A \geq c/\log(2/\epsilon)$. This implies $c \leq \frac{A 2^{2-a}}{\pi} \epsilon^{1-a} \log(2/\epsilon)$. But the right hand side has limit 0 when $\epsilon \rightarrow 0$ contradicting independence of c and ϵ .

Proof for $a > 1$.

Let $a - 1 = b > 0$, and let the positions of u and t be as in the proof for $0 < a < 1$. Then after some algebraic simplifications we get,

$$Q_r(u, t) = \left(\frac{4}{3}\right)^b \frac{1}{\pi} \frac{r^{-b}}{\epsilon^{-b} - 1} \int_0^{\pi/6} ((\theta_{uv}^+)^b - (\theta_{uv}^-)^b) d\theta_{uv}.$$

Now let $\int_0^{\pi/6} ((\theta_{uv}^+)^b - (\theta_{uv}^-)^b) d\theta_{uv} = B$ and consider the case when $r = 1/2$. We get,

$$Q_{1/2}(u, t) = \left(\frac{8}{3}\right)^b B \frac{1}{\pi} \frac{1}{\left(\frac{1}{\epsilon}\right)^b - 1} \leq \left(\frac{8}{3}\right)^b \frac{2B}{\pi \left(\frac{1}{\epsilon}\right)^b}.$$

Suppose $Q_r(u, t)$ is $\Omega(1/\log(2/\epsilon))$, i.e., there exists a constant $c > 0$ independent of ϵ such that $Q_r(u, t) \geq c/\log(2/\epsilon)$ for all r . Specifically, when $r = 1/2$ as above, $\left(\frac{8}{3}\right)^b \frac{2B}{\pi \left(\frac{1}{\epsilon}\right)^b} \geq Q_{1/2}(u, t) \geq c/\log(2/\epsilon)$. This implies that $c \leq \left(\frac{8}{3}\right)^b \frac{2B\epsilon^b}{\pi} \log(2/\epsilon)$. Again, the right hand side has limit 0 as $\epsilon \rightarrow 0$ contradicting independence of c and ϵ . \square

4 Efficient Decentralized Search

We now return to the question posed in Section 2 and prove Theorems 2.5, 2.6 and 2.7. In particular, we will show that $a = 1$ is the only value of a for which there exists a decentralized algorithm that delivers the message efficiently. Consider the following greedy algorithm.

Algorithm G^* : When the message reaches a point $u \in B$, u forwards the message to its contacts that is closest to the target point t .

Proof of Theorem 2.5 Suppose that the current message holder is at a distance r , $2 \geq r \geq 2\epsilon$ from t . The probability that the distance between the message and t is decreased by a factor of 2 (i.e., a significant improvement) in one step of G^* is at least $Q_r(u, t) \geq \frac{c}{\log \frac{2}{\epsilon}}$ from Theorem 3.3. Therefore, the expected number of steps of G^* before the distance is decreased by a factor of 2 is at most $\frac{\log \frac{2}{\epsilon}}{c}$. Since the initial distance of the message from t can be at most 2, we need at most $\log_2(2/2\epsilon)$ such significant improvements to bring the message within distance 2ϵ from t . Once the message is within 2ϵ from t , we need at most 2 more steps to bring it within ϵ from t even if the message is passed on to local contacts in these steps. Thus, we need a total of at most $b(\log(2/\epsilon))^2$ expected steps of G^* to deliver the message within ϵ of t for some constant b .

End Proof of Theorem 2.5.

Proof of Theorem 2.6. Let t be at the center of B and s be such that $\|s - t\| = 1$. Let $y \equiv (2\epsilon)(1/\epsilon)^{(1-a)/2}$ and consider *any* decentralized algorithm. Let A be the event that the algorithm delivers the message starting at s to a point that has a long-range contact within the ball $B_y(t)$ in at most $T \equiv \frac{(3^{1-a}-1)}{16 \times 6^{1-a}} (1/\epsilon)^{(1-a)/2}$ steps.

Claim 1 $P(A) \leq (1/2)$.

Proof of Claim 1 We will first show that the probability that a point $u \in B$ has a long-range

contact in $B_y(t)$ is at most

$$8 \frac{6^{1-a}}{(3^{1-a} - 1)} \left(\frac{1}{\epsilon} \right)^{(a-1)/2}.$$

We use $P(u, B_y(t))$ to denote this probability and consider two cases : $B_\epsilon(u) \cap B_y(t) = \emptyset$ and $B_\epsilon(u) \cap B_y(t) \neq \emptyset$ illustrated in Figures 5 (a) and (b) respectively. For both these cases, Lemmas A.4 and A.5 imply that

$$\begin{aligned} P(u, B_y(t)) &\leq \frac{6^{1-a}}{\pi(3^{1-a} - 1)} \int_0^{2\pi} \int_\epsilon^{y_+} \frac{1}{r_{uv}^{a+1}} (r_{uv} dr_{uv} d\theta_{uv}) \leq \frac{6^{1-a}}{\pi(3^{1-a} - 1)} \int_0^{2\pi} \int_\epsilon^{y_+} \frac{1}{\epsilon^a} (dr_{uv} d\theta_{uv}) \\ &\leq \frac{6^{1-a}}{\pi(3^{1-a} - 1) \epsilon^a} \int_0^{2\pi} \int_\epsilon^{y_+} (dr_{uv} d\theta_{uv}) \leq \frac{6^{1-a}}{\pi(3^{1-a} - 1) \epsilon^a} \int_0^{2\pi} \int_\epsilon^{\epsilon+2y} (dr_{uv} d\theta_{uv}) = 4 \frac{6^{1-a} y}{(3^{1-a} - 1) \epsilon^a} \\ &= 8 \frac{6^{1-a}}{(3^{1-a} - 1)} \left(\frac{1}{\epsilon} \right)^{\frac{a-1}{2}}. \end{aligned}$$

