

Greedy routing in small-world networks with power-law degrees

Pierre Fraigniaud, George Giakkoupis

► **To cite this version:**

Pierre Fraigniaud, George Giakkoupis. Greedy routing in small-world networks with power-law degrees. Distributed Computing, Springer Verlag, 2014, 27 (4), pp.231 - 253. 10.1007/s00446-014-0210-y . hal-01097141

HAL Id: hal-01097141

<https://hal.inria.fr/hal-01097141>

Submitted on 19 Dec 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Greedy routing in small-world networks with power-law degrees

Pierre Fraigniaud · George Giakkoupis

Received: date / Accepted: date

Abstract In this paper we study decentralized routing in small-world networks that combine a wide variation in node degrees with a notion of spatial embedding. Specifically, we consider a variant of J. Kleinberg’s grid-based small-world model in which (1) the number of long-range edges of each node is not fixed, but is drawn from a power-law probability distribution with exponent parameter $\alpha \geq 0$ and constant mean, and (2) the long-range edges are considered to be bidirectional for the purposes of routing. This model is motivated by empirical observations indicating that several real networks have degrees that follow a power-law distribution. The measured power-law exponent α for these networks is often in the range between 2 and 3. For the small-world model we consider, we show that when $2 < \alpha < 3$ the standard greedy routing algorithm, in which a node forwards the message to its neighbor that is closest to the target in the grid, finishes in an expected number of $O(\log^{\alpha-1} n \cdot \log \log n)$ steps, for any source–target pair. This is asymptotically smaller than the $O(\log^2 n)$ steps needed in Kleinberg’s original model with the same average degree, and approaches $O(\log n)$ as α approaches 2. Further, we

show that when $0 \leq \alpha < 2$ or $\alpha \geq 3$ the expected number of steps is $O(\log^2 n)$, while for $\alpha = 2$ it is $O(\log^{4/3} n)$. We complement these results with lower bounds that match the upper bounds within at most a $\log \log n$ factor.

Keywords Small worlds · Social networks · Routing · Search · Power-law degrees

1 Introduction

The study of small-world networks was initiated by the famous “six-degrees-of-separation” experiments conducted by Milgram in the 1960s [33]. These experiments quantified the so-called “small-world phenomenon,” that is, the principle that almost all people are linked by short chains of acquaintances. Milgram’s findings have been subsequently confirmed by other experiments and measurements, e.g., by Dodds, Muhamad, and Watts [12], and Backstrom, Boldi, Rosa, Ugander, and Vigna [5]. Further, it has been observed that several real networks, including social, information, technological, and biological networks, exhibit similar small-world properties; see, e.g., the surveys by Albert and Barabási [3] and Newman [34], and the book by Dorogovtsev and Mendes [13].

A striking aspect of Milgram’s experiments, pointed out by J. Kleinberg [24, 23], is that not only do short chains between people exist, but individuals are collectively very effective at finding them using only local information. To study this algorithmic aspect of the small-world phenomenon Kleinberg proposed a simple random graph model, building upon a small-world model proposed by Watts and Strogatz [38]. In Kleinberg’s model, individuals are nodes at the lattice points of a two-dimensional $n \times n$ square lattice, and acquaintance relationships between individuals are represented by directed edges. Each node u has edges to all nodes whose lattice distance from u is at most r , for some

This paper was originally invited to the special issue of Distributed Computing based on selected papers presented at PODC 2009. It appears separately due to publication delays.

P. Fraigniaud
LIAFA, Universit Paris Diderot – Paris 7, Case 7014
75205 Paris Cedex 13, France
Tel.: +33 1 57 27 92 60
Fax: +33 1 57 27 94 09
E-mail: pierre.fraigniaud@liafa.univ-paris-diderot.fr

G. Giakkoupis
IRISA/INRIA Rennes – Bretagne Atlantique
Campus Universitaire de Beaulieu
35042 Rennes Cedex, France
Tel.: +33 2 99 84 71 96
Fax: +33 2 99 84 71 71
E-mail: george.giakkoupis@inria.fr

constant r . Further, u has k random *long-range edges*, where k is another constant parameter. Each of these k edges points to a node chosen independently at random according to the 2-harmonic probability distribution, i.e., each node v is chosen with probability proportional to $1/(d_{u,v})^2$, where $d_{u,v}$ is the distance between u and v in the lattice. Kleinberg showed that a simple greedy routing algorithm, in which a node forwards the message to its neighbor that is closest to the target in the lattice, routes a message in an expected number of $O(\log^2 n)$ steps, for any source–target pair. (This bound was subsequently shown to be tight for a pair chosen uniformly at random, by Barrière, Fraigniaud, Kranakis, and Krizanc [6], and Martel and Nguyen [31].) Note that the above greedy routing algorithm is decentralized, in the sense that it does not require knowledge of the long-range edges of nodes not yet visited. Kleinberg showed also that *any* decentralized routing algorithm needs an expected number of steps that is a polynomial function in n , if the long-range edges are chosen from the h -harmonic distribution for a constant $h \neq 2$.¹

