

The Effect of Power-Law Degrees on the Navigability of Small Worlds*

[Extended Abstract]

Pierre Fraigniaud
CNRS and Univ. Paris Diderot
Paris, France
pierre.fraigniaud@liafa.jussieu.fr

George Giakkoupis
Univ. Paris Diderot
Paris, France
ggiak@liafa.jussieu.fr

ABSTRACT

We analyze decentralized routing in small-world networks that combine a wide variation in node degrees with a notion of spatial embedding. Specifically, we consider a variation of Kleinberg’s augmented-lattice model (STOC 2000), where the number of long-range contacts for each node is drawn from a power-law distribution. This model is motivated by the experimental observation that many “real-world” networks have power-law degrees. In such networks, the exponent α of the power law is typically between 2 and 3. We prove that, in our model, for this range of values, $2 < \alpha < 3$, the expected number of steps of greedy routing from any source to any target is $O(\log^{\alpha-1} n)$ steps. This bound is tight in a strong sense. Indeed, we prove that the expected number of steps of greedy routing for a uniformly-random pair of source–target nodes is $\Omega(\log^{\alpha-1} n)$ steps. We also show that for $\alpha < 2$ or $\alpha \geq 3$, greedy routing performs in $\Theta(\log^2 n)$ expected steps, and for $\alpha = 2$, $\Theta(\log^{1+\varepsilon} n)$ expected steps are required, where $1/3 \leq \varepsilon \leq 1/2$. To the best of our knowledge, these results are the first to formally quantify the effect of the power-law degree distribution on the navigability of small worlds. Moreover, they show that this effect is significant. In particular, as α approaches 2 from above, the expected number of steps of greedy routing in the augmented lattice with *power-law degrees* approaches the square-root of the expected number of steps of greedy routing in the augmented lattice with *fixed degrees*, although both networks have the same *average degree*.

Categories and Subject Descriptors

F.2.2 [Nonnumerical Algorithms and Problems]: Routing and layout; E.1 [Data Structures]: Graphs and networks; G.2.2 [Graph Theory]: Graph algorithms, Network problems; C.2.2 [Network Protocols]: Routing protocols

*Both authors are supported by the ANR project AL-ADDIN, and by the INRIA project GANG.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PODC’09, August 10–12, 2009, Calgary, Alberta, Canada.
Copyright 2009 ACM 978-1-60558-396-9/09/08 ...\$5.00.

General Terms

Algorithms, Performance, Theory

Keywords

small-world graphs, power-law degrees, greedy routing

1. INTRODUCTION

1.1 Navigability of small worlds

It has been observed that many “real-world” networks, such as social, information, technological, and biological networks, exhibit the *small-world* property; i.e., they are locally clustered, and (yet) short paths exist between almost all pairs of nodes (see [2, 9, 19] and the references therein). It is also well-established that many small-world networks (e.g., the network of acquaintances between individuals) are easy to *navigate*, provided that the nodes are able to estimate the distances to other nodes with respect to some underlying metric (e.g., geography, professions, etc.) [8, 18]. *Navigability* refers to the ability of nodes to route messages efficiently in a decentralized manner, using local information only. The most prominent example of such a routing scheme is *greedy routing*: a node handling a message destined to some target node forwards the message to its neighbor that is closest to the target, according to the underlying metric. The first formal analysis of greedy routing in a plausible model of small worlds was presented in [13]. The model studied there was the *augmented lattice*: Consider the n -node d -dimensional lattice that wraps around, where $d \geq 1$. A node has links to its $2d$ lattice-neighbors, and also to $k \geq 1$ other nodes, its *long-range contacts*. Each of the long-range contacts of a node u is chosen using an independent random trial following the d -harmonic distribution: the probability that node v is chosen in a given trial is

$$p_{u,v} \propto 1/(\text{dist}(u,v))^d, \quad (1.1)$$

where $\text{dist}(u,v)$ is the lattice distance between u and v . In [13] it was shown that, in this model, greedy routing requires $O(\frac{1}{k} \log^2 n)$ expected number of steps, for any source–target pair. (This complexity was later shown to be tight [17].) It was also shown that *any* decentralized routing algorithm performs poorly if the d -dimensional lattice is augmented using the h -harmonic distribution, for any $h \neq d$. Specifically, $\Omega(n^\gamma)$ expected steps are required, for some $\gamma > 0$ that depends on h and d .¹

¹It was recently shown [7] that for $d = 1$, the augmentation

Despite its simplicity, the augmented-lattice model seems to capture successfully the small-world and navigability properties of real-world networks. Note that in the d -dimensional lattice the d -harmonic distribution is equivalent to the “natural” distribution $p_{u,v} \propto 1/|B_u(\text{dist}(u,v))|$, where $B_u(r)$ is the ball centered at u of radius r ; this latter distribution was used in [10, 23] to extend the results of [13] to graphs of bounded ball growth, and to graphs of bounded doubling dimension. Also, the d -harmonic distribution is equivalent in the lattice to the rank-based distribution $p_{u,v} \propto 1/r_u(v)$, where $r_u(v)$ is the rank of v when nodes are sorted in increasing distance from node u ; this latter distribution was used in [15] to extend the results of [13] to non-uniform population densities. In fact, it was experimentally demonstrated that two-thirds of friendships are geographically distributed this way: the probability of befriending a particular person is inversely proportional to the number of people closer to you [16]. Finally, it was recently shown that the d -harmonic distribution of the long-range links might as well be an inherent byproduct of node mobility [4]. See also [6, 20] for other dynamics yielding the d -harmonic distribution in the lattice. Therefore, there is now a consensus that the augmented-lattice model is an appropriate framework for analyzing small-world navigability.

1.2 Power-law degree distribution

The augmented-lattice model, however, fails to capture another commonly observed property of real-world networks, the *heavy-tailed degree distribution*. Such a distribution is well approximated by a *power law*

$$\Pr[\text{deg}(u) = k] \propto 1/k^\alpha, \quad (1.2)$$

where α is a real, typically between 2 and 3 [2, 9, 19]. Nevertheless, it is straightforward to reconcile the augmented-lattice model with a power-law distribution for the node degrees, simply by drawing the number of long-range links added to each node independently at random from a power-law distribution [14]. It is reasonable to expect that this modification would reduce the lengths of shortest paths between nodes and the network diameter, since the (few) high-degree nodes should provide short-cuts between most nodes. This is typically the case in networks with power-law degree sequences [3, 5]. However, it is unclear how decentralized routing could benefit from the existence of these high-degree nodes [14].

Utilizing the heavy-tailed degree distribution in the design of decentralized routing algorithms was suggested in [1, 11, 12, 21]. In all these works, the routing algorithms only have access to information about the degrees of neighboring nodes, not to any embedding of the graph. Although some performance improvements are observed compared to routing algorithms oblivious to the node degrees, the expected number of steps remains polynomial in the network size. Also, [22] proposed a heuristic decentralized algorithm for routing in a variance of the augmented lattice where nodes have widely varying degrees. This heuristic assumes that nodes have access both to the locations of their neighbors, and to their degrees. Simulations showed that this algorithm performs better than decentralized algorithms using

using the 1-harmonic distribution is essentially optimal in the sense that for *any* augmentation distribution with k *expected* long-range contacts per node, greedy routing requires $\Omega(\frac{1}{k} \log^2 n)$ expected steps.

only one of these two sources of information. However, no formal analysis was provided.

1.3 Our framework

We consider the following variance of the augmented-lattice model. As in the original model, the long-range links are drawn independently at random according the harmonic distribution with exponent equal to the dimensionality of the lattice (cf. Eq. 1.1). Unlike the original model, however, the number of long-range contacts each node has is not fixed, but it is drawn independently at random from the power-law distribution with exponent $\alpha \geq 0$ (cf. Eq. 1.2). This distribution is scaled so that its expectation is constant and each node has at least one long-range contact.² We then remove the orientation of each of the long-range links to get a non-directed network. We study the performance of greedy routing in this network.

1.4 Our results

In this section, we ignore $O(\log \log n)$ multiplicative factors in the statement of the asymptotic bounds. The precise bounds are described in Section 2.3.

We prove that for $2 < \alpha < 3$, which is the case for most real-world networks, the expected number of steps of greedy routing *from any source to any target* is $O(\log^{\alpha-1} n)$ steps. Thus, for this range of values for α , the effect of the power-law degree distribution is significant. In particular, when α approaches 2, the expected number of steps of greedy routing in the augmented lattice with *power-law degrees* approaches the square-root of the expected number of steps of greedy routing in the augmented lattice with *fixed degrees*, although both networks have the same *average degree*. For both $\alpha < 2$ and $\alpha \geq 3$, we show that the expected number of steps of greedy routing from any source to any target is $O(\log^2 n)$ steps, which is the same order of magnitude as the performance of greedy routing in the augmented lattice with fixed degrees. For the critical value $\alpha = 2$, we prove that the expected number of steps of greedy routing from any source to any target is $O(\log^{3/2} n)$ steps.

All these upper bounds are tight (but, perhaps, for $\alpha = 2$). For $\alpha > 2$, the upper bounds are even tight in a strong sense. Indeed, we prove that the expected number of steps of greedy routing for a *uniformly-random* pair of source–target nodes is $\Omega(\log^{\alpha-1} n)$ steps if $2 < \alpha < 3$, and $\Omega(\log^2 n)$ steps if $\alpha \geq 3$. For $\alpha < 2$, we prove that there exists a source–target pair for which greedy routing requires $\Omega(\log^2 n)$ expected steps. For $\alpha = 2$, we show that the expected number of steps for a uniformly-random source–target pair is $\Omega(\log^{4/3} n)$.

We formally prove the above results for the case of the 1-dimensional lattice, i.e., the ring. Nevertheless, none of the arguments we use is specifically tied to the ring, and the *exact* same results can be easily shown for d -dimensional lattices, for constant values of d . Note that unlike the results in [13], where the critical value of the exponent depends on the dimensionality d of the lattice, our results do not depend on d .

To the best of our knowledge, these results are the first to formally quantify the effect of the power-law degree distribution on the navigability of small worlds.

²For $\alpha > 2$, even without the scaling, the expectation is constant and, with constant probability, each node has at least one long-range contact. So, the scaling makes a difference only for $\alpha \leq 2$.

