

DECENTRALIZED SEARCH ON SPHERES USING SMALL-WORLD MARKOV CHAINS: EXPECTED HITTING TIMES AND STRUCTURAL PROPERTIES

ARCHIS GHATE,* *University of Washington*

Abstract

We build a family of Markov chains on a sphere using distance-based long-range connection probabilities to model the decentralized message-passing problem that has recently gained significant attention in the small-world literature. Starting at an arbitrary source point on the sphere, the expected message delivery time to an arbitrary target on the sphere is characterized by a particular expected hitting time of our Markov chains. We prove that, within this family, there is a unique efficient Markov chain whose expected hitting time is polylogarithmic in the relative size of the sphere. For all other chains, this expected hitting time is at least polynomial. We conclude by defining two structural properties, called scale invariance and steady improvement, of the probability density function of long-range connections and prove that they are sufficient and necessary for efficient decentralized message delivery.

Keywords: Small-world problem; Markov chain

2000 Mathematics Subject Classification: Primary 60J05

Secondary 68W40; 91D30

1. Introduction

The *small-world* phenomenon was observed in the 1960s by Stanley Milgram's team [10], [13], who realized that apparent strangers are often connected by a short chain of acquaintances. This striking discovery was the result of a simple letter-forwarding experiment that started out by identifying a pair of people in the United States, unknown to each other. One of them, called the *source*, was given basic information such as the name and address of the other, termed the *target*. A letter was handed to the source and she was told that it 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 astounding observation was popularized as *six degrees of separation*.

While social experiments of this nature continue to amaze experts and nonexperts alike, several mathematicians, computer scientists, physicists, biologists, and economists have been drawn to the small-world phenomenon [1], [2], [7], [8], [9], [11], [14], [15] perhaps due to the empirical evidence that a multitude of social, natural, economic, and physical systems [12] possess the key property that they are highly clustered, yet exhibit a small diameter; indeed, a

Received 13 April 2007; revision received 30 May 2008.

* Postal address: Industrial Engineering, University of Washington, Box 352650, Seattle, WA 98195-2650, USA.

Email address: archis@u.washington.edu

friend of your friend is very likely to be your friend, and yet, you are connected to strangers by an acquaintance chain of length six on average.

The string of quantitative research in this area was pioneered by Watts and Strogatz [14] who showed that it is possible to construct networks where arbitrary nodes are connected by a short chain of edges. They achieved this by superimposing a small number of uniformly random ‘long-range connections’ on a ring lattice where each node originally has a few ‘local contacts’. However, Kleinberg [7] was the first to note that the most fascinating component of Milgram’s social experiments is not that short chains of acquaintances *exist* between strangers, but that people are able to actually *find* them by using a simple letter-forwarding procedure. Specifically, 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 [7]. In his seminal work, Kleinberg [7], [8] 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. Such networks are often termed *navigable*. This approach goes beyond the mere existence of short chains in that it attempts to actually find them, and, hence, is algorithmic. In Kleinberg’s grid 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 for which a decentralized algorithm can find short chains.

Following Kleinberg’s idea, researchers have continued employing some notion of distance while modeling Milgram’s message-passing experiment using decentralized algorithms. A recent important departure from the above modeling framework was presented by Liben-Nowell *et al.* [9], who employed rank-based connection probabilities (where the probability that person v is person u ’s long-range contact is inversely proportional to the number of persons w that are closer to u than the distance between u and v). Another interesting paradigm shift in this area includes models in Euclidean spaces. That is, instead of using the more traditional graph theoretic models like the ring lattice or the grid lattice, researchers have embedded their models in a plane [3], a square [4], or a ball [5], [6] in \mathbb{R}^2 . In some cases such as [3], these continuous-domain models do not explicitly build a network by adding random long-range contacts to an existing finite graph defined by local contacts (sometimes called the ‘substrate’), but rather model the message-passing procedure directly as a stochastic process characterized by locations of local and long-range contacts in a continuum starting at an arbitrarily selected source point.

In this paper we model the message-passing experiment using a family of Markov chains we call *small-world Markov chains* defined on the surface of a sphere. In this model, the message is initially at source s in the sphere and is to be delivered to a location close to a target point t also in the sphere (note that exact delivery is impossible in continuous-domain models [3]). When the message is at point u , it is passed on either to one of u ’s four local contacts or to its long-range contact. Given the geodesic circle through u on which u ’s long-range contact lies, the probability density function (see (2.1)) for the geodesic distance ρ between u and its long-range contact is proportional to ρ^{-a} for a nonnegative parameter a , akin to distance-based power-law models in the literature [3], [4], [5], [6], [7], [8]. Point u is only assumed to know the locations of all of its contacts (local as well as long range) and the location of target t . As a result, u ’s choice of where to send the message is entirely decentralized and the location of the message on the sphere is a Markov chain. The key quantity of interest is the expected hitting time of the message (or, equivalently, the Markov chain that defines the message location) to a small area around target t . We show in Theorem 2.1 that this expected hitting time is at most

polylogarithmic in the relative size of the sphere if and only if $a = 1$. On the other hand, for all other nonnegative values of a , this expected hitting time is at least polynomial in the relative size of the sphere (Theorem 2.2). Thus, the model with $a = 1$ performs *exponentially better* than all the other models. We also prove that $a = 1$ is the only power-law exponent that imparts two structural properties to the probability density of long-range contacts—scale-invariance (Proposition 4.1) and significant improvement (Proposition 4.2).