The above results readily extend to the analogous model based on the ℓ -dimensional lattice, for any constant $\ell \geq 1$. In this model, the $O(\log^2 n)$ expected routing time is achieved when the ℓ -harmonic distribution is used to choose the long-range edges. Further, similar results have been shown for several variants and generalizations of this model, where base structures as general as metrics of bounded doubling dimension are used in place of the lattice (cf. Related Work). The probability distribution used to choose the long-range edges in those models is similar to the harmonic distribution considered by Kleinberg [24]. Specifically, it is a variant of the following natural distribution: the probability that node u has a long-range edge to node v is inversely proportional to the number of nodes contained in the smallest ball (in the underlying metric space) that is centered at u and contains v . Interestingly, empirical results by Liben-Nowell, Novak, Kumar, Raghavan, and Tomkins [29] indicate that two-thirds of friendships are geographically distributed that way, i.e., “the probability of befriending a particular person is inversely proportional to the number of closer people.”

Kleinberg’s model and subsequent generalizations of it do not take into account the well-established fact that social and other real networks have a highly skewed distribution of node degrees. It has been observed that the degree distribution of these networks is a *power law*, i.e., the probability that a node has degree k is proportional to $1/k^\alpha$ [3, 13, 34]. Further, the value of the power-law exponent α has been measured to be between 2 and 3 for several real network, including social networks (e.g., the collaboration network

of film actors, and networks of email messages), samples of the Web and the Internet, various peer-to-peer networks, and metabolic and protein interaction biological networks (see [34, Table II]). A straightforward way to reconcile Kleinberg’s model with a power-law degree distribution is to choose the *number* of long-range edges of each node independently at random from that distribution; this approach was first proposed by Kleinberg in [26].² It is reasonable to expect that a power-law degree distribution can reduce the network diameter or the average length of shortest paths between nodes, e.g., similar to the works of Bollobás and Rioridan [7], and Chung and Lu [8]. However, prior to our work there were no results suggesting that power-law degree distributions could improve the speed of greedy routing.

1.1 Our contribution

We consider a simple variant of Kleinberg’s ℓ -dimensional small-world model [24], in which nodes have a power-law degree distribution. In this model each long-range edge is drawn independently at random from the ℓ -harmonic distribution, i.e., the distribution that yields an expected routing time of $O(\log^2 n)$ in Kleinberg’s model. The number of long-range edges of each node is drawn independently at random from a power-law distribution with exponent parameter $\alpha \geq 0$ and a fixed constant expected value. Further, we assume that each node has at least one long-range edge. (For a precise description of the model see Section 2.) For this network, we study the complexity of the same greedy routing algorithm considered by Kleinberg, except that we treat long-range edges as *bidirectional*; i.e., a node forwards the message to its out- or *in*-neighbor that is closest to the target in the grid. If we treat long-range edges as unidirectional, then we observe that greedy routing performs asymptotically the same as in Kleinberg’s original model.³ Having bidirectional long-range edges is qualitatively different, because in this case a node of high degree is easy to find, while if long-range edges are unidirectional then a node of high out-degree may have low in-degree and thus it may be difficult to find.

We now summarize our results (see also Table 1, and Figure 1). For the case of $2 < \alpha < 3$, which is the case for many real networks including social networks, we show that routing finishes in an expected number of $O(\log^{\alpha-1} n \cdot \log \log n)$ steps for any source–target pair of nodes. This is asymptotically smaller than the $O(\log^2 n)$ steps needed in Kleinberg’s original model with the same average number of long-range edges per node, and it approaches $O(\log n)$ as

¹ Paths of polylogarithmic length exist between nodes for a wide range of values for parameter h , as shown by Martel and Nguyen [31, 32]. However, short paths can be efficiently discovered by a decentralized algorithm only when $h = 2$.