Let A_i be the event that the long-range contact of the message holder in the i th step of the

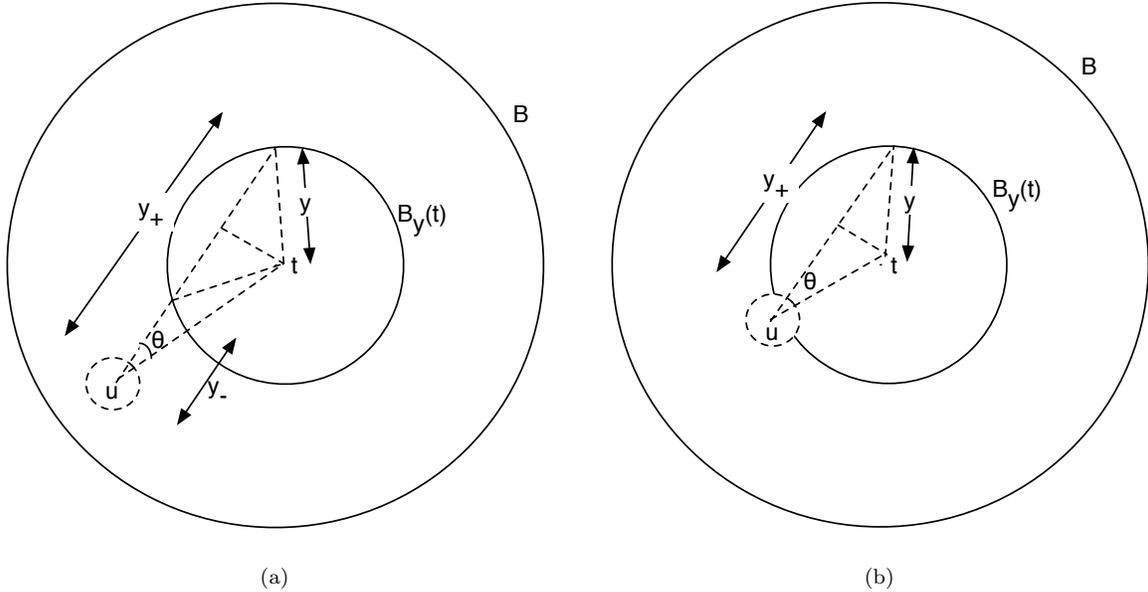


Figure 5: Illustration of proof of Claim 1 in Theorem 2.6. (a) $B_\epsilon(u) \cap B_y(t) = \emptyset$. (b) $B_\epsilon(u) \cap B_y(t) \neq \emptyset$.

algorithm is in $B_y(t)$. Note that $A = \bigcup_{i=1}^T A_i$. Therefore subadditivity implies that

$$\begin{aligned} P(A) &\leq \sum_{i=1}^T P(A_i) \leq \sum_{i=1}^T 8 \frac{6^{1-a}}{(3^{1-a} - 1)} \left(\frac{1}{\epsilon} \right)^{(a-1)/2} = 8T \frac{6^{1-a}}{(3^{1-a} - 1)} \left(\frac{1}{\epsilon} \right)^{(a-1)/2} \\ &= 8 \frac{6^{1-a}}{(3^{1-a} - 1)} \left(\frac{1}{\epsilon} \right)^{(a-1)/2} \frac{(3^{1-a} - 1)}{16 \times 6^{1-a}} \left(\frac{1}{\epsilon} \right)^{(1-a)/2} = \frac{1}{2}. \end{aligned}$$

End Proof of Claim 1

Now let B be the event that the message is ϵ -delivered to t in at most T steps.

Claim 2 B cannot occur if A does not occur.

Proof of Claim 2 Suppose not, i.e., A does not occur but the message is ϵ -delivered to t in at most T steps. Note that $T \leq \frac{1}{16} \left(\frac{1}{\epsilon}\right)^{\frac{(1-a)}{2}} \leq \left(\sqrt{\frac{1}{\epsilon}}\right)^{(1-a)} \ll \left(\frac{1}{\epsilon}\right)$. Therefore, a message that is ϵ -delivered from s to t in at most T steps must be passed to a long-range contact at least once. Moreover, the last time this happens, the long-range contact must be in $B_y(t)$ since $\frac{y}{\epsilon} = 2 \left(\frac{1}{\epsilon}\right)^{(1-a)/2} > T$. This contradicts our hypothesis that A does not occur.

End Proof of Claim 2

Now let $X_\epsilon(s, t)$ be the random number of steps required for the decentralized algorithm to ϵ -deliver the message from s to t . $E[X_\epsilon(s, t)] = E[X_\epsilon(s, t)|A]P(A) + E[X_\epsilon(s, t)|A^c]P(A^c)$. Hence, $E[X_\epsilon(s, t)] \geq E[X_\epsilon(s, t)|A^c]P(A^c) \geq \frac{1}{2}E[X_\epsilon(s, t)|A^c]$, where the last inequality follows from Claim 1. Claim 2 implies that $E[X_\epsilon(s, t)|A^c] > T$. Therefore we have,

$$E[X_\epsilon(s, t)] \geq \frac{1}{2}T \geq \frac{1}{2} \frac{(3^{1-a} - 1)}{16 \times 6^{1-a}} \left(\frac{1}{\epsilon}\right)^{(1-a)/2} \geq \frac{1}{32} \frac{(3^{1-a} - 1)}{12^{1-a}} \left(\sqrt{\frac{2}{\epsilon}}\right)^{1-a}$$

completing the proof of Theorem 2.6.