2. A family of message-forwarding models

We build our models on the surface of a sphere of radius R , denoted S_R . Choose an arbitrary point $s \in S_R$ called the source and an arbitrary point $t \in S_R$ called the target. As is common in the literature, our family of models have three main components: local as well as long-range contacts and a decentralized message delivery procedure. Suppose that the message is currently at point $u \in S_R$.

2.1. Local contacts

Let $\varepsilon > 0$ be a fixed constant such that $R/\varepsilon \gg 1$, and let $G_{\pi\varepsilon}(u)$ be the set of points in S_R that are at a geodesic distance less than or equal to $\pi\varepsilon$ from point u . Moreover, let $\partial G_{\pi\varepsilon}(u)$ be the circle of points in S_R that are at a geodesic distance $\pi\varepsilon$ from u . Point u has four local contacts equally spaced, i.e. at an angular separation $\pi/2$ each, on $\partial G_{\pi\varepsilon}(u)$. It will soon be clear that using a radius of $\pi\varepsilon$ instead of ε is without loss of generality and only motivated by the resulting simplicity of several algebraic expressions in the sequel. The model can easily be generalized to more than four local contacts; however, we focus on four local contacts for concreteness. It is not necessary for the four local contacts to be equally spaced; however, we do need them to be placed so that the conclusion of Lemma A.1 holds. Nevertheless, the assumption of equally spaced deterministically generated local contacts is consistent with models in the literature [7], [8]. We use $v_1(u)$, $v_2(u)$, $v_3(u)$, and $v_4(u)$ to denote the locations of the four local contacts of u on $\partial G_{\pi\varepsilon}(u)$.

2.2. Long-range contacts

In addition to the above four local contacts, u has one long-range contact placed on S_R outside $G_{\pi\varepsilon}(u)$ placed using geodesic distance-based probabilities parameterized by a nonnegative real number a . It is placed on a uniformly chosen great circle passing through u , at geodesic distance ρ from u , independent of the choice of the great circle and having density

$$f(\rho) = \frac{1}{C(a; R, \varepsilon)\rho^a}, \quad \pi\varepsilon < \rho \leq \pi R, \tag{2.1}$$

with the normalization constant

$$C(a; R, \varepsilon) = \int_{\pi\varepsilon}^{\pi R} \frac{1}{s^a} ds = \begin{cases} \log\left(\frac{R}{\varepsilon}\right) & \text{if } a = 1, \\ \frac{\pi^{1-a}}{1-a}(R^{1-a} - \varepsilon^{1-a}) & \text{if } a \geq 0, a \neq 1. \end{cases} \tag{2.2}$$

As the notation suggests, $C(a; R, \varepsilon)$ does not depend on the choice of the great circle. The parameter a will be called ‘the power-law exponent’. We use $v_5(u) \in \{S_R \setminus G_{\pi\varepsilon}(u)\}$ to denote u ’s long-range contact.

2.3. Decentralized search algorithm

A decentralized search algorithm is a message-passing procedure where message holder u only knows

1. the location of t ; and
2. the locations of its own local and long-range contacts, i.e. $v_1(u), v_2(u), v_3(u), v_4(u)$, and $v_5(u)$.

Most importantly, a message holder *does not* have any information about the local and long-range contacts of other points, and, hence, must forward the message to one of *its own* contacts. A message holder does not know where the message has been in the past. For notational simplicity, we focus on deterministic decentralized algorithms that can be characterized by functions $g_t : S_R^5 \rightarrow S_R$ with the property that $g_t(v_1(u), v_2(u), v_3(u), v_4(u), v_5(u)) \in \{v_1(u), v_2(u), v_3(u), v_4(u), v_5(u)\}$. That is, the function value is just one of the contact points. Specifically, if the message reaches point u with contacts $v_1(u), v_2(u), v_3(u), v_4(u)$, and $v_5(u)$, it is forwarded to contact point $g_t(v_1(u), v_2(u), v_3(u), v_4(u), v_5(u))$. We will use the terms ‘function g_t ’ and ‘decentralized algorithm g_t ’ interchangeably. We will say that the message has been ε -delivered when it reaches a point within geodesic distance ε from t , i.e. when it reaches the set $G_{\pi\varepsilon}(t)$.

Let X_n denote the location of the message after it has been forwarded n times using a decentralized algorithm. Then X_n is a Markov chain (we call this the small-world Markov chain) with state space S_R . Given the source s and target t , this Markov chain is parameterized by a and the function g_t . The main quantity of interest, the expected ε -delivery time, is the expected value $E(\tau_{s,t}(a, g_t))$ of the hitting time $\tau_{s,t}(a, g_t)$ of this Markov chain to the set $G_{\pi\varepsilon}(t)$, i.e.

$$\tau_{s,t}(a, g_t) = \min\{n \geq 1 : X_n \in G_{\pi\varepsilon}(t)\}.$$

Research question. For what values of parameter a (if any) does there exist a decentralized message-passing algorithm, i.e. a function g_t with the above-defined properties, such that

$$E(\tau_{s,t}(a, g_t)) \leq c \left(\log \left(\frac{R}{\varepsilon} \right) \right)^b \tag{2.3}$$

for any source $s \in S_R$ and any target $t \in S_R$, where b and c are constants that do not depend on R, ε, s , and t ?

A decentralized message-passing algorithm that results in a small-world Markov chain with a bound as in (2.3) on its expected hitting time will be called *polylogarithmic*. We call such a polylogarithmic algorithm *efficient*. Note that polylogarithmic performance is an appropriate measure of efficiency in this context since a model with no long-range contacts can trivially deliver the message from a source to a target in a number of steps proportional to R/ε . Consequently, the quantity R/ε is of key interest in the remainder of the paper and can be interpreted as the relative size of the sphere, denoted L for brevity.