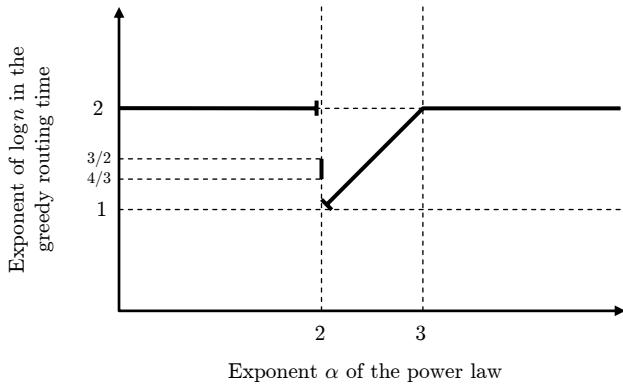


Figure 1: Summary of the results.

The following picture emerges from our analysis. For $\alpha \geq 3$, almost all nodes are of small degree, and the nodes of higher degree are too few to contribute significantly. Hence greedy routing performs essentially the same as when the degrees are fixed.

For $2 < \alpha < 3$, there are still very few nodes of high degree. However, nodes of degree roughly $\log n$ are relatively abundant, and there are more and more such nodes as α approaches 2. It is the contribution of these nodes that reduces the routing time from $\log^2 n$ to $\log^{\alpha-1} n$.

The case $\alpha = 2$ is special. All “degree scales” are present, and each is equally likely to contribute. On the one hand, this results in greater routing speed than in the case $2 < \alpha < 3$ when the current node is far from the target, since there are many high-degree nodes between the current node and the target in the lattice. On the other hand, the balance in the degree scales means that as we get closer to the target the number of high-degree nodes available decreases faster than in the case $2 < \alpha < 3$; and when we get at distance sub-polynomial from the target (essentially at distance less than $e^{\sqrt{\ln n}}$), greedy routing performs the same as when the degrees are fixed.

Finally, for $\alpha < 2$, there are many nodes of high degree, and the role of the cut-off point k_{\max} of the power law becomes critical. We assumed that $k_{\max} \sim n^\gamma$, for some $0 < \gamma \leq 1$. In this setting, only the contribution of nodes with degree close to k_{\max} is significant. However, when the current node is at distance less than k_{\max} from the target, it is very likely that greedy routing will not find a node of such degree, and from that distance it starts performing the same as when the degrees are fixed. Note that for $\alpha < 2$, nodes that are further away from the target may, in expectation, require fewer steps to reach the target than nodes closer to the target, which is not the case when $\alpha > 2$.

2. MODEL AND MAIN RESULTS

2.1 Network model

We will use the notation $[i..j] = \{k \in \mathbb{Z} : i \leq k \leq j\}$ and $[i..j) = [i..j - 1]$, for $i, j \in \mathbb{Z}$. (If $i > j$ then $[i..j] = \emptyset$.) Also, whenever we treat a real number x as an integer we will mean $\lfloor x \rfloor$.

In our analysis we will focus on the 1-dimensional lattice case. Let \mathbf{G}_n be the class of all directed graphs with set of nodes $[0..n)$ that contain as a subgraph the n -node *ring*,

i.e., the graph with set of nodes $[0..n)$, and set of edges $\{(u, u \pm 1 \bmod n) : u \in [0..n)\}$.³ Let G be a graph in \mathbf{G}_n , and E be the set of edges of G . The *out-neighbors* (*in-neighbors*) of a node u of G are all nodes v such that $(u, v) \in E$ ($(v, u) \in E$). More specifically, the nodes $u \pm 1 \bmod n$ are called the *ring-neighbors* of u , and the remaining out-neighbors (*in-neighbors*) of u are its *out-contacts* (*in-contacts*). For any two subsets of nodes A and B , we will write $A \rightarrow B$ to denote that a node in B is an out-contact of a node in A (or, equivalently, a node in A is an in-contact of a node in B). When $|A| = 1$, say $A = \{a\}$, we will often write $a \rightarrow B$, instead; and the same convention is used when $|B| = 1$. The *ring-distance* between nodes u and v , denoted $\delta(u, v)$, is the minimum number of ring edges between them in either clockwise or counter-clockwise direction, i.e.,

$$\delta(u, v) = \min\{u - v \bmod n, v - u \bmod n\}.$$

So, if v is a ring-neighbor of u then $\delta(u, v) = 1$, and if it is an in-contact or out-contact of u then $\delta(u, v) \geq 2$. We will write $\|u\|$ to denote $\delta(u, 0)$.

Random-graph models: We will study two random-graph models. Each of them is parameterized by the size n of the graph, and the exponent $\alpha \geq 0$ of a power-law distribution. In the first model, denoted $\mathcal{G}(n, \alpha)$, a random element of \mathbf{G}_n is generated by choosing the out-contacts of the nodes as follows. For each node u , we draw an integer D_u from $[1..k_{\max}]$ independently at random, such that

$$(1) \Pr[D_u = k] \propto 1/k^\alpha, \text{ for } k \neq 1, \text{ and } (2) \mathbb{E}[D_u] = 2.$$

We assume that the “cut-off point” k_{\max} of the power-law distribution is $\Theta(n^\gamma)$, for some constant $0 < \gamma \leq 1$. Then, we perform D_u independent identical random trials, such that in each trial a node $v \neq u$ is chosen with probability

$$\propto 1/\delta(u, v).$$

The out-contacts of u are all the *distinct* nodes chosen by these D_u trials that are not ring-neighbors of u ; formally, if v_i is the node chosen in the i -th trial then the out-contacts of u are the elements of the set $\bigcup_{1 \leq i \leq D_u} \{v_i\} \setminus \{u \pm 1 \bmod n\}$. The second random-graph model we consider, denoted $\mathcal{U}(n, \alpha)$, is the model in which a random graph is obtained by first generating a random graph in $\mathcal{G}(n, \alpha)$, and then taking its underlying *undirected graph*; in fact, for each (directed) edge of $\mathcal{G}(n, \alpha)$ we also add its opposite-directed edge, if it does not already exist. Formally, if E is the set of edges in $\mathcal{G}(n, \alpha)$ then the set of edges in $\mathcal{U}(n, \alpha)$ is $\{(u, v) : (u, v) \in E \text{ or } (v, u) \in E\}$.

Discussion: Recall that the long-range contacts of a node in $\mathcal{G}(n, \alpha)$ are selected using independent trials *with replacement*. This assumption simplifies the analysis, but it has the side effect that the out-degree D_u^+ of node u in $\mathcal{G}(n, \alpha)$ can be smaller than D_u , and also the distribution of D_u^+ is not *exactly* a power law. Nevertheless, the use of trials with replacement gives essentially the same results as the trials *without replacement*. This is because the discrepancy between D_u^+ and D_u is significant only for large values of D_u (e.g., of order $\Omega(\sqrt{n})$). And our analysis shows that the

³For a graph in \mathbf{G}_n , the underlying ring will be used to compute distances between nodes. Also, when we refer to nodes we will mean their integer labels.

effect of such high-degree nodes is negligible when $\alpha > 2$; while for $\alpha < 2$ our proof actually holds even if trials are without replacement. For the in-degree D_u^- of u in $\mathcal{G}(n, \alpha)$, it is easy to see that its distribution is close to a Poisson with constant expectation; so, the distribution of the (total) degree $D_u^+ + D_u^-$ of u in $\mathcal{U}(n, \alpha)$ is essentially the same as that of D_u^+ for all but very small values.

Next, recall that $D_u \leq k_{\max} = \Theta(n^\gamma)$ and $D_u \neq 0$. For $\alpha > 2$, the exact same asymptotic results hold with or without these constraints. The rest of the discussion is for the case $\alpha \leq 2$. Note that it must be $k_{\max} < \infty$, otherwise the expectation of D_u is ∞ . Also, a value for k_{\max} that is polynomial in n is consistent with real-world networks. It can be shown that if k_{\max} is poly-log in n then greedy routing performs in logarithmic time. On the lower side, a minimum value is imposed on D_u because otherwise D_u would be 0 with overwhelming probability. It can be shown that, in that case, greedy routing would require polynomial time if $\alpha < 2$, and poly-log time if $\alpha = 2$.

2.2 Greedy routing

We consider the following routing algorithm for graphs in \mathcal{G}_n . When a node u receives a message for a target node $t \neq u$, u forwards the message to an out-neighbor that is closest to t , with respect to the ring-distance. We call this routing algorithm GREEDY. We are interested in the performance of GREEDY in $\mathcal{G}(n, \alpha)$ and $\mathcal{U}(n, \alpha)$. Specifically, we study two performance measures: the *expected delivery time* of GREEDY, and the *GREEDY diameter*. Let $l_{u,v}$ be the *expected length* of the GREEDY routing path from u to v in the random graph. The *expected delivery time* is the *average* of $l_{u,v}$, taken over all possible source–target pairs, i.e.,

$$\text{Expected delivery time of GREEDY} = \frac{1}{n^2} \sum_{u,v} l_{u,v}.$$

The *GREEDY diameter* is the corresponding *maximum*, i.e.,

$$\text{GREEDY diameter} = \max_{u,v} l_{u,v}.$$

Note that the GREEDY diameter is always greater or equal to the corresponding expected delivery time. All the lower bounds we prove, except for the model $\mathcal{U}(n, \alpha)$ with $\alpha < 2$, are for the expected delivery time of GREEDY; whereas all the upper bounds are for the GREEDY diameter.

Throughout our analysis of GREEDY in $\mathcal{G}(n, \alpha)$ and $\mathcal{U}(n, \alpha)$, we will assume that the target node is node 0. We can make this assumption without loss of generality because of the symmetry of the random-graph models. Also, for the analysis in $\mathcal{U}(n, \alpha)$, instead of considering the graph $\mathcal{U}(n, \alpha)$ directly, we will consider $\mathcal{G}(n, \alpha)$, and for the purposes of routing we will ignore the direction of the links. So, whenever we refer to the in-/out-contacts, in-/out-links, etc., of a node, we will mean in $\mathcal{G}(n, \alpha)$; the same convention is used for the ‘ \rightarrow ’ notation.

In $\mathcal{G}(n, \alpha)$, the GREEDY path from a fixed source to a fixed target is a Markov chain; the next node in the path depends only on the last node visited. However, this is not the case in $\mathcal{U}(n, \alpha)$, where the next node depends on all the previously visited nodes, and also on their in- and out-links. Specifically, if $\langle Y_0, Y_1, \dots \rangle$ is the routing path from node Y_0 to 0 then for any node v with $\|v\| < \|Y_i\|$, $v \not\rightarrow \{Y_0, \dots, Y_{i-1}\}$; hence, the values of Y_0, \dots, Y_{i-1} affect the distribution of the out-contacts of v . More importantly, the distribution of

the out-contacts of Y_i is affected by whether $Y_{i-1} \rightarrow Y_i$ or not; e.g., if $Y_{i-1} \not\rightarrow Y_i$ and Y_i is not a ring-neighbor of Y_{i-1} then $Y_i \rightarrow Y_{i-1}$, which, for some values of α , changes the a-priori distribution of the out-degree of Y_i significantly.