End Proof of Theorem 2.6.

Proof of Theorem 2.7. Let s and t be such that $\|s - t\| \geq 1$. Let $z \equiv \epsilon(1/\epsilon)^{1/a}$ and consider *any* decentralized algorithm. Let A_i be the event that the message holder in step i (say u) has a long-range contact v such that $\|u - v\| > z$. Let A be the event that this happens in at most $T \equiv \frac{1}{8} (1/\epsilon)^{(a-1)/a}$ steps, i.e., $A = \bigcup_{i=1}^T A_i$.

Claim 1 $P(A) \leq (1/2)$.

Proof of Claim 1 We will first show that the probability that a point $u \in B$ has a long-range contact v with $\|u - v\| \geq z$ is at most $4(\epsilon)^{(a-1)/a}$. Denote this probability by $P(u, S_z(u))$ where $S_z(u) = (B_z(u) \cap B)^c$. We consider two cases : u is z -deep in B , i.e., $B_z(u) \subset B$ and u is not z -deep in B illustrated in Figures 6(a) and (b) respectively where $z_+ \equiv \sqrt{1 - \|u\|^2 \sin^2 \theta} + \|u\| \cos \theta$. Lemmas A.6 and A.7 imply that in both these cases,

$$\begin{aligned} P(u, S_z(u)) &\leq \frac{2(a-1)\epsilon^{(a-1)}}{\pi} \int_0^{2\pi} \int_z^{z_+} \frac{1}{r_{uv}^{a+1}} (r_{uv} dr_{uv} d\theta_{uv}) \\ &= \frac{2(a-1)\epsilon^{(a-1)}}{\pi} \int_0^{2\pi} \int_z^{z_+} \frac{1}{r_{uv}^a} (dr_{uv} d\theta_{uv}) = \frac{2\epsilon^{(a-1)}}{\pi} \int_0^{2\pi} \left(\frac{1}{z^{(a-1)}} - \frac{1}{z_+^{(a-1)}} \right) d\theta_{uv} \\ &\leq \frac{2\epsilon^{(a-1)}}{\pi} \int_0^{2\pi} \left(\frac{1}{z^{(a-1)}} \right) d\theta_{uv} = 4 \left(\frac{\epsilon}{z} \right)^{(a-1)} = 4(\epsilon)^{(a-1)/a}. \end{aligned}$$

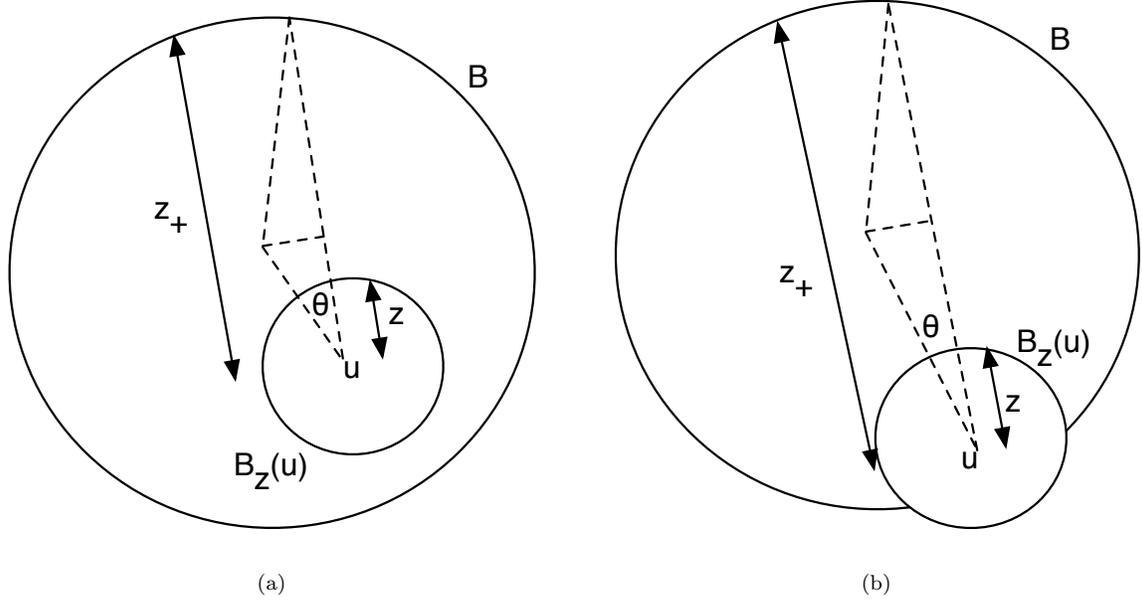


Figure 6: Illustration of proof of Claim 1 in Theorem 2.7. (a) u is z -deep in B , i.e., $B_z(u) \subset B$. (b) u is not z -deep in B .

Now subadditivity implies that

$$P(A) \leq \sum_{i=1}^T P(A_i) \leq \sum_{i=1}^T 4(\epsilon)^{(a-1)/a} = 4T(\epsilon)^{(a-1)/a} = 4\frac{1}{8}(1/\epsilon)^{(a-1)/a} (\epsilon)^{(a-1)/a} = \frac{1}{2}.$$

End Proof of Claim 1

Now let B be the event that the message is ϵ -delivered to t in at most T steps.

Claim 2 B cannot occur if A does not occur.