Consider the following greedy decentralized algorithm g_t^* : whenever the message reaches a point u , u forwards the message to one of its contacts that is closest to the target point t (breaking ties arbitrarily). For example, when ties are broken by the smallest index rule, the greedy decentralized algorithm is characterized by a function g_t^* with the property that, when the message reaches a point u , $g_t^*(v_1(u), v_2(u), v_3(u), v_4(u), v_5(u)) = v_i(u)$, where

$$i = \min\{j : d(v_j(u), t) = \min\{d(v_1(u), t), d(v_2(u), t), d(v_3(u), t), d(v_4(u), t), d(v_5(u), t)\}\}$$

and $d(x, y)$ denotes the geodesic distance between $x, y \in S_R$. Lemma A.1 in Appendix A shows that our greedy algorithm is well defined in the sense that point u has at least one local contact that is closer to the target than u itself. This has the important implication that, under g_t^* , the message is moved from one point to another (sometimes the negligible distance ε) at every step and, in particular, never gets stuck. In the next section we prove the following key theorems that answer the above research question.

Theorem 2.1. *For any $s, t \in S_R$, $E(\tau_{s,t}(1, g_t^*)) \leq c(\log L)^2$, where c is a constant that does not depend on R, ε, s , or t . That is, for any source s and target t on the sphere S_R , the expected ε -delivery time of the greedy algorithm g_t^* is at most polylogarithmic in $L = R/\varepsilon$ if $a = 1$.*

Theorem 2.2. *If $a \neq 1$, there exist $s, t \in S_R$ for which the expected ε -delivery time of every decentralized algorithm is at least polynomial in $L = R/\varepsilon$. In particular, the following statements hold.*

- (i) *If $0 \leq a < 1$, there exist $s, t \in S_R$ for which $E(\tau_{s,t}(a, g_t)) \geq \alpha L^{(1-a)/2}$ for every decentralized algorithm g_t , where α is a constant that does not depend on R, ε , or g_t .*
- (ii) *If $a > 1$, there exist $s, t \in S_R$ for which $E(\tau_{s,t}(a, g_t)) \geq \beta L^{(a-1)/a}$ for every decentralized algorithm g_t , where β is a constant that does not depend on R, ε , or g_t .*

3. Proofs of expected hitting time bounds

The two key theorems in Section 2 are proven here. The preliminary results required for these proofs are deferred to Appendix A. To geometrically indicate the location of point u 's long-range contact $v_5(u)$, we use the geodesic coordinate system (θ, ρ) centered at u . Here θ is the angle between two great circles passing through u —one through the north pole and the other through $v_5(u)$. Also, ρ is the geodesic distance between u and $v_5(u)$. Recall that $v_5(u)$ is determined by randomly selecting both θ and ρ —the former uniformly from $[0, 2\pi]$ while the latter according to the density in (2.1). We begin with a simple lemma.

Lemma 3.1. *Let $u \in S_R$, let $\pi\varepsilon \leq r_1 < r_2 \leq \pi R$, and let $0 \leq \theta_1 < \theta_2 \leq \pi$. Let S denote the spherical segment between geodesic distances r_1 and r_2 from u and the great circles at spherical angles θ_1 and θ_2 with respect to the great circle passing through u and the north pole. Let $P(u, S)$ be the probability that u 's long-range contact $v_5(u)$ is in S . Then*

$$P(u, S) = \frac{\theta_2 - \theta_1}{2\pi} \int_{r_1}^{r_2} \frac{1}{C(a; R, \varepsilon)\rho^a} d\rho.$$

Proof. The probability of interest is given by

$$\begin{aligned} P(u, S) &= P(\theta_1 \leq \theta \leq \theta_2 \cap r_1 \leq \rho \leq r_2) \\ &= P(\theta_1 \leq \theta \leq \theta_2) P(r_1 \leq \rho \leq r_2) \quad (\text{since } \theta \text{ and } \rho \text{ are chosen independently}) \\ &= \frac{\theta_2 - \theta_1}{2\pi} P(r_1 \leq \rho \leq r_2) \quad (\text{since } \theta \text{ is chosen uniformly from } [0, 2\pi]) \\ &= \frac{\theta_2 - \theta_1}{2\pi} \int_{r_1}^{r_2} \frac{1}{C(a; R, \varepsilon)\rho^a} d\rho \quad (\text{from (2.1)}). \end{aligned}$$

This completes the proof.

3.1. Proof of Theorem 2.1

For this proof, we use a geometric trials argument. Suppose that the current message holder u is at a geodesic distance r , $\pi R \geq r \geq 2\pi\epsilon$, from t . Let $P_r(u, t; 1)$ be the probability that u 's long-range contact is within a geodesic distance $r/2$ from t , where the 1 emphasizes that we are considering the case in which $a = 1$. Equation (2.2), Lemma 3.1, Corollary A.1 in Appendix A, and spherical symmetry imply that

$$P_r(u, t; 1) \geq 2 \frac{\pi/12 - 0}{2\pi} \int_{3r/4}^r \frac{1}{(\log L)\rho} d\rho = \frac{\log(4/3)}{12 \log L}.$$

Therefore, the probability that the distance between the message and t is decreased by a factor of 2 in one step of g_t^* is at least $\log(4/3)/12 \log L$. As a result, the expected number of steps of g_t^* before the distance is decreased by a factor of 2 is at most $12 \log L / \log(4/3)$. Since the initial distance of the message from t can be at most πR , we need at most $\log_2 L$ such improvements to bring the message within distance $2\pi\epsilon$ from t . Once the message is within $2\pi\epsilon$ from t , we need at most two more steps to bring it within $\pi\epsilon$ from t , even if the message is passed on to local contacts in these steps. Thus, we need a total of at most $c(\log L)^2$ expected number of steps of g_t^* to ϵ -deliver the message to t for a constant c independent of R and ϵ . This completes the proof.