² The extent to which this model resembles real social networks has yet to be evaluated empirically.

³ For the case of $\ell = 1$, this follows also from a general lower bound by Dietzfelbinger and Woelfel [11].

Power-law exponent α	Upper bound (for any pair)	Lower bound	
$0 \leq \alpha < 2$	$O(\log^2 n)$	$\Omega(\log^2 n)$	(for worst pair)
$\alpha = 2$	$O(\log^{4/3} n)$	$\Omega(\log^{4/3} n)$	(for worst pair)
$2 < \alpha < 3$	$O(\log^{\alpha-1} n \cdot \log \log n)$	$\Omega(\log^{\alpha-1} n)$	(for random pair)
$\alpha = 3$	$O(\log^2 n)$	$\Omega(\log^2 n / \log \log n)$	(for random pair)
$\alpha > 3$	$O(\log^2 n)$	$\Omega(\log^2 n)$	(for random pair)

Table 1 Our bounds on the expected routing time of greedy routing, for the small-world model with a power-law degree distribution of exponent α , and bidirectional long-range edges. The upper bounds hold for any source–target pair. The lower bounds hold for a randomly chosen pair when $\alpha > 2$, and for the worst pair when $\alpha \leq 2$. The upper bounds are stated formally in Theorem 1, and the lower bounds in Theorem 2.

α approaches 2. We complement this upper bound with an almost matching lower bound of $\Omega(\log^{\alpha-1} n)$ steps, on the expected routing time for a uniformly random pair. For the case of $\alpha \geq 3$ or $\alpha < 2$, we show that an expected number of $O(\log^2 n)$ steps suffices for any pair. Further, we show a matching lower bound of $\Omega(\log^2 n)$ expected steps for a random pair when $\alpha > 3$, and for the worst pair when $\alpha < 2$; for $\alpha = 3$ we show a lower bound that is weaker by a $\log \log n$ factor. Finally, for the critical value $\alpha = 2$, we show an upper bound of $O(\log^{4/3} n)$ on the expected number of steps for any pair, and a matching lower bound of $\Omega(\log^{4/3} n)$ for the worst pair.

1.2 Related work

Several papers have extended Kleinberg’s work [24] to small-world models based on structures other than the lattice (or the grid graph). In a follow-up work, Kleinberg [25] extended his results to a hierarchical model, in which nodes are the leaves of a complete b -ary tree, and the distance between two nodes is the length of the path between them in the tree. Further, he proposed a model that generalizes both the lattice-based and the hierarchical models, in which the distances between nodes are induced by certain families of node sets. Small-world networks on base structures similar to the grid graph were studied by Martel and Nguyen [31, 32], who computed the diameter of these networks. Duchon, Hanusse, Lehar, and Schabanel [14] considered small-world networks in which the base graph has a “bounded growth rate” property. Fraigniaud [16] studied base graphs of bounded tree-width. Slivkins [37] considered the case in which nodes are embedded in a metric space of bounded doubling dimension. Finally, Abraham and Gavoille [1] studied base graphs that exclude a fixed minor. In all these settings, greedy routing (with respect to distances in the base structure) finishes in a polylogarithmic expected number of steps, for suitable distributions of long-range edges similar to the ℓ -harmonic distribution on the ℓ -dimensional lattice. More recently, in [19] we considered the case in which the base structure is an arbitrary graph, and we showed that routing in sub-polynomial expected time can be

achieved by a slight variant of greedy routing, for long-range edges drawn from an adaptation of the harmonic distribution.

A long line of work has studied *lower bounds* for greedy routing on small-world networks, e.g., the works by Aspnes, Diamadi, and Shah [4], Flammini, Moscardelli, Navarra, and Pérennes [15], Giakkoupis and Hadzilacos [20], and Dietzfelbinger and Woelfel [10, 11]. These bounds are for an extension of Kleinberg’s one-dimensional model, which assumes that the distribution used to choose the number of long-range edges per node and the length of these edges can be arbitrary, but is the same for all nodes. In the most recent of these works [11], it was shown that for any distribution with constant expected number of long-range edges per node, greedy routing needs an expected number of $\Omega(\log^2 n)$ steps for a random source–target pair of nodes. This result does not contradict our results, as it assumes *unidirectional* long-range edges.