2.3 Statement of the results

In all the results below, the asymptotic notation is as $n \rightarrow \infty$, and α is not a function of n .

THEOREM 2.1. *The expected delivery time of GREEDY in $\mathcal{G}(n, \alpha)$ is $\Omega(\ln^2 n)$.*

THEOREM 2.2. *The GREEDY diameter of $\mathcal{G}(n, \alpha)$ is $O(\ln^2 n)$.*

THEOREM 2.3. *The expected delivery time of GREEDY in $\mathcal{U}(n, \alpha)$ is*

$$\begin{cases} \Omega(\ln^{4/3} n), & \text{if } \alpha = 2; \\ \Omega(\ln^{\alpha-1} n), & \text{if } 2 < \alpha < 3; \\ \Omega(\ln^2 n / \ln \ln n), & \text{if } \alpha = 3; \\ \Omega(\ln^2 n), & \text{if } \alpha > 3. \end{cases}$$

Also, for $0 \leq \alpha < 2$, the GREEDY diameter is $\Omega(\ln^2 n)$.

THEOREM 2.4. *The GREEDY diameter of $\mathcal{U}(n, \alpha)$ is*

$$\begin{cases} O(\ln^2 n), & \text{if } 0 \leq \alpha < 2; \\ O(\ln^{3/2} n), & \text{if } \alpha = 2; \\ O(\ln^{\alpha-1} n \ln \ln n), & \text{if } 2 < \alpha < 3; \\ O(\ln^2 n), & \text{if } \alpha \geq 3. \end{cases}$$

3. DEFINITIONS AND BASIC FACTS

3.1 Distribution of out-contacts

Recall that the out-contacts of a node u in $\mathcal{G}(n, \alpha)$ are chosen using independent identical trials, such that, in each trial, node $v \neq u$ is picked with probability $\propto 1/\delta(u, v)$. So, the probability that v is picked in a given trial is

$$\frac{1}{\nu \delta(u, v)}, \quad \text{where } \nu = \sum_{v \neq u} \frac{1}{\delta(u, v)} = 2 \ln n + O(1).$$

Also, the number D_u of trials used is chosen independently at random from $[1, k_{\max}]$, where $k_{\max} = \Theta(n^\gamma)$, such that $\mathbb{P}\mathbb{r}[D_u = k] \propto 1/k^\alpha$, for $k \neq 1$, and $\mathbb{E}[D_u] = 2$. Let

$$q_k = \mathbb{P}\mathbb{r}[D_u = k] = \begin{cases} \frac{1}{\beta k^\alpha}, & \text{if } k \neq 1; \\ 1 - \sum_{j \neq 1} q_j, & \text{if } k = 1, \end{cases}$$

where β is the normalizing factor such that $\mathbb{E}[D_u] = \sum_k k q_k = 2$. It is easy to see that

$$\beta = \sum_k \frac{k-1}{k^\alpha} = \begin{cases} \Theta(1), & \text{if } \alpha > 2; \\ \Theta(\ln n), & \text{if } \alpha = 2; \\ \Theta(k_{\max}^{2-\alpha}), & \text{if } 0 \leq \alpha < 2. \end{cases}$$

Also, the probability that $D_u = 1$ is $q_1 = \Theta(1)$, if $\alpha > 2$, and $q_1 = 1 - o(1)$, if $0 \leq \alpha \leq 2$.

3.2 Simple facts about $\mathcal{G}(n, \alpha)$

We now state without proof some simple facts that we use repeatedly in the analysis. In all these facts, the underlying graph is $\mathcal{G}(n, \alpha)$ and $u, v \in [0..n)$.

FACT 3.1. If $U \subseteq [0..n)$ and $p = \mathbb{P}\mathbb{r}[u \rightarrow U \mid D_u = 1]$,

$$\frac{1}{2} \min\{1, kp\} \leq \mathbb{P}\mathbb{r}[u \rightarrow U \mid D_u = k] \leq \min\{1, kp\}.$$

FACT 3.2. If v is not a ring-neighbor of u^4 then

$$\frac{1}{\nu\delta(u,v)} \leq \mathbb{P}\mathbb{r}[u \rightarrow v] \leq \frac{2}{\nu\delta(u,v)}.$$

FACT 3.3. If $U = \{u + d \bmod n : d \in [a..b]\}$ or $U = \{u - d \bmod n : d \in [a..b]\}$, where $2 \leq a \leq b \leq n/2$, then

$$\frac{1}{\nu} \ln \frac{b+1}{a} \leq \mathbb{P}\mathbb{r}[u \rightarrow U \mid D_u = 1] \leq \frac{1}{\nu} \ln \frac{b}{a-1}.$$

FACT 3.4. If $U_1, U_2 \subseteq [0..n)$ and $U_1 \cap U_2 = \emptyset$ then

$$\mathbb{P}\mathbb{r}[u \rightarrow U_1 \mid \{D_u = k\} \cap \{u \not\rightarrow U_2\}] = \frac{\mathbb{P}\mathbb{r}[u \rightarrow U_1 \mid D_u = k]}{\mathbb{P}\mathbb{r}[u \not\rightarrow U_2 \mid D_u = 1]}.$$

For the next fact we need to introduce some notation, which we also use throughout the analysis. Let R_x be the set of all nodes at ring-distance at most x from 0; i.e.,

$$R_x = \{u : \|u\| \leq x\}.$$

By \mathbf{H}_u we denote the set of all sets $H \subseteq [0..n) \setminus R_{\|u\|}$, such that for any two distinct $v_1, v_2 \in H$, $\|v_1\| \neq \|v_2\|$. Note that for any graph in \mathbf{G}_n and node s , every prefix path of the routing path from s to 0 that contains no nodes in $R_{\|u\|}$ belongs to \mathbf{H}_u .

FACT 3.5. If $H \in \mathbf{H}_u$ and $d = \min_{v \in H} \delta(u, v)$ then

- (a) $\mathbb{P}\mathbb{r}[u \rightarrow H \mid D_u = k] \leq \mathbb{P}\mathbb{r}[0 \rightarrow [d..d + |H|] \mid D_0 = k]$;
(b) $\mathbb{P}\mathbb{r}[u \not\rightarrow H \mid D_u = k] \geq 2^{-k}$; and (c) $\mathbb{P}\mathbb{r}[u \not\rightarrow H] \geq q_1/2$.

4. PROOF OF THE LOWER BOUNDS

We begin with an auxiliary lemma that bounds from below the average length of any process that approaches 0 with jumps that follow a distribution of a specific form. We use this result in the proofs of all the lower bounds. We prove the lower bound for $\mathcal{G}(n, \alpha)$ in Section 4.1, and for $\mathcal{U}(n, \alpha)$ with $\alpha > 2$, $\alpha < 2$, and $\alpha = 2$ in Sections 4.2–4.4, respectively.

The next lemma provides a lower bound on the expected number of steps of an arbitrary process on the non-negative integers, which is non-increasing, and the length of the jump in each step is bounded by a distribution of a certain form. We will use this result in the proofs of all the lower bounds.

LEMMA 4.1. If $\langle X_0, X_1, \dots \rangle$ is a non-increasing, non-negative, integer-valued random process with $X_0 > \rho \geq 1$, such that for all j with $\rho < j \leq X_0$,

$$\mathbb{P}\mathbb{r}[X_{i+1} = j' \mid X_i = j] \leq \begin{cases} c \frac{(j/j')^\epsilon}{\rho(j-j')}, & \text{if } 0 < j' \leq j - 2; \\ c \frac{1}{\rho \ln j}, & \text{if } j' = 0, \end{cases}$$

where $0 \leq \epsilon < 1$, then the expected number of steps to reach 0 is at least $c' \rho \ln(X_0/\rho)$, where $c' = c'(c, \epsilon) > 0$.

The proof of Lemma 4.1 is similar to the proof of the lower bound for the augmented lattice with fixed degrees, described in [17] (Theorem 7). Roughly, we consider the sequence of $\ln X_i$, show that the average reduction in each step is at most $c''(c, \epsilon)/\rho$, and use an expectation argument to obtain the lower bound. The full proof is omitted due to space limitations.

⁴If v is a ring-neighbor of u then $\mathbb{P}\mathbb{r}[u \rightarrow v] = 0$, by the definition of ‘ \rightarrow ’.

4.1 Proof of Theorem 2.1

It is a straightforward application of Lemma 4.1. Let $\langle Y_0, Y_1, \dots \rangle$ be the routing path from Y_0 to 0 in $\mathcal{G}(n, \alpha)$. For all u, v with $\|v\| \leq \|u\| - 2$,

$$\mathbb{P}\mathbb{r}[Y_{i+1} = v \mid Y_i = u] \leq \mathbb{P}\mathbb{r}[u \rightarrow v] \leq \frac{2}{\nu\delta(u,v)},$$

by Fact 3.2. From this and Lemma 4.1, applied for $n/4 \leq X_0 \leq n/2$, $\rho = \nu$, $\epsilon = 0$, and $X_i = \|Y_i\|$, we obtain that the expected length of the routing path from u to 0 is $\Omega(\nu \ln n) = \Omega(\ln^2 n)$, for all u with $\|u\| \geq n/4$; the theorem then follows.

4.2 Proof of Theorem 2.3 case $\alpha > 2$

We describe a random process \mathcal{N} , which we prove approaches zero faster than GREEDY (Section 4.2.1), and we derive a lower bound on its expected length (Section 4.2.2). Combining these two results we obtain the theorem (Section 4.2.3). Unlike GREEDY, \mathcal{N} is a Markov chain, so, it is easier to analyze.