Proof of Claim 2 Since A does not occur, the message covers at most distance $z = \epsilon(1/\epsilon)^{1/a}$ in each of first $T = \frac{1}{8}(1/\epsilon)^{(a-1)/a}$ steps of the decentralized algorithm. Thus, the total distance covered is at most

$$\epsilon(1/\epsilon)^{1/a} \frac{1}{8}(1/\epsilon)^{(a-1)/a} = 1/8.$$

Thus, the message cannot be ϵ -delivered from s to t in at most T steps since $\|s - t\| = 1$.

End Proof of Claim 2

Now let $X_\epsilon(s, t)$ be the random number of steps required to ϵ -deliver the message from s to t . By calculations similar to the proof of Theorem 2.6 we get,

$$E[X_\epsilon(s, t)] \geq E[X_\epsilon(s, t)|A^c]P[A^c] \geq \frac{T}{2} \geq \frac{1}{16} \left(\frac{1}{\epsilon}\right)^{(a-1)/a} = \frac{1}{16} \left(\frac{1}{2}\right)^{\frac{a-1}{a}} \left(\frac{2}{\epsilon}\right)^{\frac{a-1}{a}}$$

completing the proof of Theorem 2.7.

End Proof of Theorem 2.7.

5 Conclusions

We constructed a family of Small World networks within a ball in R^2 using a Markov chain approach. In particular, we employed a family of Hit-and-Run algorithms parameterized by a single non-negative real parameter a . These network models were put in the context of Milgram's Small World experiments. Specifically, we argued that $a = 1$ is the unique value of this parameter for which there exists a decentralized algorithm that ϵ -delivers the message efficiently from any point in the world to any other point. The existence of an efficient decentralized algorithm was shown to be equivalent to two fundamental properties of the network models: scale invariance and steady improvement. Even though we focussed on the two-dimensional case, our results can be generalized to a ball in R^n in a straightforward manner. Informally, the capability of Hit-and-Run to generate a family of Small World networks on the fly may partly explain its recent emergence as the most efficient method to solve convex programs defined by oracles [6, 7].

A Preliminary Lemmas

Lemma A.1. *Let $u \in B$ and the local contacts of u be as defined in section 2. Let $t \in B \setminus B_\epsilon(u)$. Then there exists at least one local contact of u , say w , such that $\|u - t\| > \|w - t\|$.*

Proof. If u has a local contact w on the segment joining u and t then clearly $\|u - t\| > \|u - w\|$. Therefore we focus on the case where u does not have a local contact on segment ut . See Figure 7. Let the line joining u and t intersect $D_\epsilon(u)$ in point x . Since $\mathcal{P}(D_\epsilon(u))$ is a partition of $D_\epsilon(u)$, there exists some block V of the partition such that $x \in V$. Let $w \in V$ be a local contact of u . Also let y be the foot of the perpendicular from w to segment ux . Consider triangles uwy and uwt . Notice that since the geodesic diameter of V is at most $\pi/4$, and $w, x \in V$, $\angle wux \leq \pi/4$. Therefore, $\|u - y\| \geq \epsilon/\sqrt{2}$ and $\|w - y\| \leq \epsilon/\sqrt{2}$. We have,

$$\|u - t\| = \|u - y\| + \|y - t\| \geq (\epsilon/\sqrt{2}) + \|y - t\| \geq \|w - y\| + \|y - t\| > \|w - t\|,$$

where the last inequality is the triangle inequality for the Euclidean norm in R^n since w , y and t are not collinear. \square

We compute $C_{u,\epsilon}(\theta)$, i.e., the integral in the denominator of Equation (1) in the text here. Recall that point $u \in B$ is ϵ -deep inside B if $B_\epsilon(u) \subset B$.

Lemma A.2. *Let u be ϵ -deep in B . Then the values of $C_{u,\epsilon}(\theta) = \int_{\mu \in l_u(\theta)} \frac{1}{|\mu|^a} d|\mu|$ are as follows:*

$$\begin{aligned} & l_+ + l_- - 2\epsilon \quad \text{for } a = 0 \\ & 2 \log \left(\frac{\sqrt{l_+ \times l_-}}{\epsilon} \right) \quad \text{for } a = 1 \\ & \frac{1}{(1-a)} \left(\left[\frac{1}{l_+^{a-1}} - \frac{1}{\epsilon^{a-1}} \right] + \left[\frac{1}{l_-^{a-1}} - \frac{1}{\epsilon^{a-1}} \right] \right) \quad \text{for } a > 0, a \neq 1, \end{aligned}$$

where $l_+ \equiv (\sqrt{1 - \|u\|^2 \sin^2(\theta)} + \|u\| \cos(\theta))$ and $l_- \equiv (\sqrt{1 - \|u\|^2 \sin^2(\theta)} - \|u\| \cos(\theta))$ for brevity (see Figure 8(a)).

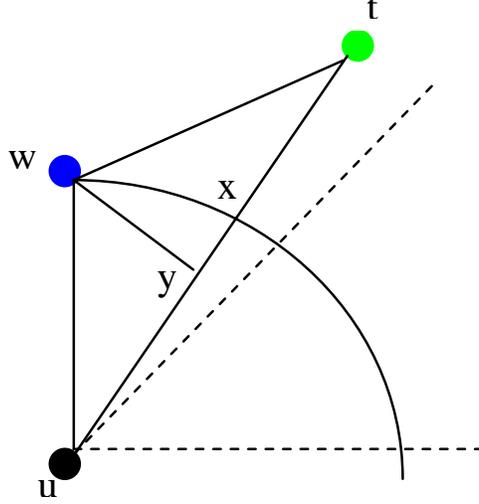


Figure 7: Illustration of proof of Lemma A.1. The line joining u and t intersects $D_\epsilon(u)$ at point x . $\angle wux = \theta \leq \pi/4$. The dotted lines show angle $\pi/4$ for comparison.