Several decentralized routing algorithms, other than the standard greedy algorithm, have been proposed for Kleinberg’s network. Manku, Naor, and Wieder [30] considered a greedy algorithm assuming that each node knows also the neighbors of its neighbors; their results built upon the work of Coppersmith, Gamarnik, and Sviridenko [9]. Lehar and Schabanel [28], Fraigniaud, Gavoille, and Paul [17], Martel and Nguyen [31], and Giakkoupis and Schabanel [21] proposed algorithms that construct shorter paths than greedy routing, but they visit (or “consult”) an additional small number of nearby nodes before they decide the next node in the path. For more related work on decentralized routing (or search) in small-world networks, see the survey by Kleinberg [26]. See also the work by Lattanzi, Panconesi, and Sivakumar [27] for a different model on search in social networks.

Exploiting the degree distribution in the design of decentralized search algorithms for networks with power-law degree distributions has been investigated by Adamic, Lukose, Puniyani, and Huberman [2], Kim, Yoon, Han, and Jeong [22], and Sarshar, Boykin, and Roychowdhury [35]. In these works, the search algorithm has access only to information on the degrees of neighboring nodes, and not to

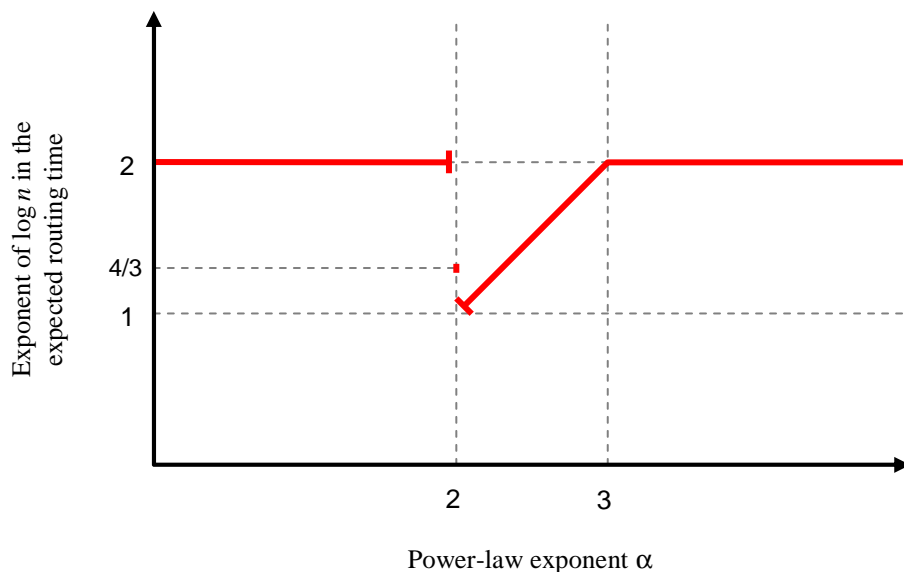


Fig. 1 Summary of our results.

any type of spacial embedding of nodes as in Kleinberg’s model. Simsek and Jensen [36] proposed a heuristic routing algorithm for a variant of Kleinberg’s model [24] similar to ours, in which nodes have widely varying degrees. Their algorithm assumes that nodes know *both* the location and the degree of their neighbors, and their simulation results suggest that it is faster than decentralized algorithms that use only *one* of these two types of information. A main difference between this work and ours is that the greedy routing algorithm we consider does not take into account node degrees. Despite that, we show that it can benefit from power-law degree distributions. Further, in our work we give provable bounds as opposed to Simsek and Jensen who provide an experimental evaluation.

2 Model and results

Below we describe the small-world model and the routing algorithms we consider, then state our main theorems, and finally give some intuition for the results.

2.1 Small-world graph

For integers $\ell, n \geq 1$ and a real $\alpha \geq 0$, we denote by $SW(\ell, n, \alpha)$ a random directed graph generated as follows. We start with an ℓ -dimensional grid graph that wraps around (i.e., a torus), in which edges are bidirectional and every dimension has size n . Thus the total number of nodes is n^ℓ . We will refer to this graph as the *grid*, and call the 2ℓ neighbors of each node its *grid-neighbors*. Further for any two nodes

u, v we define the *grid-distance*, or simply *distance*, between u and v to be their shortest-path distance in the grid, and denote it by $d(u, v)$ or $d_{u,v}$.⁴ This deterministic grid is then augmented by adding random edges as follows. For each node u , we draw independently at random an integer C_u from a power-law distribution with exponent α and mean 2; the details of this distribution are given below. Then, we choose C_u nodes independently at random (with replacement) according to a distribution that assigns to each node $v \neq u$ a probability proportional to $d(u, v)^{-\ell}$; the details are given below. The resulting *multi-set* of nodes is the set of *out-contacts* of u , and we draw an edge from u to each of these out-contacts, allowing parallel edges. These edges are the *long-range* edges of u . We say that node v is an *in-contact* of u , if u is an out-contact of v .