4.2.1 Process \mathcal{N}

Process \mathcal{N} is parameterized by n , α , and s , where $s \in [0..n)$, and it resembles GREEDY routing in $\mathcal{U}(n, \alpha)$ from source s to target 0. Roughly speaking, \mathcal{N} differs from GREEDY mainly in that: (1) each time the message is forwarded to an *in-contact*, say v , of the current node, the message is next forwarded to an *out-neighbor* of v closest to 0, and these *two* forwardings count as a *single* step of \mathcal{N} ; and (2) the random graph is regenerated in each step of \mathcal{N} . In addition, instead of the contacts of the current node, say u , the out-contacts of a node a_1 and the in-contacts of a (possibly different) node a_2 are used to determine the next node. The a_i are functions on u , have $\|a_i\| \geq \|u\|$, and they are such that they minimize the expected length of \mathcal{N} . We introduce \mathcal{N} because its expected length is a lower bound for the expected steps GREEDY takes to route a message from s to 0, and because \mathcal{N} is a Markov chain, hence, it is easier to analyze than GREEDY. Another useful property of \mathcal{N} is that its expected length is a non-decreasing function of $\|s\|$.

We now define \mathcal{N} formally. Let $a_1 : [0..n) \rightarrow [0..n)$, $A_1 : [0..n) \rightarrow 2^{[0..n)}$, $a_2 : [0..n)^2 \rightarrow [0..n)$, and $A_2 : [0..n)^2 \rightarrow 2^{[0..n)}$ be functions such that for all nodes u, r ,

$$\begin{aligned} \|a_1(u)\| &\geq \|u\|, & A_1(u) &\in \mathbf{H}_{a_1(u)}, \\ \|a_2(u, r)\| &\geq \|u\|, & A_2(u, r) &\in \mathbf{H}_{a_2(u, r)}. \end{aligned}$$

Recall from Section 3.2 that for any graph in \mathbf{G}_n and node u' , \mathbf{H}_u contains every prefix path of the routing path from u' to 0 such that no node in this prefix path is in $R_{\|u\|} = \{v : v \leq \|u\|\}$. The a_i and A_i should also satisfy an additional condition, which we specify later.

Let $u \neq 0$ be the current node in \mathcal{N} . (Initially $u = s$, and \mathcal{N} finishes when $u = 0$.) The next node, denoted N_u , is a node closest to 0 among the two ring-neighbors of u , and the nodes $N_{u,1}, N_{u,2}$ which are determined as follows. First we choose the out-contacts of $a_1(u)$ as in $\mathcal{G}(n, \alpha)$, conditioned on the event $\{a_1(u) \not\rightarrow A_1(u)\}$. We let $N_{u,1}$ be an out-contact of $a_1(u)$ that is closest to 0; or, if $a_1(u)$ has no out-contacts, $N_{u,1}$ is a randomly chosen node among u and the ring-neighbors of u (this ensures that $\mathbb{P}\mathbb{r}[N_{u,1} = v] > 0$, for all v). Suppose that $N_{u,1} = r$. $N_{u,2}$ is then chosen as follows. We choose the out-contacts of the nodes in $R_{\|a_2(u, r)\| - 1}$ as in $\mathcal{G}(n, \alpha)$, conditioned on the event $\{R_{\|a_2(u, r)\| - 1} \not\rightarrow A_2(u, r)\}$. (If $a_1(u) \in R_{a_2(u, r) - 1}$ then the out-links of $a_1(u)$ generated

earlier to determine $N_{u,1}$ are deleted, and replaced by new ones.) Let Z be the set of the in-contacts of $a_2(u, r)$ that are in $R_{\|a_2(u, r)\| - 1}$ and are closest to 0 ($0 \leq |Z| \leq 2$). If $Z = \emptyset$, $N_{u,2} = a_2(u, r)$; if $Z = \{0\}$, $N_{u,2} = 0$; otherwise, $N_{u,2}$ is a node closest to 0 among the out-neighbors of the nodes in Z .

Functions a_i and A_i should satisfy the following optimization condition. Roughly speaking, this condition says that given the values of a_i and A_i for all u with $\|u\| < \|v\|$, their values for $u = v$ are such that they minimize the expected length of \mathcal{N} when starting from $s = u$. Formally, let $L_u^{\mathcal{N}}$ denote the expected number of steps of \mathcal{N} for $s = u$. The condition is described inductively as: for $\|u\| = 1, 2, \dots$,

$$\left\{ \begin{array}{l} \text{for all } r, a_2(u, r) \text{ and } A_2(u, r) \text{ are such that they} \\ \quad \text{minimize } \mathbb{E}[L_u^{\mathcal{N}} \mid N_{u,1} = r]; \\ a_1(u) \text{ and } A_1(u) \text{ are such that they minimize } \mathbb{E}[L_u^{\mathcal{N}}]. \end{array} \right.$$

The next two lemmata state the two properties of \mathcal{N} we described at the beginning, that $\mathbb{E}[L_s^{\mathcal{N}}]$ is a non-decreasing function of $\|s\|$, and it is a lower bound for the expected value of the number of steps L_s that GREEDY requires to route a message from s to 0. Due to space limitations, the proof of Lemma 4.3 is omitted.

LEMMA 4.2. *If $\|u\| \geq \|u'\|$ then $\mathbb{E}[L_u^{\mathcal{N}}] \geq \mathbb{E}[L_{u'}^{\mathcal{N}}]$.*

PROOF. By induction on $\|u\|$. We show that if $\|u\| \geq \|u'\|$ and $\|r\| \geq \min\{\|r'\|, \|u'\| - 1\}$ then

$$\mathbb{E}[L_u^{\mathcal{N}} \mid N_{u,1} = r] \geq \mathbb{E}[L_{u'}^{\mathcal{N}} \mid N_{u',1} = r'], \quad (4.1)$$

$$\mathbb{E}[L_u^{\mathcal{N}}] \geq \mathbb{E}[L_{u'}^{\mathcal{N}}]. \quad (4.2)$$

Clearly, both relations hold if $u' = 0$. Below we assume that $u' \neq 0$ and, thus, $u \neq 0$. The induction hypothesis (i.h.) is that for all v, v', w, w' , such that $\|u\| > \|v\| \geq \|v'\|$ and $\|w\| \geq \min\{\|w'\|, \|v'\| - 1\}$, $\mathbb{E}[L_v^{\mathcal{N}} \mid N_{v,1} = w] \geq \mathbb{E}[L_{v'}^{\mathcal{N}} \mid N_{v',1} = w']$ and $\mathbb{E}[L_v^{\mathcal{N}}] \geq \mathbb{E}[L_{v'}^{\mathcal{N}}]$. From the i.h. it is immediate that

$$\mathbb{E}[L_v^{\mathcal{N}} \mid N_{v,1} = w] = \mathbb{E}[L_v^{\mathcal{N}} \mid N_{v,1} = \|w\|], \quad (4.3)$$

$$\mathbb{E}[L_v^{\mathcal{N}}] = \mathbb{E}[L_{\|v\|}^{\mathcal{N}}]. \quad (4.4)$$

We derive (4.1) as follows. By (4.4),

$$\begin{aligned} & \mathbb{E}[L_u^{\mathcal{N}} \mid N_{u,1} = r] \\ &= 1 + \sum_v \mathbb{E}[L_{\min\{\|u\| - 1, \|r\|, \|v\|\}}^{\mathcal{N}}] \cdot \mathbb{P}\mathbb{r}[N_{u,2} = v \mid N_{u,1} = r]. \end{aligned}$$

For $\mathbb{E}[L_{u'}^{\mathcal{N}} \mid N_{u',1} = r']$, suppose that we compute $L_{u'}^{\mathcal{N}}$ using $a_2(u, r)$ and $A_2(u, r)$ in place of $a_2(u', r')$ and $A_2(u', r')$, respectively, and let M be the resulting quantity. By the optimality of the a_2 and A_2 ,

$$\begin{aligned} & \mathbb{E}[L_{u'}^{\mathcal{N}} \mid N_{u',1} = r'] \leq \mathbb{E}[M \mid N_{u',1} = r'] \\ &= 1 + \sum_v \mathbb{E}[L_{\min\{\|u'\| - 1, \|r'\|, \|v\|\}}^{\mathcal{N}}] \cdot \mathbb{P}\mathbb{r}[N_{u,2} = v \mid N_{u,1} = r]. \end{aligned}$$

From the two results above and the second part of the i.h., we obtain (4.1). We now derive (4.2).

$$\mathbb{E}[L_u^{\mathcal{N}}] = \sum_r \mathbb{E}[L_u^{\mathcal{N}} \mid N_{u,1} = r] \cdot \mathbb{P}\mathbb{r}[N_{u,1} = r].$$

For $\mathbb{E}[L_{u'}^{\mathcal{N}}]$, suppose that when computing $L_{u'}^{\mathcal{N}}$ we replace $a_1(u')$ and $A_1(u')$ by $a_1(u)$ and $A_1(u)$, respectively, and let

M' be the resulting quantity. By the optimality of the a_1 and A_1 , and (4.3) (for $v = u'$),

$$\mathbb{E}[L_{u'}^{\mathcal{N}}] \leq \mathbb{E}[M'] = \sum_r \mathbb{E}[L_{u'}^{\mathcal{N}} \mid N_{u',1} = r] \cdot \mathbb{P}\mathbb{r}[N_{u,1} = r].$$

Combining the two results above and applying the first part of the i.h. (for $v = u$), yields (4.2). \square

LEMMA 4.3. $\mathbb{E}[L_s^{\mathcal{N}}] \leq \mathbb{E}[L_s]$.

4.2.2 Expected length of \mathcal{N}

The next lemma provides lower bounds on the expected length of \mathcal{N} , for $\alpha > 2$.

LEMMA 4.4.