Proof. Let $a = 0$. Observe that

$$\int_{\mu \in l_u(\theta)} d|\mu| = \int_{\epsilon}^{l_+} d\mu + \int_{\epsilon}^{l_-} d\mu = l_+ + l_- - 2\epsilon$$

Let $a = 1$. Then,

$$\int_{\mu \in l(\theta)} \frac{1}{|\mu|} d|\mu| = \int_{\epsilon}^{l_+} \frac{1}{\mu} d\mu + \int_{\epsilon}^{l_-} \frac{1}{\mu} d\mu = \log(l_+) - \log(\epsilon) + \log(l_-) - \log(\epsilon) = 2 \log \left(\frac{\sqrt{l_+ \times l_-}}{\epsilon} \right).$$

Let $a > 0$, $a \neq 1$. We have,

$$\begin{aligned} \int_{\mu \in l(\theta)} \frac{1}{|\mu|^a} d|\mu| &= \int_{\epsilon}^{l_+} \frac{1}{\mu^a} d\mu + \int_{\epsilon}^{l_-} \frac{1}{\mu^a} d\mu = \frac{1}{(1-a)} \left(\left[\frac{1}{\mu^{a-1}} \right]_{\epsilon}^{l_+} + \left[\frac{1}{\mu^{a-1}} \right]_{\epsilon}^{l_-} \right) \\ &= \frac{1}{(1-a)} \left(\left[\frac{1}{l_+^{a-1}} - \frac{1}{\epsilon^{a-1}} \right] + \left[\frac{1}{l_-^{a-1}} - \frac{1}{\epsilon^{a-1}} \right] \right). \end{aligned}$$

□

This lemma can be adapted to the case when u is not ϵ -deep in B . The calculations become more complicated. Toward that end, let y and z be the points of intersection of $\partial B_\epsilon(u)$ and B and let α be the angle between $y - u$ and $-u$ as well as between $z - u$ and $-u$. See Figure 8(b).

Lemma A.3. *Suppose that u is not ϵ -deep in B , i.e., $1 - \epsilon < \|u\| < 1$ and let α be as defined above. (Notice that $0 < \alpha \leq \pi$, and in fact, it is not hard to prove using the cosine rule that*

$\alpha > \pi/2$ if $1 - \epsilon < \|u\| \leq \sqrt{1 - \epsilon^2}$, and $\alpha \leq \pi/2$ if $\|u\| > \sqrt{1 - \epsilon^2}$). Let l_+ and l_- be as in Lemma A.2. Then the values of $C_{u,\epsilon}(\theta) = \int_{\mu \in l_u(\theta)} \frac{1}{|\mu|^a} d\mu$ are as follows:

For $a = 0$,

1. $\alpha \leq \pi/2$.

$$\begin{aligned} l_+ - \epsilon & \text{ for } 0 \leq \theta \leq \alpha \\ l_u(\theta) = \emptyset & \text{ for } \alpha < \theta \leq \pi - \alpha \\ l_- - \epsilon & \text{ for } \pi - \alpha < \theta \leq \pi + \alpha \\ l(\theta) = \emptyset & \text{ for } \pi + \alpha < \theta \leq 2\pi - \alpha \\ l_+ - \epsilon & \text{ for } 2\pi - \alpha < \theta \leq 2\pi \end{aligned}$$

2. $\alpha > \pi/2$.

$$\begin{aligned} l_+ - \epsilon & \text{ for } 0 \leq \theta \leq \pi - \alpha \\ l_+ + l_- - 2\epsilon & \text{ for } \pi - \alpha < \theta \leq \alpha \\ l_- - \epsilon & \text{ for } \alpha < \theta \leq 2\pi - \alpha \\ l_+ + l_- - 2\epsilon & \text{ for } 2\pi - \alpha < \theta \leq \pi + \alpha \\ l_+ - \epsilon & \text{ for } \pi + \alpha < \theta \leq 2\pi \end{aligned}$$

For $a = 1$,

1. $\alpha \leq \pi/2$.

$$\begin{aligned} \log(l_+/\epsilon) & \text{ for } 0 \leq \theta \leq \alpha \\ l_u(\theta) = \emptyset & \text{ for } \alpha < \theta \leq \pi - \alpha \\ \log(l_-/\epsilon) & \text{ for } \pi - \alpha < \theta \leq \pi + \alpha \\ l_u(\theta) = \emptyset & \text{ for } \pi + \alpha < \theta \leq 2\pi - \alpha \\ \log(l_+/\epsilon) & \text{ for } 2\pi - \alpha < \theta \leq 2\pi \end{aligned}$$

2. $\alpha > \pi/2$.

$$\begin{aligned} \log(l_+/\epsilon) & \text{ for } 0 \leq \theta \leq \pi - \alpha \\ 2\log(\sqrt{(l_+ \times l_-)}/\epsilon) & \text{ for } \pi - \alpha < \theta \leq \alpha \\ \log(l_-/\epsilon) & \text{ for } \alpha < \theta \leq 2\pi - \alpha \\ 2\log(\sqrt{(l_+ \times l_-)}/\epsilon) & \text{ for } 2\pi - \alpha < \theta \leq \pi + \alpha \\ \log(l_+/\epsilon) & \text{ for } \pi + \alpha < \theta \leq 2\pi \end{aligned}$$

For $a > 0$, $a \neq 1$.