Next we specify the two probability distributions used in the above construction. The power-law distribution from which the number C_u of the out-contacts of u is drawn is defined as follows. If $\alpha > 2$ then C_u is chosen from the set $\{1, 2, \dots\}$ of positive integers such that

$$\Pr(C_u = k) = q_k := \begin{cases} k^{-\alpha}/\nu, & \text{if } k \geq 2; \\ 1 - \sum_{i \geq 2} i^{-\alpha}/\nu, & \text{if } k = 1, \end{cases}$$

⁴ In Kleinberg’s original model, a node has edges to all nodes at distance at most r ; in our model we assume that $r = 1$. Further, in Kleinberg’s model the grid does not wrap around; this assumption, however, is used in many subsequent works, e.g., by Martel and Nguyen [31, 32]. We expect that these two assumptions are not critical for our results.

where ν is the normalizing constant that yields $\mathbf{E}[C_u] = 2$, i.e., $\sum_{k \geq 1} kq_k = 2$. It follows that

$$\nu = \sum_{i \geq 2} (i^{1-\alpha} - i^{-\alpha}).$$

For $\alpha \leq 2$ this definition does not work, as in this case the sum above is unbounded. For this reason we impose a maximum value of k_{\max} on C_u , i.e., C_u is now chosen from the finite set $\{1, \dots, k_{\max}\}$; other than that, the distribution we use is exactly the same as in the case of $\alpha > 2$. We will assume that $k_{\max} = n^\gamma$ for some constant $0 < \gamma \leq \ell$.

We point out that our model assumptions that $C_u \geq 1$ with probability 1 and $\mathbf{E}[C_u] = 2$ were made just to simplify exposition. The results of this paper hold as long as $\Pr(C_u \geq 1) = \Omega(1)$ and $\mathbf{E}[C_u] = \Theta(1)$. Hence, for the case of $\alpha > 2$, we could use instead the more natural distribution that $\Pr(C_u = k) = k^{-\alpha}/\nu$ for any $k \geq 1$ (i.e., including $k = 1$) for a constant ν , and $k = 0$ with the remaining probability. For $\alpha \leq 2$, we cannot assume that $\Pr(C_u = k) = k^{-\alpha}/\nu$ for $k = 1$, as then C_u must be zero with probability $1 - o(1)$ (and thus $\Pr(C_u \geq 1) = o(1)$) in order to have $\mathbf{E}[C_u] = O(1)$.

The distribution from which we choose each of the C_u out-contacts of u is the standard distribution proposed by Kleinberg: Each node $v \neq u$ is picked with probability

$$p_{d_{u,v}} := (d_{u,v})^{-\ell} / \eta,$$

where ℓ is the dimension of the grid and η is the normalizing factor

$$\eta = \sum_{v \neq u} (d_{u,v})^{-\ell} = \Theta(\ln n).$$

The assumption that the out-contacts of a node are chosen *with* replacement is standard and is convenient for the analysis. We expect that the same asymptotic results should hold even if the out-contacts are chosen *without* replacement. The reason is that the fraction of duplicate out-contacts is significant only for very large values of C_u , and our analysis shows that the role of nodes u with such large C_u is negligible, because they are so rare.

Finally, for the number of in-contacts of a node, it can be shown that it follows a distribution that is close to a binomial distribution with constant mean.

2.2 Greedy routing

We consider two versions of greedy routing. The first is the standard greedy protocol: A node u that receives a message for target node $t \neq u$ forwards the message to the node v among its grid-neighbors and out-contacts that is closest to t in the grid, i.e., $d(v, t)$ is minimal. If there are more than one such v , then any of them can be used, but we assume that the choice is deterministic (e.g., we can use the first of those

v in the lexicographic order of their d -dimensional vector of coordinates). We call this algorithm **GreedyUniDir**, as the messages can be sent through an edge only along the direction of the edge. The second algorithm we consider is called **GreedyBiDir**, and ignores the direction of edges: node u forwards the message to the node v among its grid-neighbors, out-contacts, and in-contacts that is closest to t in the grid.