Recall that $d(x, y)$ denotes the geodesic distance between any two points $x, y \in S_R$, and $G_r(x)$ is the geodesic cap $\{y \in S_R : d(x, y) \leq r\}$.

3.2. Proof of Theorem 2.2

Let s be at the north pole, let t be at the south pole, and consider any decentralized algorithm g_t .

(i) Here $0 \leq a < 1$. The proof uses Lemmas 3.2 and 3.3, below.

Lemma 3.2. For $u, v \in S_R$, let $P(u, G_r(v))$ denote the probability that u 's long-range contact is inside a geodesic cap $G_r(v)$ of radius r around v , where $r = \pi\epsilon L^{(1-a)/2}$. Then

$$P(u, G_r(v)) \leq 4L^{(a-1)/2}.$$

Proof. Equation (2.2), Lemma 3.1, and spherical symmetry imply that

$$\begin{aligned} P(u, G_r(v)) &\leq 2 \frac{\pi - 0}{2\pi} \int_{\pi\epsilon}^{2r+\pi\epsilon} \frac{1 - a}{\pi^{1-a}(R^{1-a} - \epsilon^{1-a})} \frac{1}{\rho^a} d\rho \\ &\leq \frac{2r(1 - a)}{\pi^{1-a}(R^{1-a} - \epsilon^{1-a})(\pi\epsilon)^a} \\ &= \frac{2r(1 - a)}{\pi\epsilon^a(R^{1-a} - \epsilon^{1-a})} \\ &\leq \frac{4(1 - a)}{\pi R} L^a r \\ &= 4(1 - a)L^{(a-1)/2} \\ &\leq 4L^{(a-1)/2}. \end{aligned}$$

Lemma 3.3. Let A be the event that a message starting at s is delivered by g_t to a point u that has its long-range contact in $G_r(t)$ in at most $T = \frac{1}{8}L^{(1-a)/2}$ steps, where r is as defined in Lemma 3.2. Then $P(A) \leq \frac{1}{2}$.

Proof. Let A_i be the event that the long-range contact of the message holder in the i th step of decentralized algorithm g_t is in $G_r(t)$. Note that $A = \bigcup_{i=1}^T A_i$, and $P(A_i) \leq 4L^{(a-1)/2}$ by Lemma 3.2. Therefore, subadditivity implies that

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

This proves the lemma.

Let B be the event that the message starting at s is ε -delivered to t in at most T steps. The event B cannot occur if A (as defined in Lemma 3.3) does not occur. Suppose that A does not occur, but the message is ε -delivered to t in at most T steps. Note that $T < L = \pi R/\pi\varepsilon$. Thus, a message that is ε -delivered to t starting at s 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 $G_r(t)$ since $r/\pi\varepsilon = 8T > T$. This contradicts our hypothesis that A does not occur. Note that $E(\tau_{s,t}(a, g_t)) = E(\tau_{s,t}(a, g_t) \mid A)P(A) + E(\tau_{s,t}(a, g_t) \mid A^c)P(A^c)$. Hence, $E(\tau_{s,t}(a, g_t)) \geq E(\tau_{s,t}(a, g_t) \mid A^c)P(A^c) \geq \frac{1}{2}E(\tau_\varepsilon(s, t) \mid A^c)$, where the last inequality follows because $P(A) \leq \frac{1}{2}$ by Lemma 3.3. Note that $E(\tau_{s,t}(a, g_t) \mid A^c) > T$. As a result,

$$E(\tau_{s,t}(a, g_t)) > \frac{1}{2}T = \frac{1}{2} \frac{1}{8}L^{(1-a)/2} = \frac{1}{16}L^{(1-a)/2}.$$

This completes the proof of Theorem 2.2(i).

(ii) Here $a > 1$. The proof uses Lemmas 3.4 and 3.5, below.

Lemma 3.4. For $u, v \in S_R$, let $P(u, G_r^c(u))$ denote the probability that u 's long-range contact is outside a geodesic cap of radius r around u , where $r = \pi\varepsilon L^{1/a}$. Then $P(u, G_r^c(u)) \leq 2(1/L)^{(a-1)/a}$.

Proof. Equation (2.2), Lemma 3.1, and spherical symmetry imply that

$$\begin{aligned} P(u, G_r^c(u)) &\leq 2 \frac{\pi - 0}{2\pi} \int_r^{\pi R} \frac{1 - a}{\pi^{1-a}(R^{1-a} - \varepsilon^{1-a})} \frac{1}{\rho^a} d\rho \\ &= \frac{1}{\pi^{1-a}(R^{1-a} - \varepsilon^{1-a})} ((\pi R)^{1-a} - r^{1-a}) \\ &= \frac{(r^{1-a} - (\pi R)^{1-a})\varepsilon^{a-1} R^{a-1}}{(R^{a-1} - \varepsilon^{a-1})\pi^{1-a}} \\ &\leq \frac{r^{1-a} \varepsilon^{a-1} R^{a-1}}{(R^{a-1} - \varepsilon^{a-1})\pi^{1-a}} \\ &\leq 2 \frac{r^{1-a} \varepsilon^{a-1}}{\pi^{1-a}} \\ &= 2 \frac{1}{L}^{(a-1)/a}. \end{aligned}$$

Lemma 3.5. Let A be the event that a message starting at s is delivered by g_t to a point u that has its long-range contact in $G_r^c(u)$ in at most $T = \frac{1}{4}L^{(a-1)/a}$ steps, where r is as defined in Lemma 3.4. Then $P(A) \leq \frac{1}{2}$.