- (a) *If $\alpha > 3$ then $\mathbb{E}[L_{n/4}^{\mathcal{N}}] = \Omega(\ln^2 n)$.*
- (b) *If $\alpha = 3$ then $\mathbb{E}[L_{n/4}^{\mathcal{N}}] = \Omega\left(\frac{\ln^2 n}{\ln \ln n}\right)$.*
- (c) *If $2 < \alpha < 3$ then for $\lambda = e^{\ln^{\alpha-2} n}$, $\mathbb{E}[L_{\lambda}^{\mathcal{N}}] = \Omega\left(\frac{\ln^{\alpha-1} n}{\ln \ln n}\right)$.*

PROOF. (a) We show below that for all u, j such that $0 \leq j \leq \|u\| - 2$,

$$\mathbb{P}\mathbb{r}[\|N_u\| = j] = O\left(\frac{1}{\nu(\|u\| - j)}\right). \quad (4.5)$$

From this and Lemma 4.1, applied for $X_0 = n/4$, $X_{i+1} = \|N_{X_i}\|$, $\rho = \nu$, and $\epsilon = 0$, it follows that $\mathbb{E}[L_{n/4}^{\mathcal{N}}] = \Omega(\ln^2 n)$. We now prove (4.5).

$$\begin{aligned} \mathbb{P}\mathbb{r}[\|N_u\| = j] &\leq \mathbb{P}\mathbb{r}[\|N_{u,1}\| = j] \\ &\quad + \max_r \mathbb{P}\mathbb{r}[\|N_{u,2}\| = j \mid N_{u,1} = r]. \end{aligned} \quad (4.6)$$

Below we will write a_1 and A_1 instead of $a_1(u)$ and $A_1(u)$, respectively.

$$\begin{aligned} \mathbb{P}\mathbb{r}[\|N_{u,1}\| = j] &\leq \mathbb{P}\mathbb{r}[a_1 \rightarrow \{j, n - j\} \mid a_1 \not\rightarrow A_1] \\ &\leq \frac{\mathbb{P}\mathbb{r}[a_1 \rightarrow \{j, n - j\}]}{\mathbb{P}\mathbb{r}[a_1 \not\rightarrow A_1]} = O\left(\frac{1}{\nu(\|a_1\| - j)}\right), \end{aligned} \quad (4.7)$$

by Facts 3.2 and 3.5(c). Next we bound the second term on the right-hand side of (4.6). We will need the following definitions. Let S_v , for $v \neq 0$, be an *out-neighbor* of v in $\mathcal{G}(n, \alpha)$ that is closest to 0 (there may be two such nodes); and $S_0 = 0$. Let also Z_v be the set of the *in-contacts* of v that are in $R_{\|v\| - 1}$ and are closest to 0 ($0 \leq |Z_u| \leq 2$).

$$\begin{aligned} & \mathbb{P}\mathbb{r}[\|N_{u,2}\| = j \mid N_{u,1} = r] \\ &\leq \sum_{v: j \leq \|v\| < \|a_2\|} \sum_k \mathbb{P}\mathbb{r}[\{v \in Z_{a_2}\} \cap \{\|S_v\| = j\} \cap \{D_v = k\} \\ &\quad \mid R_{\|a_2\| - 1} \not\rightarrow A_2], \end{aligned} \quad (4.8)$$

where again we write a_2 and A_2 instead of $a_2(u, r)$ and $A_2(u, r)$, respectively. For $j + 2 \leq \|v\| < \|a_2\|$,

$$\begin{aligned} & \mathbb{P}\mathbb{r}[\|S_v\| = j \mid \{v \in Z_{a_2}\} \cap \{D_v = k\} \cap \{R_{\|a_2\| - 1} \not\rightarrow A_2\}] \\ &= \mathbb{P}\mathbb{r}[\|S_v\| = j \mid \{D_v = k - 1\} \cap \{v \not\rightarrow A_2\}] \\ &\leq \mathbb{P}\mathbb{r}[v \rightarrow \{j, n - j\} \mid \{D_v = k - 1\} \cap \{v \not\rightarrow A_2\}] \\ &= \frac{\mathbb{P}\mathbb{r}[v \rightarrow \{j, n - j\} \mid D_v = k - 1]}{\mathbb{P}\mathbb{r}[v \not\rightarrow A_2 \mid D_v = 1]} \\ &= O\left(\frac{k - 1}{\nu(\|v\| - j)}\right), \end{aligned} \quad (4.9)$$

where the second-to-last line was obtained using Fact 3.4, and the last using Facts 3.1 and 3.5(b); also,

$$\begin{aligned} \Pr[v \in Z_{a_2} \mid \{D_v = k\} \cap \{R_{\|a_2\|-1} \not\rightarrow A_2\}] \\ \leq \Pr[v \rightarrow a_2 \mid \{D_v = k\} \cap \{v \not\rightarrow A_2\}] \\ = O\left(\frac{k}{\nu(\|a_2\| - \|v\|)}\right), \end{aligned} \quad (4.10)$$

similarly to (4.7); and

$$\begin{aligned} \Pr[D_v = k \mid R_{\|a_2\|-1} \not\rightarrow A_2] &= \Pr[D_v = k \mid v \not\rightarrow A_2] \\ &= O(q_k), \end{aligned} \quad (4.11)$$

by Fact 3.5(c). Combining (4.8)–(4.11), we obtain

$$\begin{aligned} \Pr[\|N_{u,2}\| = j \mid N_{u,1} = r] \\ = O\left(\sum_{v:j+2 \leq \|v\| < \|a_2\|} \sum_k \frac{k^{2-\alpha}}{\nu^2(\|v\| - j)(\|a_2\| - \|v\|)} \right. \\ \left. + \frac{1}{\nu(\|a_2\| - j)}\right) \\ = O\left(\frac{\ln(\|a_2\| - j)}{\nu^2(\|a_2\| - j)} + \frac{1}{\nu(\|a_2\| - j)}\right) \\ = O\left(\frac{1}{\nu(\|a_2\| - j)}\right). \end{aligned} \quad (4.12)$$

Applying (4.7) and (4.12) to (4.6), yields (4.5).

(b) We consider an “early-stopping” variance of \mathcal{N} that differs from \mathcal{N} as follows: Let $u \neq 0$ be the current node, suppose $N_{u,1} = r$, and let Z be the set of the in-contacts of $a_2(u, r)$ that are in $R_{\|a_2(u,r)\|-1}$ and are closest to 0 (see the definition of \mathcal{N} in Section 4.2.1); if $D_v > \ln^2 n$ for some $v \in Z$ then the process jumps to node 0 in the next step.

Let M_u denote the next node after node u in this new process, and L_u^M be the number of steps to reach 0 from u . Clearly, $\mathbb{E}[L_u^M] \geq \mathbb{E}[L_u^N]$, so, it suffices to bound $\mathbb{E}[L_{n/4}^M]$.

We show that for all u with $\|u\| \geq \rho = \frac{\ln n}{\ln \ln n}$,

$$\Pr[\|M_u\| = j] = \begin{cases} O\left(\frac{1}{\rho(\|u\| - j)}\right), & \text{if } 0 < j \leq \|u\| - 2; \\ O\left(\frac{1}{\rho \ln n}\right), & \text{if } j = 0. \end{cases} \quad (4.13)$$

From this and Lemma 4.1, applied for $X_0 = n/4$, $X_{i+1} = \|M_{X_i}\|$, and $\epsilon = 0$, it follows that $\mathbb{E}[L_{n/4}^M] = \Omega\left(\frac{\ln^2 n}{\ln \ln n}\right)$. The proof of (4.13) is similar to that of (4.5) and is omitted.

(c) We show that for all u with $\nu < \|u\| \leq \lambda$,

$$\Pr[\|N_u\| = j] = \begin{cases} O\left(\frac{(\|u\|j)^{3-\alpha}}{\nu(\|u\| - j)}\right), & \text{if } 0 < j \leq \|u\| - 2; \\ O\left(\frac{1}{\nu \ln \|u\|}\right), & \text{if } j = 0. \end{cases} \quad (4.14)$$

From this and Lemma 4.1, applied for $X_0 = \lambda$, $X_{i+1} = \|N_{X_i}\|$, $\rho = \nu$, and $\epsilon = 3 - \alpha$, it follows that $\mathbb{E}[L_\lambda^N] = \Omega(\nu \ln \lambda) = \Omega(\ln^{\alpha-1} n)$. The derivation of (4.14) is similar to that of (4.5) — but more computationally involved. The main difference is that a more accurate bound is used in place of (4.9). \square

4.2.3 Putting the pieces together

If $\alpha > 3$ then, by Lemmata 4.4(a), 4.2, and 4.3, $\mathbb{E}[L_u] \geq \mathbb{E}[L_u^N] = \Omega(\ln^2 n)$, for all u with $\|u\| \geq n/4$. Hence, the expected delivery time is $\Omega(\ln^2 n)$. For the cases $\alpha = 3$ and $2 < \alpha < 3$ the theorem follows similarly, using Lemmata 4.4(b) and 4.4(c), respectively, in place of Lemma 4.4(a).

4.3 Proof of Theorem 2.3 case $\alpha < 2$

The theorem follows from the fact that for some λ that is polynomial in n , all nodes in R_λ have out-degree (at most) 1, with probability $\Theta(1)$. Specifically, the probability that $D_u = 1$ is

$$q_1 = \begin{cases} 1 - \Theta(1/k_{\max}^{2-\alpha}), & \text{if } 1 < \alpha < 2; \\ 1 - \Theta(\ln n/k_{\max}), & \text{if } \alpha = 1; \\ 1 - \Theta(1/k_{\max}), & \text{if } 0 \leq \alpha < 1. \end{cases} \quad (4.15)$$

Let $\lambda = \min\{\frac{n}{2}, \frac{1}{1-q_1}\}$, and $\mathcal{E} = \bigcap_{u \in R_\lambda} \{D_u = 1\}$. Then, $\Pr[\mathcal{E}] = q_1^{2\lambda+1} = \Theta(1)$. Let $\langle Y_0, Y_1, \dots \rangle$ be the routing path from λ to 0. If $\|v\| \leq \|u\| - 2$,

$$\begin{aligned} \Pr[Y_{i+1} = v \mid \{Y_i = u\} \cap \{\langle Y_j \rangle_{j=0}^{i-1} = H\} \cap \mathcal{E}] \\ \leq \Pr[\{u \rightarrow v\} \cup \{v \rightarrow u\} \mid \{u, v \not\rightarrow H\} \cap \{D_u = D_v = 1\}] \\ = O\left(\frac{1}{\nu \delta(u, v)}\right), \end{aligned}$$

by Fact 3.5(b). This and Lemma 4.1, applied for $\rho = \nu$, $\epsilon = 0$, and $X_i = \|Y_i\| \mid \mathcal{E}$, yields $\mathbb{E}[L_\lambda \mid \mathcal{E}] = \Omega(\nu \ln \lambda)$. By (4.15) and the fact that k_{\max} is polynomial in n , $\ln \lambda = \Theta(\ln k_{\max}) = \Theta(\ln n)$, so, $\mathbb{E}[L_\lambda \mid \mathcal{E}] = \Omega(\ln^2 n)$. And since $\Pr[\mathcal{E}] = \Theta(1)$, $\mathbb{E}[L_\lambda] \geq \mathbb{E}[L_\lambda \mid \mathcal{E}] \cdot \Pr[\mathcal{E}] = \Omega(\ln^2 n)$; hence, the GREEDY diameter is $\Omega(\ln^2 n)$.