We consider two measures for the performance of the above routing algorithms. The first is the *expected routing time for the worst-case source–target pair*, i.e., we measure the expected routing time for each pair and then take the largest of these expected times. The second measure is the *expected routing time for a source–target pair chosen uniformly at random*. This is equivalent to measuring the expected routing time for each pair and then taking the average. Clearly, the first measure (for the worst pair) is always greater or equal to the second measure (for the random pair). Our upper bounds are for the worst pair, and our lower bounds for a random pair, except for the lower bounds for **GreedyBiDir** when $\alpha \leq 2$, which are for the worst pair.

2.3 Results

Next we state our main theorems. The asymptotic notation is for $n \rightarrow \infty$, and we assume that ℓ and α are not functions of n .

Theorem 1 (Upper bounds) *The expected routing time of GreedyUniDir for the worst pair in $SW(\ell, n, \alpha)$ is $O(\log^2 n)$. For GreedyBiDir, the expected routing time for the worst pair is*

$$\begin{cases} O(\log^2 n), & \text{if } 0 \leq \alpha < 2 \text{ or } \alpha \geq 3; \\ O(\log^{\alpha-1} n \cdot \log \log n), & \text{if } 2 < \alpha < 3; \\ O(\log^{4/3} n), & \text{if } \alpha = 2. \end{cases}$$

Theorem 2 (Lower bounds) *The expected routing time of GreedyUniDir for a random pair in $SW(\ell, n, \alpha)$ is $\Omega(\log^2 n)$. For GreedyBiDir, the expected routing time for a random pair is*

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

and the expected routing time for the worst pair is

$$\begin{cases} \Omega(\log^2 n), & \text{if } 0 \leq \alpha < 2; \\ \Omega(\log^{4/3} n), & \text{if } \alpha = 2. \end{cases}$$

2.4 Intuition

We give now some informal intuition for the results above. For the $O(\log^2 n)$ bound to hold it suffices that each node has just one out-contact. Achieving smaller routing times requires that nodes with $\omega(1)$ out-contacts are encountered sufficiently often. For GreedyUniDir, the probability that the next node in the routing path has k out-contacts equals $q_k = \Pr(C_u = k)$. We show that this probability is not high enough to yield routing times below $\Omega(\log^2 n)$. In the case of GreedyBiDir, nodes with many out-contacts are more likely to be found than in GreedyUniDir, as they can be reached through their long-range edges. The probability that a given node has an in-contact with k out-contacts is proportional to the fraction of long-range edges in the network starting from nodes with k out-contacts, and is roughly kq_k . The actual probability that the next node in the routing path has k out-contacts is smaller than that, and the main reason is that only nodes that are closer to the target than the node who has the message are relevant to routing. This reduces the above probability of kq_k by roughly a factor of $\ln d / \ln n$, where d is the current distance to the target.

For the case of $2 < \alpha < 3$, the above reasoning gives that the probability of the next node in the path to have $\Theta(\ln n)$ out-contacts (more concretely, say, between $\ln n$ and $2 \ln n$) is roughly $\ln n \cdot (\ln n \cdot q_{\ln n}) \cdot (\ln d / \ln n) \approx \ln d / \ln^{\alpha-1} n$; and from such a node, the distance to the target halves in the next step with probability $\Omega(1)$. Based on that we show the $O(\ln^{\alpha-1} n \cdot \log \log n)$ bound. To prove the lower bound we further observe that the contribution of nodes with either $o(\ln n)$ or $\omega(\ln n)$ out-contacts is negligible compared to that of nodes with $\Theta(\ln n)$ out-contacts. We point out that for $\alpha \rightarrow 2$ the above upper bound approaches $O(\ln n)$, which is the expected routing time in Kleinberg's model for $k = \Theta(\ln n)$ out-contacts per node.

In the case of $\alpha = 2$, roughly equal fractions of long-range edges start from nodes with $\Theta(2^i)$ out-contacts for each i . When the distance d to the target is sufficiently large, the distance decreases to $d^{1-\varepsilon}$ in a step with probability roughly $\ln^2 d / \ln^2 n$, and the nodes that contribute the most to that reduction are those with a number of out-contacts roughly between $d^{\varepsilon'}$ and d , for constants ε and ε' . It follows that the distance to the target decreases quickly when d is large, but when the distance drops below roughly $2^{(\ln n)^{1/3}}$, we have essentially no speedup. Then it takes

$$\Theta(\log n \cdot \log 2^{(\log n)^{1/3}}) = \Theta(\log^{4/3} n)$$

steps to reach the target from that distance.