1. $\alpha \leq \pi/2$.

$$\begin{aligned}
& \frac{1}{(1-a)} \left(\left\{ \frac{1}{l_+^{a-1}} - \frac{1}{\epsilon^{a-1}} \right\} \right) \quad \text{for } 0 \leq \theta \leq \alpha \\
& \quad \quad \quad l_u(\theta) = \emptyset \quad \text{for } \alpha < \theta \leq \pi - \alpha \\
& \frac{1}{(1-a)} \left(\left\{ \frac{1}{l_-^{a-1}} - \frac{1}{\epsilon^{a-1}} \right\} \right) \quad \text{for } \pi - \alpha < \theta \leq \pi + \alpha \\
& \quad \quad \quad l_u(\theta) = \emptyset \quad \text{for } \pi + \alpha < \theta \leq 2\pi - \alpha \\
& \frac{1}{(1-a)} \left(\left\{ \frac{1}{l_+^{a-1}} - \frac{1}{\epsilon^{a-1}} \right\} \right) \quad \text{for } 2\pi - \alpha < \theta \leq 2\pi
\end{aligned}$$

2. $\alpha > \pi/2$.

$$\begin{aligned}
& \frac{1}{(1-a)} \left(\left\{ \frac{1}{l_+^{a-1}} - \frac{1}{\epsilon^{a-1}} \right\} \right) \quad \text{for } 0 \leq \theta \leq \pi - \alpha \\
& \frac{1}{(1-a)} \left(\left\{ \frac{1}{l_+^{a-1}} - \frac{1}{\epsilon^{a-1}} \right\} + \left\{ \frac{1}{l_-^{a-1}} - \frac{1}{\epsilon^{a-1}} \right\} \right) \quad \text{for } \pi - \alpha < \theta \leq \alpha \\
& \frac{1}{(1-a)} \left(\left\{ \frac{1}{l_-^{a-1}} - \frac{1}{\epsilon^{a-1}} \right\} \right) \quad \text{for } \alpha < \theta \leq 2\pi - \alpha \\
& \frac{1}{(1-a)} \left(\left\{ \frac{1}{l_+^{a-1}} - \frac{1}{\epsilon^{a-1}} \right\} + \left\{ \frac{1}{l_-^{a-1}} - \frac{1}{\epsilon^{a-1}} \right\} \right) \quad \text{for } 2\pi - \alpha < \theta \leq \pi + \alpha \\
& \frac{1}{(1-a)} \left(\left\{ \frac{1}{l_+^{a-1}} - \frac{1}{\epsilon^{a-1}} \right\} \right) \quad \text{for } \pi + \alpha < \theta \leq 2\pi
\end{aligned}$$

Proof. Very similar to Lemma A.2 hence omitted. \square

Lemma A.4. Let $0 \leq a < 1$. Let $y \equiv 2\epsilon(1/\epsilon)^{(1-a)/2}$. Let $u \in B$ be such that $r_u \geq y$. Then for $0 \leq \theta \leq \arcsin(y/r_u)$,

$$C_{u,\epsilon}(\theta) \geq \frac{(3^{1-a} - 1)}{6^{1-a}}.$$

Proof. Without loss of generality we assume that ϵ is small enough such that $\epsilon < (\sqrt{3/16})^{1/(1+a)}$ and $\epsilon < 1/6$. From Lemma A.2 and Lemma A.3,

$$\begin{aligned}
C_{u,\epsilon}(\theta) & \geq \frac{1}{1-a} \left(\sqrt{1 - \|u\|^2 \sin^2 \theta}^{1-a} - \epsilon^{1-a} \right) \\
& \geq \left(\sqrt{1 - y^2}^{1-a} - \epsilon^{1-a} \right) \\
& = \left(\sqrt{1 - 4\epsilon(\epsilon)^a}^{1-a} - \epsilon^{1-a} \right) \geq \left(\sqrt{1 - 3/4}^{1-a} - \epsilon^{1-a} \right) \\
& = \left((1/2)^{1-a} - \epsilon^{1-a} \right) \geq \left((1/2)^{1-a} - (1/6)^{1-a} \right) = \frac{(3^{1-a} - 1)}{6^{1-a}}
\end{aligned}$$

\square

Lemma A.5. Let $0 \leq a < 1$. Let $y \equiv 2\epsilon(1/\epsilon)^{(1-a)/2}$ as in Lemma A.4. Let $u \in B$ be such that $r_u < y$. Then for $0 \leq \theta \leq 2\pi$,