4.4 Proof of Theorem 2.3 case $\alpha = 2$

The proof consists of two parts, which are roughly as follows. First we show that for any λ and node s with $\|s\| \geq \lambda$, with probability $\Theta(1)$, the routing path from s to 0 contains some node u such that $\lambda^{1/3} \leq \|u\| \leq \lambda$ and u has a small expected out-degree. Next, we show that if $\lambda = e^{\ln^{1/3} n}$ then the expected number of remaining steps from u to 0 is $\Omega(\ln^{4/3} n)$. The two lemmata we state below correspond to these two parts. Let $\langle Y_0, Y_1, \dots \rangle$ be the routing path from node Y_0 to 0.

LEMMA 4.5. *If $\|Y_0\| > \lambda = \omega(1)$ and $K = \min\{i : \|Y_i\| \leq \lambda\}$ then*

$$\Pr[\{\|Y_K\| \geq \lambda^{1/3}\} \cap (\{Y_K \not\rightarrow Y_{K-1}\} \cup \{D_{Y_K} = 1\})] = \Theta(1).$$

LEMMA 4.6. *For $\lambda = e^{\ln^{1/3} n}$ and $\lambda^{1/3} \leq \|u\| \leq \lambda$,*

$$\mathbb{E}[L_{Y_i} \mid Y_0, \dots, Y_{i-1}, \{Y_i = u\} \cap (\{Y_i \not\rightarrow Y_{i-1}\} \cup \{D_{Y_i} = 1\})] \text{ is } \Omega(\ln^{4/3} n).$$

Let $\mathcal{E} = \{\|Y_K\| \geq \lambda^{1/3}\} \cap (\{Y_K \not\rightarrow Y_{K-1}\} \cup \{D_{Y_K} = 1\})$. We prove Lemma 4.5 by showing that \mathcal{E} occurs with probability $\Theta(1)$, for any fixed K and Y_0, \dots, Y_{K-1} , and conditionally on the event that for all v with $\|v\| > \lambda$, $D_v \leq \|v\|$. Since this last event occurs with probability $\Theta(1)$, the lemma follows. The proof of Lemma 4.6 is analogous to that of Lemma 4.4(b). We analyze an early-stopping variation of GREEDY, where if in some step we visit a node $v \in R_{\|u\|}$ that has an in-contact $v' \in R_{\|v\|-1}$ with $D_{v'} > 1$ then we jump to 0 in the next step. The full proofs of Lemmata 4.5 and 4.6 are omitted due to space limitations.

The theorem now follows easily. For $\lambda = e^{\ln^{1/3} n}$, K as in the statement of Lemma 4.6, and \mathcal{E} as above,

$$\mathbb{E}[L_{Y_0}] \geq \mathbb{E}[L_{Y_0} \mid \mathcal{E}] \cdot \Pr[\mathcal{E}] \geq \mathbb{E}[L_{Y_K} \mid \mathcal{E}] \cdot \Pr[\mathcal{E}] = \Omega(\ln^{4/3} n),$$

by Lemmata 4.5 and 4.6.

5. PROOF OF THE UPPER BOUNDS

As in the proof of the lower bounds in Section 4, we start with a simple lemma that bound from above the length of any process that approaches 0 with jumps that follow a distribution of a specific form. In Section 5.1, we show that $O(\ln^2 n)$ steps are required in all models. In Sections 5.2 and 5.3, we prove tighter upper bounds for $\mathcal{U}(n, \alpha)$, for $2 < \alpha < 3$ and $\alpha = 2$, respectively.

Lemma 5.1(a) below is an analogue of Lemma 4.1, and we will use it in the proofs of all the upper bounds. Lemma 5.1(b) provides a with-high-probability bound for the length of the process; we will use it in Sections 5.2 and 5.3. The proof is straightforward and is omitted.

LEMMA 5.1. *If $\langle X_0, X_1, \dots \rangle$ is a non-increasing, non-negative, integer-valued random process, such that for all j ,*

$$\Pr[X_{i+1} \leq j/2 \mid X_0, \dots, X_{i-1}, \{X_i = j\}] \geq 1/\rho,$$

then for $\kappa = \lceil \log(X_0 + 1) \rceil$,

- (a) *The expected number of steps to reach 0 is at most $\rho\kappa$.*
- (b) *The number of steps to reach 0 is greater than $t \geq 4\rho\kappa$ with probability at most $e^{-\frac{t}{4\rho}}$.*

We will also use the following simple fact, which we state without proof.

FACT 5.2. *If $Q_1, Q_2, \dots, Q_\kappa$ are independent 0-1 random variables and $Q = \sum_i Q_i$ then (a) $\Pr[Q = 0] \leq e^{-\mathbb{E}[Q]}$; and (b) if for all i , $\mathbb{E}[Q_i] \leq 1/2$ then $\Pr[Q = 0] \geq e^{-\frac{3}{2}\mathbb{E}[Q]}$.*

5.1 Proof of an $O(\ln^2 n)$ bound for all models

Let $\langle Y_0, Y_1, \dots \rangle$ be the routing path from node Y_0 to 0 in $\mathcal{G}(n, \alpha)$ or $\mathcal{U}(n, \alpha)$. We will show that for all u ,

$$\Pr[\|Y_{i+1}\| \leq \|u\|/2 \mid \{Y_j\}_{j=0}^{i-1}, \{Y_i = u\}] = \Omega(1/\nu). \quad (5.1)$$

From this and Lemma 5.1(a), applied for $X_i = \|Y_i\|$ and $\rho = \Theta(\nu)$, we obtain that the expected length of the routing path from Y_0 to 0 is $O(\nu \ln(\|Y_0\| + 1)) = O(\ln^2 n)$. We now prove (5.1). For $\|u\| \leq 2$, it obviously holds; so, suppose that $\|u\| > 2$.

- In $\mathcal{G}(n, \alpha)$, the left-hand side of (5.1) equals

$$\Pr[u \rightarrow R_{\|u\|/2}] \geq q_1 \Pr[u \rightarrow R_{\|u\|/2} \mid D_u = 1] = \Omega\left(\frac{1}{\nu}\right),$$

by Fact 3.3.

- In $\mathcal{U}(n, \alpha)$, we have

$$\begin{aligned} \Pr[\|Y_{i+1}\| \leq \|u\|/2 \mid \{Y_i = u\} \cap \{Y_j\}_{j=0}^{i-1} = H] \\ \geq \Pr[R_{\|u\|/2} \rightarrow u \mid R_{\|u\|/2} \not\rightarrow H]. \end{aligned} \quad (5.2)$$

But for any $v \in R_{\|u\|/2}$,

$$\begin{aligned} \Pr[v \rightarrow u \mid R_{\|u\|/2} \not\rightarrow H] \\ = \Pr[v \rightarrow u \mid v \not\rightarrow H] \geq \Pr[\{v \rightarrow u\} \cap \{v \not\rightarrow H\}] \\ \geq \Pr[\{v \rightarrow u\} \cap \{D_v = 1\}] = q_1 \Pr[v \rightarrow u \mid D_v = 1] \\ \geq \frac{q_1}{2\nu\|u\|}. \end{aligned}$$

So, $\sum_{v \in R_{\|u\|/2}} \Pr[v \rightarrow u \mid R_{\|u\|/2} \not\rightarrow H] \geq \frac{1}{4\nu}$; and since the events $\{v \rightarrow u\}$ are independent (conditionally on $R_{\|u\|/2} \not\rightarrow H$), we have, by Fact 5.2(a), that $\Pr[R_{\|u\|/2} \rightarrow u \mid R_{\|u\|/2} \not\rightarrow H] \geq 1 - e^{-\frac{1}{4\nu}} = \Theta(1/\nu)$. Combining this and (5.2) yields (5.1).

5.2 Proof of Theorem 2.4 case $2 < \alpha < 3$

We will use the following result, which is analogous to Lemma 5.1(a).

LEMMA 5.3. *If $\langle X_0, X_1, \dots \rangle$ is a non-increasing, non-negative, integer-valued random process with $X_0 > \lambda \geq 1$, such that for all j with $\lambda < j \leq X_0$,*

$$\Pr[X_{i+1} \leq j/2 \mid X_0, \dots, X_{i-1}, \{X_i = j\}] \geq \frac{\log j}{\rho} \quad (5.3)$$

then the expected number of steps until the process' value reduces to at most λ is at most $\rho(\ln \log X_0 + 1)$.

PROOF. Let T_k , for $k \geq 0$, be the number of steps until the process' value is reduced to at most $2^k \lambda$; i.e., $T_k = \min\{i : X_i \leq 2^k \lambda\}$. (Note that smaller k correspond to larger T_k .) To prove the lemma we must show that $\mathbb{E}[T_0] \leq \rho(\ln \log X_0 + 1)$. For $k \geq \log X_0 - \log \lambda$, $T_k = 0$. For $0 \leq k < \log X_0 - \log \lambda$,

$$\begin{aligned} \Pr[T_k = i + 1 \mid X_0, \dots, X_i, \{T_{k+1} \leq i < T_k\}] \\ = \Pr[X_{i+1} \leq 2^k \lambda \mid X_0, \dots, X_i, \{2^k \lambda < X_i \leq 2^{k+1} \lambda\}] \\ \geq \frac{\log \lambda + k}{\rho}, \end{aligned}$$

by (5.3). So, $T_k - T_{k+1}$ is stochastically smaller than a geometric random variable with probability parameter $\frac{\log \lambda + k}{\rho}$. Therefore,

$$\begin{aligned} \mathbb{E}[T_0] &= \mathbb{E} \left[\sum_{k=0}^{\log X_0 - \log \lambda - 1} (T_k - T_{k+1}) \right] \\ &\leq \sum_{k=0}^{\log X_0 - \log \lambda - 1} \frac{\rho}{\log \lambda + k} \leq \rho(\ln \log X_0 + 1). \quad \square \end{aligned}$$

Roughly, the proof of the theorem proceeds as follows. We show that in every three steps of GREEDY the ring-distance to 0 is halved with probability $\Omega\left(\frac{\ln \|u\|}{\ln^{\alpha-1} n}\right)$, provided that we are not too close to 0 and not too many steps have been taken so far. Also, by the analysis in Section 5.1, the ring-distance to 0 is halved with probability $\Omega(1/\ln n)$ in each step, independently of the previous steps. By combining these two results and applying Lemmata 5.1 and 5.3 we obtain the theorem.

The next lemma gives a lower bound on the speed of GREEDY when the length of the prefix of the routing path so far is much smaller than the current ring-distance to the target. Two steps at a time are considered instead of just one. Interestingly, this bound is obtained by counting only the contribution of nodes with out-degree $\Theta(\ln n)$. Let $\langle Y_0, Y_1, \dots \rangle$ be the routing path from node Y_0 to 0.