In the case of $\alpha > 3$, nodes with $\omega(1)$ out-contacts are very rare, and only a $o(1)$ fraction of the long-range edges starts from those nodes. It follows that the $O(\log^2 n)$ bound is tight for this case. On the other hand, when $\alpha < 2$, a large fraction of the long-range edges starts from nodes that have

a very high number of out-contacts. However, nodes with $\omega(1)$ out-contacts are so rare that, with significant probability, no such node exists within distance $\Theta(n^\varepsilon)$ from the target, for a sufficiently small constant ε . It follows that $\Theta(\log n \cdot \log n^\varepsilon) = \Theta(\log^2 n)$ steps are needed to cover that distance.

The lower bounds we provide for GreedyBiDir for $\alpha \leq 2$ are shown for the worst pair only rather than a random pair. In [18] we showed that the $\Omega(\log^{4/3} n)$ bound for $\alpha = 2$ holds also for a random pair, but here we consider just the worst pair as the proof for the random pair is much more involved. For $\alpha < 2$, we do not know whether the $O(\log^2 n)$ bounds holds for a random pair.

3 Preliminaries

In this section we first introduce some notation, and then prove a collection of results on the distribution of edges to be used later in the analysis.

For two sets of nodes U_1 and U_2 we write $U_1 \rightarrow U_2$ to denote that some node from U_2 is an out-contact of a node from U_1 . If U_1 or U_2 is a singleton set, we may write its element instead of the set in this notation, e.g., we will write $u \rightarrow v$ to denote $\{u\} \rightarrow \{v\}$. A ball $B_u(r)$ centered at node u with radius r is the set of nodes v with $d(u, v) \leq r$. A sphere $S_u(r)$ is the set $B_u(r) \setminus B_u(r-1)$.

We denote by $\mathcal{K}(\ell, n, k)$ Kleinberg's model obtained in the same way as $\mathcal{SW}(\ell, n, \alpha)$ except that each node u has k out-contacts, instead of C_u . It follows that $\mathcal{SW}(\ell, n, \alpha)$ can be obtained from $\mathcal{K}(\ell, n, 1)$, by adding $C_u - 1$ additional long-range edges from each node u .

Next we provide some bounds on the probability that a given node in $\mathcal{K}(\ell, n, 1)$ has out-contacts or in-contacts from a given set of nodes.

Observation 1 *Let u, t be two distinct nodes in $\mathcal{K}(\ell, n, 1)$, and U be a nonempty set of nodes.*

- (a) $\Pr(u \rightarrow B_t(r)) = O\left(\ln\left(\frac{d_{u,t}}{d_{u,t}-r}\right) / \ln n\right)$, for $1 \leq r \leq d_{u,t} - 1$.
- (b) $\Pr(u \rightarrow B_t(r)) = \begin{cases} \Omega\left(\left(\frac{r}{d_{u,t}}\right)^\ell / \ln n\right), & \text{if } 1 \leq r \leq d_{u,t}/2; \\ \Omega\left(\ln\left(\frac{d_{u,t}}{d_{u,t}-r}\right) / \ln n\right), & \text{if } d_{u,t}/2 \leq r \leq d_{u,t} - 1. \end{cases}$
- (c) $\Pr(u \rightarrow B_t(d_{u,t}/2)) = \Theta(1 / \ln n)$.
- (d) $\Pr(u \rightarrow B_t(d_{u,t} - 1)) = \Theta(\ln(d_{u,t}) / \ln n)$.
- (e) $\Pr(u \rightarrow S_t(r)) = O\left(\frac{1}{\ln n \cdot (d_{u,t}-r)}\right)$, for $0 \leq r \leq d_{u,t} - 1$.
- (f) $\Pr(u \rightarrow U) = o(1)$, if $|U| = n^{o(1)}$.
- (g) $\Pr(U \rightarrow u) \in [\Pr(u \rightarrow U)/2, \Pr(u \rightarrow U)]$.
- (h) $\Pr(U \not\rightarrow u) \geq 1/4$.