$$C_{u,\epsilon}(\theta) \geq \frac{(3^{1-a} - 1)}{6^{1-a}}.$$

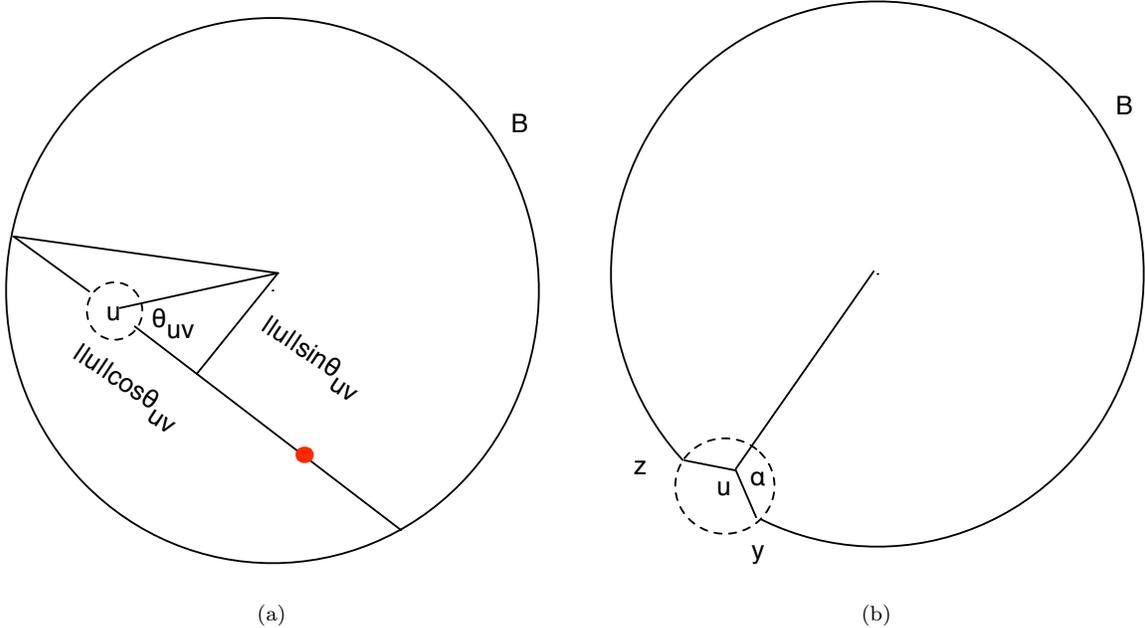


Figure 8: (a) Illustration of calculations in Lemma A.2. (b) Illustration of angle α when u is not ϵ deep in B as in Lemma A.3.

Proof. Very similar to Lemma A.4 hence omitted. \square

Lemma A.6. *Let $a > 1$. Let $z \equiv \epsilon(1/\epsilon)^{1/a}$. Then for any $u \in B$ with $B_z(u) \subseteq B$ (i.e., u is z -deep), $C_{u,\epsilon}(\theta) \geq \frac{1}{2(a-1)} \frac{1}{\epsilon^{(a-1)}}$ for $0 \leq \theta \leq 2\pi$.*

Proof. See Figure 9 (a). Since u is z -deep, it is also ϵ -deep since $z > \epsilon$. From Lemma A.2 we have,

$$\begin{aligned} C_{u,\epsilon}(\theta) &= \frac{1}{a-1} \left(\frac{2}{\epsilon^{(a-1)}} - \left(\frac{1}{l_+^{(a-1)}} + \frac{1}{l_-^{(a-1)}} \right) \right) \geq \frac{1}{a-1} \left(\frac{2}{\epsilon^{(a-1)}} - \frac{2}{z^{a-1}} \right) \\ &= \frac{2}{(a-1)\epsilon^{(a-1)}} \left(1 - \frac{1}{(1/\epsilon)^{(a-1)/a}} \right) \geq \frac{1}{(a-1)} \frac{1}{\epsilon^{(a-1)}}, \end{aligned}$$

where we have assumed without loss of generality that $(1/\epsilon)^{(a-1)/a} > 2$ to get the last inequality. \square

Lemma A.7. *Let $a > 1$ and $z \equiv \epsilon(1/\epsilon)^{1/a}$ as in Lemma A.6. Suppose $u \in B$ is not z -deep and angle β be as shown in Figure 9 (b). Then $C_{u,\epsilon}(\theta) \geq \frac{1}{2(a-1)} \frac{1}{\epsilon^{(a-1)}}$ for $0 \leq \theta \leq \beta$, and $2\pi - \beta \leq \theta \leq 2\pi$.*

Proof. Observe that since u is not z -deep, it may not be ϵ -deep. Recall the definition of α in Lemma A.3 and note that since $z > \epsilon$, $\alpha \geq \beta$. For any θ that satisfies the range in the hypothesis, we have,

$$C_{u,\epsilon}(\theta) \geq \frac{1}{(a-1)} \left(\frac{1}{\epsilon^{(a-1)}} - \frac{1}{z^{(a-1)}} \right) = \frac{1}{(a-1)\epsilon^{(a-1)}} \left(1 - \frac{1}{(1/\epsilon)^{(a-1)/a}} \right) \geq \frac{1}{2(a-1)} \frac{1}{\epsilon^{(a-1)}},$$

where we have assumed without loss of generality that $(1/\epsilon)^{(a-1)/a} > 2$ as in Lemma A.6 to get the last inequality. \square