Proof. Let A_i be the event that the message holder in step i (say u) of decentralized algorithm g_t has a long-range contact in $G_r^c(u)$. Lemma 3.4 implies that $P(A_i) \leq 2(1/L)^{(a-1)/a}$. Since $A = \bigcup_{i=1}^T A_i$, we have

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

This proves the lemma.

Now let B be the event that the message is ε -delivered to t in at most T steps. The event B cannot occur if A (as defined in Lemma 3.5) does not occur. To see this, if A does not occur, the message covers at most distance $r = \pi \varepsilon L^{1/a}$ in each of the first $T = \frac{1}{4} L^{(a-1)/a}$ steps of the decentralized algorithm. Thus, the total distance covered is at most $\pi \varepsilon L^{1/a} \frac{1}{4} L^{(a-1)/a} = \pi R/4$. Thus, the message cannot be ε -delivered from s to t in at most T steps since $d(s, t) = \pi R$. By calculations similar to the proof of part (i) above we obtain

$$E(\tau_{s,t}(a, g_t)) \geq E(\tau_{s,t}(a, g_t) \mid A^c) P(A^c) \geq \frac{T}{2} \geq \frac{1}{8} L^{(a-1)/a}.$$

This completes the proof of Theorem 2.2(ii).

4. Structural properties: scale invariance and significant improvement

In this section we define two structural properties of long-range contacts, i.e. properties that do not depend on any algorithmic message-passing procedure, but rather are intrinsic to probability density (2.1). These are scale invariance and significant improvement. We show that the model with $a = 1$ is the only model within our family that possesses these properties. Thus, these two structural properties are equivalent to the existence of an efficient decentralized search algorithm.

Definition 4.1. Without loss of generality, let $\varepsilon \ll 1$ and $R = 2^M$ for some integer $M \gg 1$. For $j = 0, 1, 2, \dots, M - 1$ and $a \geq 0$, let $P_j^u(a)$ be the probability that the geodesic distance between u and its long-range contact $v_5(u)$ is between $\pi 2^j$ and $\pi 2^{j+1}$. We say that the probability density in (2.1) is *scale invariant* if and only if $P_j^u(a)$ does not depend on j .

The intuition behind scale invariance is similar to Kleinberg’s original grid lattice model and its extensions. The long-range contact of u is equally likely to belong to any one of the ‘rings’ of width 2^j at geodesic distance 2^j from u if the probability density in (2.1) is scale invariant.

Proposition 4.1. *The probability density in (2.1) is scale invariant if and only if $a = 1$.*

Proof. From Lemma 3.1 we have

$$P_j^u(a) = \frac{2\pi - 0}{2\pi} \int_{\pi 2^j}^{\pi 2^{j+1}} \frac{1}{C(a; R, \varepsilon) \rho^a} d\rho.$$

Then, using (2.2), we obtain

$$P_j^u(1) = \frac{1}{\log L} \int_{\pi 2^j}^{\pi 2^{j+1}} \frac{1}{\rho} d\rho = \frac{\log 2}{\log L}.$$

On the other hand, for $a \geq 0, a \neq 1$, we have

$$P_j^u(a) = \frac{1 - a}{\pi^{1-a}(R^{1-a} - \varepsilon^{1-a})} \int_{\pi 2^j}^{\pi 2^{j+1}} \frac{1}{\rho^a} d\rho = \frac{2^{j(1-a)}(2^{1-a} - 1)}{R^{1-a} - \varepsilon^{1-a}},$$

proving that the probability density in (2.1) is scale invariant if and only if $a = 1$.

Definition 4.2. Let u and v be two points in S_R such that the geodesic distance between u and v is $l \geq 2\pi\varepsilon$. Let $P_l(u, v; a)$ be the probability that this long-range contact is within geodesic distance $l/2$ from v . We say that the probability density in (2.1) exhibits *steady improvement* if and only if $P_l(u, v; a) \geq \alpha/\log L$, where α is a constant that does not depend on u, v , or distance l .

Proposition 4.2. *The probability density function in (2.1) exhibits steady improvement if and only if $a = 1$.*

Proof. Lemma 3.1 and Corollary A.1 in Appendix A imply that

$$P_l(u, v; 1) \geq \frac{\log(4/3)}{12 \log L},$$

i.e. the probability density in (2.1) exhibits steady improvements if $a = 1$. From Lemma 3.1 we obtain, for $a \neq 1$,

$$P_l(u, v; a) \leq \frac{1 - a}{\pi^{1-a}(R^{1-a} - \varepsilon^{1-a})} \int_{l/2}^l \frac{1}{\rho^a} d\rho = \frac{l^{1-a}}{\pi^{1-a}(R^{1-a} - \varepsilon^{1-a})} \left(1 - \frac{1}{2^{1-a}}\right),$$

which is bounded above by $(l/\pi R)^{1-a}$ when $0 \leq a < 1$. On the other hand, when $a > 1$, it is bounded above by $(2\pi R/l)^{a-1}(2/L^{a-1})$, implying that

$$P_{\pi R}(u, t; a) \leq 2\left(\frac{2}{L}\right)^{a-1} \text{ for } a > 1.$$

In summary, when $a \neq 1$, there cannot exist a constant α such that $P_l(u, v; a) \geq \alpha/\log L$. Therefore, the density does not exhibit steady improvements when $a \neq 1$.