LEMMA 5.4. *If $\|u\| \geq 8(i^2 + 1)$ then*

$$\Pr[\|Y_{i+2}\| \leq \|u\|/2 \mid Y_0, \dots, Y_{i-1}, \{Y_i = u\} \cap \{Y_i \not\rightarrow Y_{i-1}\}]$$

is $\Omega\left(\frac{\ln \|u\|}{\ln^{\alpha-1} n}\right)$.

PROOF. We describe an event \mathcal{E} such that if \mathcal{E} occurs and $Y_i = u$ then $\|Y_{i+2}\| \leq \|u\|/2$, and we bound \mathcal{E} 's conditional probability instead. Informally, if $\mathcal{E} \cap \{Y_i = u\}$ occurs then the following statements are true about Y_{i+1} : (1) it is an in-contact of Y_i ; (2) it has out-degree $\Theta(\ln n)$; (3) $\|u\|/2 < \|Y_{i+1}\| \leq \|u\| - \|u\|^{1/2}$; and (4) at least one of its out-contacts is in $R_{\|u\|/2}$. Formally, we define the following four events. Let

$$\mathcal{E}_0 = \{u \not\rightarrow R_{\|u\| - \|u\|^{1/2}}\}.$$

Define the sets $C = R_{\|u\| - \|u\|^{1/2}} \setminus R_{\|u\|/2}$ and $C^* = \{v \in C : \nu \leq D_v \leq 2\nu\}$, and let

$$\mathcal{E}_1 = \{C^* \rightarrow u\}, \quad \mathcal{E}_2 = \{R_{\|u\|-1} \setminus C^* \not\rightarrow u\}.$$

Last, if \mathcal{E}_1 occurs, let Z be the in-contact of u in C^* that is closest to 0 (if there are two such nodes then Z is the one that GREEDY would choose), and let

$$\mathcal{E}_3 = \{Z \rightarrow R_{\|u\|/2}\}.$$

We define

$$\mathcal{E} = \mathcal{E}_0 \cap \mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3.$$

It is easy to see that $\mathcal{E} \cap \{Y_i = u\} \subseteq \{\|Y_{i+2}\| \leq \|u\|/2\}$. So, to prove the lemma it suffices to show that

$$\mathbb{P}\mathbb{r}[\mathcal{E} \mid \{(Y_j)_{j=0}^{i-1} = H\} \cap \{Y_i = u\} \cap \{Y_i \not\rightarrow Y_{i-1}\}] = \Omega\left(\frac{\ln \|u\|}{\ln^{\alpha-1} n}\right).$$

But the left-hand side is equal to $\mathbb{P}\mathbb{r}[\mathcal{E} \mid R_{\|u\|} \not\rightarrow H]$; so, we will show that $\mathbb{P}\mathbb{r}[\mathcal{E} \mid R_{\|u\|} \not\rightarrow H] = \Omega\left(\frac{\ln \|u\|}{\ln^{\alpha-1} n}\right)$. Let $\mathcal{A} = \{R_{\|u\|} \not\rightarrow H\}$. Since \mathcal{E}_0 is independent of the other three events,

$$\begin{aligned} \mathbb{P}\mathbb{r}[\mathcal{E} \mid \mathcal{A}] &= \mathbb{P}\mathbb{r}[\mathcal{E}_0 \mid \mathcal{A}] \cdot \mathbb{P}\mathbb{r}[\mathcal{E}_1 \mid \mathcal{A}] \cdot \mathbb{P}\mathbb{r}[\mathcal{E}_2 \mid \mathcal{E}_1 \cap \mathcal{A}] \\ &\quad \cdot \mathbb{P}\mathbb{r}[\mathcal{E}_3 \mid \mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{A}]. \end{aligned} \quad (5.4)$$

We compute lower bounds for the four probabilities on the right-hand side.

$$\begin{aligned} \mathbb{P}\mathbb{r}[\mathcal{E}_0 \mid \mathcal{A}] &\geq \mathbb{P}\mathbb{r}[u \not\rightarrow R_{\|u\| - \|u\|^{1/2}} \cup H] \\ &\geq q_1 \mathbb{P}\mathbb{r}[u \not\rightarrow R_{\|u\| - \|u\|^{1/2}} \cup H \mid D_u = 1] \\ &\geq q_1 \mathbb{P}\mathbb{r}[0 \rightarrow [2 \cdot \|u\|^{1/2}] \mid D_0 = 1] \\ &\quad + q_1 \mathbb{P}\mathbb{r}[0 \rightarrow [|H| + 2 \cdot n/2 - \|u\|] \mid D_0 = 1] \\ &= \Omega\left(\frac{1}{\nu} \ln \|u\|^{1/2} + \frac{1}{\nu} \ln \frac{n/2 - \|u\| + 1}{|H| + 1}\right) \\ &= \Omega(1). \end{aligned} \quad (5.5)$$

The third relation was obtained using Fact 3.5(a); the second-to-last was obtained using Fact 3.3; the last using the facts that $\|u\| + |H| < n/2 + 1$ and $\|u\| > i = |H|$. Next we bound $\mathbb{P}\mathbb{r}[\mathcal{E}_1 \mid \mathcal{A}]$. Let \mathcal{D}_v denote the event $\{\nu \leq D_v \leq 2\nu\}$, and Q_v be the indicator random variable of the event $\mathcal{D}_v \cap \{v \rightarrow u\}$. For all $v \in C$,

$$\begin{aligned} \mathbb{E}[Q_v \mid \mathcal{A}] &= \mathbb{P}\mathbb{r}[v \rightarrow u \mid \mathcal{D}_v \cap \{v \not\rightarrow H\}] \\ &\quad \cdot \frac{\mathbb{P}\mathbb{r}[v \not\rightarrow H \mid \mathcal{D}_v]}{\mathbb{P}\mathbb{r}[v \not\rightarrow H]} \cdot \mathbb{P}\mathbb{r}[\mathcal{D}_v] \\ &= \Omega\left(\frac{1}{\nu^{\alpha-1} \delta(u, v)}\right), \end{aligned} \quad (5.6)$$

where the last relation holds because: $\mathbb{P}\mathbb{r}[v \rightarrow u \mid \mathcal{D}_v \cap \{v \not\rightarrow H\}] \geq \mathbb{P}\mathbb{r}[v \rightarrow u \mid D_v = \nu] = \Theta\left(\frac{1}{\delta(u, v)}\right)$; $\mathbb{P}\mathbb{r}[v \rightarrow H \mid \mathcal{D}_v] \leq |H| \cdot \frac{2\nu}{\nu \|u\|^{1/2}} \leq \frac{1}{\sqrt{2}}$, since $\|u\| > 8|H|^2$; and $\mathbb{P}\mathbb{r}[\mathcal{D}_v] = \Theta\left(\frac{1}{\nu^{\alpha-1}}\right)$. So, for $Q = \sum_{v \in C} Q_v$,

$$\mathbb{E}[Q \mid \mathcal{A}] = \sum_{v \in C} \mathbb{E}[Q_v \mid \mathcal{A}] = \Omega\left(\frac{\ln \|u\|}{\nu^{\alpha-1}}\right).$$

And since the Q_i are independent (conditionally on \mathcal{A}), we have, by Fact 5.2(a), that $\mathbb{P}\mathbb{r}[Q \neq 0 \mid \mathcal{A}] \geq 1 - e^{-\Omega\left(\frac{\ln \|u\|}{\nu^{\alpha-1}}\right)} = \Omega\left(\frac{\ln \|u\|}{\nu^{\alpha-1}}\right)$. Finally, since $\mathcal{E}_1 = \{Q \neq 0\}$,

$$\mathbb{P}\mathbb{r}[\mathcal{E}_1 \mid \mathcal{A}] = \Omega\left(\frac{\ln \|u\|}{\nu^{\alpha-1}}\right). \quad (5.7)$$

Next,

$$\begin{aligned} \mathbb{P}\mathbb{r}[\mathcal{E}_2 \mid \mathcal{E}_1 \cap \mathcal{A}] &\geq \mathbb{P}\mathbb{r}[\mathcal{E}_2 \mid \mathcal{A}] \geq \mathbb{P}\mathbb{r}[R_{\|u\|-1} \not\rightarrow u \mid \mathcal{A}] \\ &= \Theta(1), \end{aligned} \quad (5.8)$$

where the last relation is obtained as follows. For all $v \in R_{\|u\|-1}$,

$$\mathbb{P}\mathbb{r}[v \rightarrow u \mid \mathcal{A}] = \mathbb{P}\mathbb{r}[v \rightarrow u \mid v \not\rightarrow H] = O(\mathbb{P}\mathbb{r}[v \rightarrow u]),$$

by Fact 3.5(c); so, $\sum_{v \in R_{\|u\|-1}} \mathbb{P}\mathbb{r}[v \rightarrow u \mid \mathcal{A}] = O(\sum_v \mathbb{P}\mathbb{r}[v \rightarrow u]) = O(1)$, since the expected number of in-contacts a node has is constant; and since the events $\{v \rightarrow u\}$ are independent, we have, by Fact 5.2(b), that $\mathbb{P}\mathbb{r}[R_{\|u\|-1} \not\rightarrow u \mid \mathcal{A}] \geq e^{-O(1)} = \Theta(1)$. The last bound we need is

$$\begin{aligned} \mathbb{P}\mathbb{r}[\mathcal{E}_3 \mid \mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{A}] &\geq \mathbb{P}\mathbb{r}[u \rightarrow R_{\|u\|/2} \mid D_u = \nu] \\ &= \Theta(1), \end{aligned} \quad (5.9)$$

by Facts 3.1 and 3.3. Combining (5.4), (5.5), (5.7), (5.8), and (5.9), yields $\mathbb{P}\mathbb{r}[\mathcal{E} \mid \mathcal{A}] = \Omega\left(\frac{\ln \|u\|}{\nu^{\alpha-1}}\right)$. \square

We will now use Lemma 5.4 to show that if $\|u\| > \lambda = e^{\ln^{\alpha-2} n}$ and $i = o(\lambda^{1/2})$ then

$$\mathbb{P}\mathbb{r}[\|Y_{i+3}\| \leq \|u\|/2 \mid Y_0, \dots, Y_{i-1}, \{Y_i = u\}] = \Omega\left(\frac{\ln \|u\|}{\nu^{\alpha-1}}\right). \quad (5.10)$$