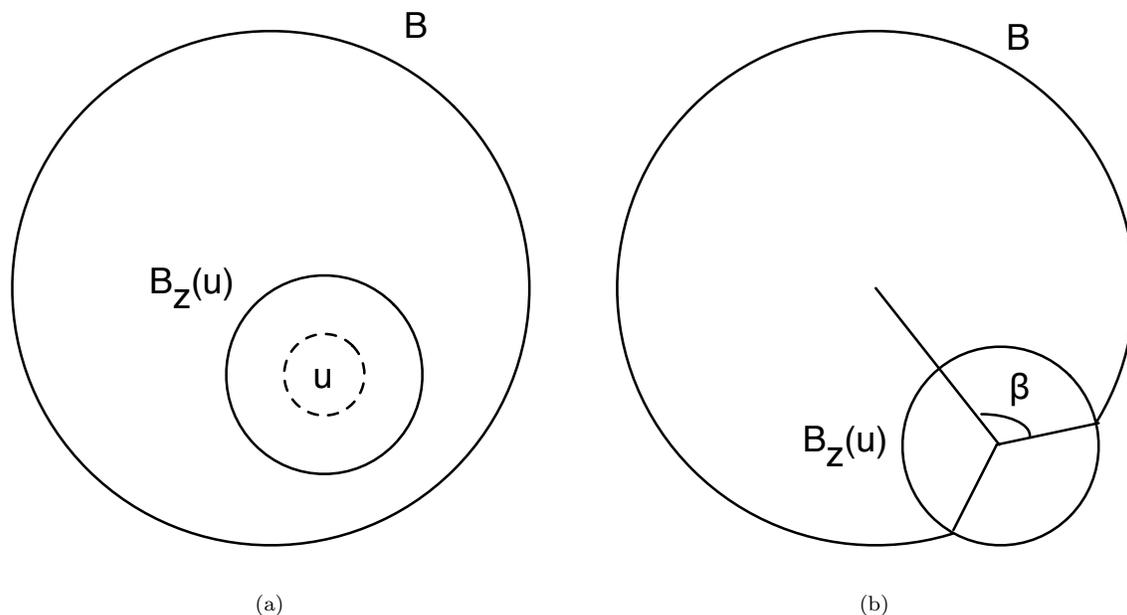


Figure 9: Illustration of proof of Lemma A.7. (a) Point u shown as a black dot is z -deep (hence ϵ -deep) in B . (b) u is not z -deep in B .

References

- [1] Albert, R. and Barabasi A. L., Statistical mechanics of complex networks, 2002, *Rev. Modern Phys.*, 74, 4797.
- [2] Barabasi, A. L. and Albert, R., Emergence of scaling in random networks, *Science*, 286:509–512, 1999.
- [3] Belisle, C. J. P., Convergence Theorems for a Class of Simulated Annealing Algorithms on R^d , *Journal of Applied Probability*, 29:885–895, 1992.
- [4] Belisle, C.J.P., Romeijn, H.E., and Smith, R. L., Hit-and-run algorithms for generating multivariate distributions, *Mathematics of Operations Research*, 18:255–266, 1993.
- [5] Berbee, H.C.P., Boender, C.G.E., Rinnooy Kan, A.H.G., Scheffer, C.L., Smith, R.L. and Telgen, J., Hit-and-run algorithm for the identification of nonredundant linear inequalities, *Mathematical Programming*, 37:184–270, 1987.
- [6] Bertsimas, D., and Vempala, S., Solving convex programs by random walks, *Journal of the ACM (JACM)*, 51(4): 540–556, 2004.
- [7] Kalai, A., and Vempala, S., Simulated Annealing for convex optimization, *Mathematics of Operations Research*, 31 (2), 253-266, May 2006.
- [8] Kleinberg, J. M., Navigation in a small world, *Nature*, 406, 845 2000.
- [9] Kleinberg, J. M., The small world phenomenon: An algorithmic perspective, Proceedings of the 32nd Annual ACM Symposium on Theory of Computing, 163-170, 2000.

- [10] Kleinberg, J. M., Small World Phenomena and the Dynamics of Information, Advances in Neural Information Processing Systems, 2001.
- [11] Kleinberg, J. M., Complex Networks and Decentralized Search Algorithms, Proceedings of the International Congress on Mathematics, 2006.
- [12] Lovasz, L., Hit-and-run mixes fast, *Mathematical Programming, Series A*, 86:443–461, 1999.
- [13] Lovasz, L., and Vempala, S., Hit-and-run is fast and fun, *Microsoft Research Tech. Rep. MSR-TR-2003-05*. Available at <http://wwwmath.mit.edu/vempala/papers/logcon-hitrun.ps>.
- [14] Lovasz, L., and Vempala, S., Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithm, Proc. of the 44th IEEE Foundations of Computer Science (FOCS '03), Boston, 2003.
- [15] Lovasz, L., and Vempala, S., Hit-and-run from a corner, Proc. of the 36th ACM Symposium on the Theory of Computing (STOC '04), Chicago, 2004. To appear in SIAM J. Computing. Available at <http://www-math.mit.edu/vempala/papers/start.pdf>
- [16] Milgram, S., The small world problem, *Psychology Today*, 2, 60-67, 1967.
- [17] Newman, M. E. J., Models of the small world, *J. Statis. Phys.*, 101, 819-941, 2000.
- [18] Newman, M. E. J., The structure and function of complex networks, *SIAM Review*, 45(2), 167-256, 2003.
- [19] Posamentier, A. S., Excursions in advanced Euclidean geometry, Addison-Wesley, Menlo Park, California, 1984.
- [20] Reaume, D. J., Romeijn, H. E., and Smith, R. L., Implementing pure adaptive search for global optimization, *Journal of Global Optimization*, 20(1):33–47, 2001.
- [21] Romeijn, H. E., and Smith, R. L., Simulated annealing for constrained global optimization, *Journal of Global Optimization*, 5:101–126, 1994.
- [22] Smith, R. L., Efficient Monte Carlo procedures for generating points uniformly distributed over bounded regions, *Operations Research*, 32:1296–1308, 1984.
- [23] Travers, J., and Milgram, S., An experimental study of the small world problem, *Sociometry*, 32, 425-443, 1969.
- [24] Watts, D., and Strogatz, S., Collective dynamics of small world networks, *Nature*, 393, 440-442, 1998.
- [25] Watts, D., Dodds, P. S., and Newman, M. E. J., Identity and search in social networks, *Science*, 296, 1302-1305, 2002.
- [26] Zabinsky, Z. B., Smith, R. L., McDonald, J. F., Romeijn, H. E., and Kaufman, D. E., Improving hit-and-run for global optimization, *Journal of Global Optimization*, 3:171–192, 1993.