Appendix A. Auxiliary geometric results

Recall that $d(x, y)$ denotes the geodesic distance between any two points $x, y \in S_R$ and that $0 \leq d(x, y) \leq \pi R$. The great circle that passes through points $x, y \in S_R$ will be denoted by $O(x, y)$, and the (shorter) great circle arc connecting x and y will be denoted by $A(x, y)$.

Lemma A.1. *Let $u \in S_R$, and let $v_1(u), v_2(u), v_3(u)$, and $v_4(u)$ be its local contacts evenly spaced (in that order) on $\partial G_{\pi\varepsilon}(u)$, as described in Section 2. Let z be any point in $S_R \setminus G_{\pi\varepsilon}(u)$. Then*

$$\min\{d(v_1(u), z), d(v_2(u), z), d(v_3(u), z), d(v_4(u), z)\} \leq d(u, z).$$

Proof. The claim is obvious when $d(u, z) = \pi R$. For the other cases, it suffices to show that there exists a local contact of u such that the geodesic distance between this local contact and z is less than or equal to the distance between u and z . Let a local contact of u (say $v_1(u)$) be such that the angle between $O(u, z)$ and $O(u, v_1(u))$ is strictly positive (the claim is obvious when this angle is 0) and at most $\pi/4$. Such a local contact clearly exists because the angular

separation between great circle pairs $O(u, v_1(u)) - O(u, v_2(u))$, $O(u, v_2(u)) - O(u, v_3(u))$, and $O(u, v_3(u)) - O(u, v_4(u))$ is $\pi/2$ each. We will show that $d(v_1(u), z) \leq d(u, z)$. Toward this end, let m be any point on the geodesic arc $A(u, z)$. Then $d(u, z) = d(u, m) + d(m, z)$. Now if m is such that $d(u, m) \geq d(v_1(u), m)$ then we have $d(u, z) \geq d(v_1(u), m) + d(m, z) \geq d(v_1(u), z)$ by the triangle inequality, completing the proof. Thus, it suffices to show the existence of an m on $A(u, z)$ such that $d(u, m) \geq d(v_1(u), m)$. We show that choosing m as the point of an orthogonal intersection between $O(u, z)$ and a great circle through $v_1(u)$ serves this purpose. In particular, we show that, for m chosen this way, $d(u, m) \leq d(u, v_1(u)) = \pi \varepsilon$ (claim 1), proving that m is on $A(u, z)$ since $d(u, z) \geq \pi \varepsilon = d(u, v_1(u))$, and that $d(u, m) \geq d(v_1(u), m)$ (claim 2). The proof of claim 1 is as follows. Consider the geodesic triangle $\Delta uv_1(u)m$. With some abuse of notation, we will use the vertices u , $v_1(u)$, and m of this triangle to denote the measure of the corresponding angles. Using the geodesic law of cosines for this triangle implies that

$$\begin{aligned} \cos\left(\frac{\pi \varepsilon}{R}\right) &= \sin\left(\frac{d(v_1(u), m)}{R}\right) \cos\left(\frac{\pi}{2}\right) \sin\left(\frac{d(u_1(u), m)}{R}\right) \\ &\quad + \cos\left(\frac{d(v_1(u), m)}{R}\right) \cos\left(\frac{d(u, m)}{R}\right) \\ &= \cos\left(\frac{d(v_1(u), m)}{R}\right) \cos\left(\frac{d(u, m)}{R}\right). \end{aligned}$$

Since $\varepsilon/R \ll \frac{1}{2}$, $\cos(d(v_1(u), m)/R) \neq 0$; that is, $d(v_1(u), m)/R \neq \pi/2$, implying that

$$\cos\left(\frac{d(u, m)}{R}\right) = \frac{\cos(\pi \varepsilon/R)}{\cos(d(v_1(u), m)/R)}. \tag{A.1}$$

Now using the geodesic law of cosines one more time yields

$$\frac{\sin(d(u, v_1(u))/R)}{\sin(d(u, m)/R)} = \frac{\cos(u)}{\cos(d(v_1(u), m)/R)}.$$

But the geodesic law of sines implies that

$$\frac{\sin(d(u, v_1(u))/R)}{\sin(d(u, m)/R)} = \frac{\sin(m)}{\sin(v_1(u))} = \frac{1}{\sin(v_1(u))} > 0,$$

where the last strict inequality follows from the fact that $0 < d(u, m) < \pi$ implies that $0 < v_1(u) < \pi$. As a result,

$$\frac{\cos(u)}{\cos(d(v_1(u), m)/R)} > 0,$$

implying that $d(v_1(u), m)/R < \pi/2$, since $\cos(u) > 0$ as $0 \leq u \leq \pi/4$. Then (A.1) implies that $\cos(d(u, m)/R) \geq \cos(\pi \varepsilon/R)$. Therefore, $d(u, m) \leq \pi \varepsilon$, proving claim 1.