By Lemma 5.4, it suffices to show that

$$\begin{aligned} \mathbb{P}\mathbb{r}[\|Y_{i+3}\| \leq \|u\|/2 \mid Y_0, \dots, Y_{i-1}, \{Y_i = u\} \cap \{Y_i \rightarrow Y_{i-1}\}] \\ \text{is } \Omega\left(\frac{\ln \|u\|}{\nu^{\alpha-1}}\right). \text{ Let } \mathcal{H} \text{ be the } \sigma\text{-algebra generated by } Y_0, \dots, \\ Y_{i-1}, \{Y_i = u\} \cap \{Y_i \rightarrow Y_{i-1}\}. \\ \mathbb{P}\mathbb{r}[\|Y_{i+3}\| \leq \|u\|/2 \mid \mathcal{H}] \\ \geq \mathbb{P}\mathbb{r}[\|Y_{i+3}\| \leq \|u\|/2 \mid \mathcal{H}, \{Y_{i+1} \not\rightarrow Y_i\}] \cdot \mathbb{P}\mathbb{r}[Y_{i+1} \not\rightarrow Y_i \mid \mathcal{H}] \\ \geq \mathbb{P}\mathbb{r}[\|Y_{i+3}\| \leq \|u\|/2 \mid \mathcal{H}, \{Y_{i+1} \not\rightarrow Y_i\} \cap \{\|Y_{i+1}\| > \|u\|/2\}] \\ \quad \cdot \mathbb{P}\mathbb{r}[R_{\|u\|-1} \not\rightarrow u \mid \mathcal{H}] \\ = \Omega\left(\frac{\ln \|u\|}{\nu^{\alpha-1}}\right), \end{aligned}$$

because the probability in the third line is $\Omega\left(\frac{\ln(\|u\|/2)}{\ln^{\alpha-1} n}\right)$, by Lemma 5.4, and the probability in the second-to-last line is $\Theta(1)$, similarly to the last relation in (5.8). We can now obtain the theorem as follows:

- If $Y_0 \leq \lambda$ then, by (5.1) and Lemma 5.1(a), applied for $X_i = \|Y_i\|$ and $\rho = \Theta(\ln n)$, we have $\mathbb{E}[L_{Y_0}] = O(\ln n \ln \lambda) = O(\ln^{\alpha-1} n)$.
- If $Y_0 > \lambda$, let T_1 be the number of steps from Y_0 until we reach a node within ring-distance λ of 0, and T_2 be the number of remaining steps to 0. Similarly to the case $Y_0 \leq \lambda$, $\mathbb{E}[T_2] = O(\ln^{\alpha-1} n)$; so, we just need to show that $\mathbb{E}[T_1] = O(\ln^{\alpha-1} n \ln \ln n)$. By (5.1) and Lemma 5.1(b), applied for $X_i = \|Y_i\|$, $\rho = \Theta(\ln n)$, and $t = 4\rho \ln n$, we have that

$$\mathbb{P}\mathbb{r}[T_1 \geq \ln^3 n] < 1/n.$$

Also, by (5.10) and Lemma 5.3, applied for $X_i = \|Y_{3i}\|$ if $i < \ln^3 n/3$, $X_i = 0$ if $i \geq \ln^3 n/3$, and $\rho = \Theta(\ln^{\alpha-1} n)$, we obtain that $\mathbb{E}[\min\{T_1, \ln^3 n\}] = O(\ln^{\alpha-1} n \ln \ln n)$; so,

$$\begin{aligned} \mathbb{E}[T_1 \mid T_1 < \ln^3 n] &= \mathbb{E}[\min\{T_1, \ln^3 n\} \mid T_1 < \ln^3 n] \\ &\leq \mathbb{E}[\min\{T_1, \ln^3 n\}] \\ &= O(\ln^{\alpha-1} n \ln \ln n) \end{aligned}$$

Therefore,

$$\begin{aligned}\mathbb{E}[T_1] &= \mathbb{E}[T_1 \mid T_1 < \ln^3 n] \cdot \Pr[T_1 < \ln^3 n] \\ &\quad + \mathbb{E}[T_1 \mid T_1 \geq \ln^3 n] \cdot \Pr[T_1 \geq \ln^3 n] \\ &\leq \mathbb{E}[T_1 \mid T_1 < \ln^3 n] + n \Pr[T_1 \geq \ln^3 n] \\ &= O(\ln^{\alpha-1} n \ln \ln n).\end{aligned}$$

5.3 Proof of Theorem 2.4 case $\alpha = 2$

It is similar to the proof of case $2 < \alpha < 3$. The next two lemmata are the analogues of Lemmata 5.3 and 5.4, respectively. Their proofs are omitted.

LEMMA 5.5. *If $\langle X_0, X_1, \dots \rangle$ is a non-increasing, non-negative, integer-valued random process with $X_0 > \lambda \geq 2$, such that for all j with $\lambda < j \leq X_0$,*

$$\Pr[X_{i+1} \leq j^{1-\epsilon} \mid X_0, \dots, X_{i-1}, \{X_i = j\}] \geq \frac{\log^2 j}{\rho},$$

where $0 < \epsilon < 1$, then the expected number of steps until the process' value reduces to at most λ is at most $c \frac{\rho}{\log^2 \lambda}$, where $c = c(\epsilon)$.

Let $\langle Y_0, Y_1, \dots \rangle$ be the routing path from Y_0 to 0.

LEMMA 5.6. *If $\|u\| \geq 4^6(i^6 + 1)$ then*

$$\Pr[\|Y_{i+2}\| \leq \|u\|^{2/3} \mid Y_0, \dots, Y_{i-1}, \{Y_i = u\} \cap \{Y_i \neq Y_{i-1}\}]$$

is $\Omega\left(\frac{\ln^2 \|u\|}{\ln^2 n}\right)$.

Note that, unlike in case $2 < \alpha < 3$ where only nodes of out-degree $\Theta(\ln n)$ contribute to routing significantly, now the contribution of nodes with out-degrees in a wider range is significant.

The rest of the proof is completely analogous to that of case $2 < \alpha < 3$. Instead of (5.10), we show (using Lemma 5.6) that if $\|u\| > \lambda = e^{\sqrt{\ln n}}$ and $i = o(\lambda^{1/6})$ then

$$\Pr[\|Y_{i+3}\| \leq \|u\|^{2/3} \mid Y_0, \dots, Y_{i-1}, \{Y_i = u\}] = \Omega\left(\frac{\ln^2 \|u\|}{\ln^2 n}\right);$$

and instead of Lemma 5.3, we use Lemma 5.5, for $\rho = \Theta(\ln^2 n)$ and $\epsilon = 1/3$.

6. ACKNOWLEDGMENTS

The second author thanks Vassos Hadzilacos and Philipp Woelfel for helpful discussions.

7. REFERENCES

- [1] L. A. Adamic, R. M. Lukose, A. R. Puniyani, and B. A. Huberman. Search in power-law networks. *Phys. Rev. E*, 64:46135, 2001.
- [2] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47–97, 2002.
- [3] B. Bollobás and O. Riordan. The diameter of a scale-free random graph. *Combinatorica*, 24(1):5–34, 2004.
- [4] A. Chaintreau, P. Fraigniaud, and E. Lebar. Networks become navigable as nodes move and forget. In *Proc. 35th Int. Colloq. on Automata, Languages and Programming (ICALP)*, pages 133–144, 2008.
- [5] F. R. K. Chung and L. Lu. The average distance in a random graph with given expected degrees. *Internet Mathematics*, 1(1):91–113, 2003.
- [6] A. Clauset and C. Moore. How do networks become navigable? preprint at arxiv.org, 2003.
- [7] M. Dietzfelbinger and P. Woelfel. Tight lower bounds for greedy routing in uniform small world rings. In *Proc. 41st ACM Symp. on Theory of Computing (STOC)*, 2009.
- [8] P. S. Dodds, R. Muhamad, and D. J. Watts. An experimental study of search in global social networks. *Science*, 301(5634):827–829, 2003.
- [9] S. N. Dorogovtsev and J. F. F. Mendes. *Evolution of networks: From biological networks to the Internet and WWW*. Oxford University Press, 2003.
- [10] P. Duchon, N. Hanusse, E. Lebar, and N. Schabanel. Could any graph be turned into a small-world? *Theoretical Computer Science*, 355(1):96–103, 2006.
- [11] P. Fraigniaud, P. Gauron, and M. Latapy. Combining the use of clustering and scale-free nature of user exchanges into a simple and efficient P2P system. In *Proc. 11th European Conf. on Parallel Computing (Euro-Par)*, pages 1163–1172, 2005.
- [12] B. Kim, C. Yoon, S. Han, and H. Jeong. Path finding strategies in scale-free networks. *Phys. Rev. E*, 65:027103, 2002.
- [13] J. Kleinberg. The small-world phenomenon: An algorithm perspective. In *Proc. 32nd ACM Symp. on Theory of Computing (STOC)*, pages 163–170, 2000.
- [14] J. Kleinberg. Complex networks and decentralized search algorithms. In *Proc. International Congress of Mathematicians (ICM)*, 2006.
- [15] R. Kumar, D. Liben-Nowell, and A. Tomkins. Navigating low-dimensional and hierarchical population networks. In *Proc. 14th Ann. European Symp. on Algorithms (ESA)*, pages 480–491, 2006.
- [16] D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, and A. Tomkins. Geographic routing in social networks. *Proc. National Academy of Sciences of the USA*, 102(33):11623–11628, 2005.
- [17] C. Martel and V. Nguyen. Analyzing Kleinberg's (and other) small-world models. In *Proc. 23rd ACM Symp. on Principles of Distributed Computing (PODC)*, pages 179–188, 2004.
- [18] S. Milgram. The small world problem. *Psychology Today*, 67(1):60–67, 1967.
- [19] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45(2):167–256, 2003.
- [20] O. Sandberg and I. Clarke. The evolution of navigable small-world networks. preprint at arxiv.org, 2006.
- [21] N. Sarshar, P. O. Boykin, and V. P. Roychowdhury. Percolation search in power law networks: Making unstructured peer-to-peer networks scalable. In *Proc. 4th Int. Conf. on Peer-to-Peer Computing*, pages 2–9, 2004.
- [22] Ö. Simsek and D. Jensen. Decentralized search in networks using homophily and degree disparity. In *Proc. 19th Int. Joint Conf. on Artificial Intelligence (IJCAI)*, pages 304–310, 2005.
- [23] A. Slivkins. Distance estimation and object location via rings of neighbors. In *Proc. 24th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 41–50, 2005.