The proof of claim 2 is as follows. The only nontrivial situation is when $d(v_1(u), m) > 0$. The geodesic law of sines implies that

$$\frac{\sin(d(u, m)/R)}{\sin(d(v_1(u), m)/R)} = \frac{\sin(v_1(u))}{\sin(u)}. \tag{A.2}$$

Observe that $u + v_1(u) + m \geq \pi$, since $\Delta uv_1(u)m$ is geodesic. Recall that $m = \pi/2$ and $u \leq \pi/4$. Thus, $v_1(u) \geq \pi/4$. Since $d(u, m)/R$ and $d(v_1(u), m)/R$ are both acute, (A.2) then

implies that $d(u, m) \geq d(v_1(u), m)$, proving claim 2, as long as $v_1(u) \leq \pi/2$. Thus, it remains to show that $v_1(u)$ is acute. The geodesic law of cosines yields

$$\cos(v_1(u)) = \frac{\cos(d(u, m)/R) - \cos(d(v_1(u), m)/R) \cos(\pi \varepsilon/R)}{\sin(d(v_1(u), m)/R) \sin(\pi \varepsilon/R)}.$$

Substituting for $\cos(d(u, m)/R)$ using (A.1) along with some algebraic simplification leads to

$$\cos(v_1(u)) = \frac{\cos(d(v_1(u), m)/R) \cos(\pi \varepsilon/R)}{\sin(d(v_1(u), m)/R) \sin(\pi \varepsilon/R)} > 0,$$

where the last strict inequality follows since $0 < d(v_1(u), m)/R < \pi/2$. This shows that $v_1(u)$ is acute, completing the proof of the lemma.

Lemma A.2. *The lemma has two parts.*

1. Define $h_1(\phi, \theta) \equiv \cos(\theta) \sin^2(\phi) + \cos^2(\phi) - \cos(\phi/2)$. Then $h_1(\phi, \theta) \geq 0$ for all $0 \leq \phi \leq \pi$ if $0 \leq \theta \leq \arccos(\frac{7}{8})$. In particular, since $\pi/12 \leq \arccos(\frac{7}{8})$,

$$\cos(\theta) \sin^2(\phi) + \cos^2(\phi) \geq \cos(\phi/2) \tag{A.3}$$

for all $0 \leq \phi \leq \pi$ if $0 \leq \theta \leq \pi/12$.

2. Define $h_2(\phi, \theta) \equiv \cos(\theta) \sin(\phi) \sin(3\phi/4) + \cos(\phi) \cos(3\phi/4) - \cos(\phi/2)$. Then $h_2(\phi, \theta) \geq 0$ for all $0 \leq \phi \leq \pi$ if $0 \leq \theta \leq \arccos(\frac{7}{8})$. In particular, since $\pi/12 \leq \arccos(\frac{7}{8})$,

$$\cos(\theta) \sin(\phi) \sin(3\phi/4) + \cos(\phi) \cos(3\phi/4) \geq \cos(\phi/2) \tag{A.4}$$

for all $0 \leq \phi \leq \pi$ if $0 \leq \theta \leq \pi/12$.

Proof. We begin with the first claim. Note that

$$\begin{aligned} h_1(\phi, \theta) &= \cos(\theta) \sin^2(\phi) + 1 - \sin^2(\phi) - \cos(\phi/2) \\ &= \sin^2(\phi)(\cos(\theta) - 1) + 1 - \cos(\phi/2) \\ &= \sin^2(\phi)(\cos(\theta) - 1) + 2 \sin^2(\phi/4) \\ &= [2 \sin(\phi/2) \cos(\phi/2)]^2(\cos(\theta) - 1) + 2 \sin^2(\phi/4) \\ &= [4 \sin(\phi/4) \cos(\phi/4) \cos(\phi/2)]^2(\cos(\theta) - 1) + 2 \sin^2(\phi/4) \\ &= 2 \sin^2(\phi/4) 8 \cos^2(\phi/4) \cos^2(\phi/2)(\cos(\theta) - 1) + 2 \sin^2(\phi/4) \\ &= 2 \sin^2(\phi/4) [8 \cos^2(\phi/4) \cos^2(\phi/2)(\cos(\theta) - 1) + 1]. \end{aligned}$$

Thus, $h_1(\phi, \theta) \geq 0$ if and only if $\bar{h}_1(\phi, \theta) \equiv 8 \cos^2(\phi/4) \cos^2(\phi/2)(\cos(\theta) - 1) + 1 \geq 0$. We will show that if $0 \leq \theta \leq \arccos(\frac{7}{8})$, i.e. if $1 \geq \cos(\theta) \geq \frac{7}{8}$, then $\bar{h}_1(\phi, \theta) \geq 0$. For in this case, $\cos(\theta) - 1 \geq -\frac{1}{8}$, implying that $8 \cos^2(\phi/4) \cos^2(\phi/2)(\cos(\theta) - 1) \geq -\cos^2(\phi/4) \cos^2(\phi/2)$, which in turn is bounded below by -1 , proving the first claim.

For the second claim, note that

$$\cos(\phi/4) = \cos(\phi - 3\phi/4) = \cos(\phi) \cos(3\phi/4) + \sin(\phi) \sin(3\phi/4).$$

As a result, $\cos(\phi) \cos(3\phi/4) = \cos(\phi/4) - \sin(\phi) \sin(3\phi/4)$. Therefore,

$$\begin{aligned} h_2(\phi, \theta) &= \sin(\phi) \sin(3\phi/4)(\cos(\theta) - 1) + \cos(\theta/4) - \cos(\theta/2) \\ &= \sin(\phi) \sin(3\phi/4)(\cos(\theta) - 1) + 2 \sin(3\phi/8) \sin(\phi/8) \\ &= 2 \sin(\phi) \sin(3\phi/8) \cos(3\phi/8)(\cos(\theta) - 1) + 2 \sin(3\phi/8) \sin(\phi/8) \\ &= 2 \sin(3\phi/8)[\sin(\phi) \cos(3\phi/8)(\cos(\theta) - 1) + \sin(\phi/8)]. \end{aligned}$$

Since $\sin(3\phi/8)$ is nonnegative for $0 \leq \phi \leq \pi$, $h_2(\phi, \theta)$ is nonnegative if and only if $\sin(\phi) \cos(3\phi/8)(\cos(\theta) - 1) + \sin(\phi/8)$ is nonnegative. That is, if and only if $\sin(\phi/8) \geq \sin(\phi) \cos(3\phi/8)(1 - \cos(\theta))$. We will show that this inequality holds when $0 \leq \theta \leq \arccos(\frac{7}{8})$. For in this case, $(1 - \cos(\theta)) \geq \frac{1}{8}$. Moreover, $0 \leq \phi \leq \pi$ implies that $\sin(\phi) \geq 0$ and $\cos(3\phi/8) \geq 0$. Therefore, $\sin(\phi) \cos(3\phi/8)/8 \geq \sin(\phi) \cos(3\phi/8)(1 - \cos(\theta))$. In addition, $\sin(\phi)/8 \geq \sin(\phi) \cos(3\phi/8)/8$. Thus, it suffices to prove that $8 \sin(\phi/8) \geq \sin(\phi)$. To see that this inequality holds, observe that, for $0 \leq \phi \leq \pi$,

$$\begin{aligned} 8 \sin(\phi/8) &\geq 8 \sin(\phi/8) \cos(\phi/8) \cos(\phi/4) \cos(\phi/2) \\ &= 4 \sin(\phi/4) \cos(\phi/4) \cos(\phi/2) \\ &= 2 \sin(\phi/2) \cos(\phi/2) \\ &= \sin(\phi), \end{aligned}$$

completing the proof.

Corollary A.1. Consider a geodesic triangle $\Delta uv t$ on a sphere of radius R such that the geodesic angle $\angle vut = \theta \leq \pi/12$.

1. If $d(u, v) = d(u, t) = s$ then point v is inside the geodesic cap of radius $s/2$ centered at t .
2. If $d(u, t) = s$ and $d(u, v) = 3s/4$ then point v is inside the geodesic cap of radius $s/2$ centered at t .

Proof. Let $d(v, t) = x$. For the first case where $d(u, v) = d(u, t) = s$, the spherical law of cosines implies that

$$\cos(x/R) = \cos(\theta) \sin^2(s/R) + \cos^2(s/R).$$

Equation (A.3) then yields

$$\cos(x/R) \geq \cos\left(\frac{s/2}{R}\right).$$

In other words, $x \leq s/2$, proving the first claim. For the second case, where $d(u, t) = s$ and $d(u, v) = 3s/4$, the spherical law of cosines implies that

$$\cos(x/R) = \cos(\theta) \sin(s/R) \sin(3s/4R) + \cos(s/R) \cos(3s/4R).$$

Equation (A.4) yields

$$\cos(x/R) \geq \cos\left(\frac{s/2}{R}\right).$$

In other words, $x \leq s/2$, proving the claim.

Acknowledgements

The author thanks the University of Washington for financial support. Several suggestions by two anonymous referees and an editor improved the presentation as compared to an earlier version of this manuscript.

References

- [1] ALBERT, R. AND BARABASI, A. L. (2002). Statistical mechanics of complex networks. *Rev. Modern Phys.* **74**, 47–97.
- [2] BARABASI, A. L. AND ALBERT, R. (1999). Emergence of scaling in random networks. *Science* **286**, 509–512.
- [3] FRANCESCETTI, M. AND MEESTER, R. (2006). Navigation in small-world networks: a scale free continuum model. *J. Appl. Prob.* **43**, 1173–1180.
- [4] GANESH, A. J. AND DRAIEF, M. (2006). Efficient routing in Poisson small world networks. *J. Appl. Prob.* **43**, 678–686.
- [5] GHATE, A. (2006). Markov chains, game theory and infinite programming: three paradigms for optimization of complex systems. Doctoral Thesis, The University of Michigan.
- [6] GHATE, A. AND SMITH, R. L. (2008). A hit-and-run approach to generating scale invariant small world networks. To appear in *Networks*.
- [7] KLEINBERG, J. M. (2000). Navigation in a small world. *Nature* **406**, 845.
- [8] KLEINBERG, J. M. (2000). The small world phenomenon: an algorithmic perspective. In *Proc. 32nd Annual ACM Symp. Theory Comput.* ACM, New York, pp. 163–170.
- [9] LIBEN-NOWELL, D. *et al.* (2005). Geographic routing in social networks. *Proc. Nat. Acad. Sci.* **102**, 11623–11628.
- [10] MILGRAM, S. (1967). The small world problem. *Psychology Today* **2**, 60–67.
- [11] NEWMAN, M. E. J. (2000). Models of the small world. *J. Statist. Phys.* **101**, 819–941.
- [12] NEWMAN, M. E. J. (2003). The structure and function of complex networks. *SIAM Rev.* **45**, 167–256.
- [13] TRAVERS, J. AND MILGRAM, S. (1969). An experimental study of the small world problem. *Sociometry* **32**, 425–443.
- [14] WATTS, D. AND STROGATZ, S. (1998). Collective dynamics of small world networks. *Nature* **393**, 440–442.
- [15] WATTS, D., DODDS, P. S. AND NEWMAN, M. E. J. (2002). Identity and search in social networks. *Science* **296**, 1302–1305.