Proof We give first some useful facts. The maximum distance between any two nodes is $d_{\max} = \ell \cdot \lfloor n/2 \rfloor$. The number of nodes at distance r from node t is

$$|S_t(r)| = \begin{cases} \Theta(r^{\ell-1}), & \text{if } 1 \leq r \leq n/2; \\ O(r^{\ell-1}), & \text{if } n/2 \leq r \leq d_{\max}. \end{cases} \quad (1)$$

For the case of $1 \leq r \leq n/2$, the result follows from [31, Fact 19(i)]. For the case of $n/2 \leq r \leq d_{\max}$, we just apply the result for the previous case to the larger grid in which every dimension has size d_{\max} (instead of n), and we observe that $|S_t(r)|$ in the original grid is bounded by the corresponding quantity in the larger grid. We can use the above bound for $|S_t(r)|$ to bound the size of ball $B_t(r)$,

$$|B_t(r)| = \sum_{i=0}^r |S_t(i)| = \Theta(r^\ell). \quad (2)$$

(a): For each node $v \in B_t(r)$, its distance $d_{u,v}$ from u satisfies $d_{u,t} - r \leq d_{u,v} \leq d_{u,t} + r$. It follows that $B_t(r) \subseteq \bigcup_{i=d_{u,t}-r}^{d_{u,t}+r} S_u(i)$, and thus

$$\begin{aligned} \Pr(u \rightarrow B_t(r)) &\leq \sum_{i=d_{u,t}-r}^{d_{u,t}+r} \Pr(u \rightarrow S_u(i)) = \sum_{i=d_{u,t}-r}^{d_{u,t}+r} |S_u(i)| \cdot p_i \\ &\stackrel{(1)}{=} O\left(\sum_{i=d_{u,t}-r}^{d_{u,t}+r} \frac{i^{\ell-1}}{\ln n \cdot i^\ell}\right) \\ &= O\left(\ln\left(\frac{d_{u,t}+r}{d_{u,t}-r}\right) / \ln n\right). \end{aligned}$$

To complete the proof we argue that $\ln\left(\frac{d_{u,t}+r}{d_{u,t}-r}\right) \leq 2\ln\left(\frac{d_{u,t}}{d_{u,t}-r}\right)$: We must show $\frac{d_{u,t}+r}{d_{u,t}-r} \leq \left(\frac{d_{u,t}}{d_{u,t}-r}\right)^2$ which is equivalent to $(d_{u,t}+r) \cdot (d_{u,t}-r) \leq (d_{u,t})^2$, and this is equivalent to $(d_{u,t})^2 - r^2 \leq (d_{u,t})^2$, which is true.

(b): First we consider the case of $1 \leq r \leq (1-\varepsilon) \cdot d_{u,t}$, for an arbitrary small constant $\varepsilon > 0$. For each node $v \in B_t(r)$ we have $d_{u,v} \leq d_{u,t} + r \leq 2d_{u,t}$, and thus

$$\begin{aligned} \Pr(u \rightarrow B_t(r)) &\geq |B_t(r)| \cdot p_{2d_{u,t}} \\ &\stackrel{(2)}{=} \Omega\left(\frac{r^\ell}{\ln n \cdot (2d_{u,t})^\ell}\right) = \Omega\left(\frac{r^\ell}{\ln n \cdot (d_{u,t})^\ell}\right). \end{aligned}$$

This proves the case of $r \leq d_{u,t}/2$, and also the case of $d_{u,t}/2 \leq r \leq (1-\varepsilon) \cdot d_{u,t}$, because in the latter case both quantities $\ln\left(\frac{d_{u,t}}{d_{u,t}-r}\right)$ and $\left(\frac{r}{d_{u,t}}\right)^\ell$ are $\Theta(1)$.

Next we consider the case of $(1-\varepsilon) \cdot d_{u,t} < r \leq d_{u,t} - 1$, for a sufficiently small constant $\varepsilon > 0$; in particular, we will need that $\varepsilon \leq 1/(4\ell^2)$. Further, we assume that $d_{u,t} > 1/\varepsilon$, otherwise the above range for r is empty. We represent each node by its ‘grid coordinates’, i.e., an ℓ -vector from $\{-\lfloor n/2 \rfloor, \dots, \lfloor n/2 \rfloor\}^\ell$; we assume that $u = (0, 0, \dots, 0)$ and $t = (x_1, \dots, x_\ell)$. Since $\sum_{i=1}^\ell |x_i| = d_{u,t}$, it follows that for

some index i it holds $|x_i| \geq d_{u,t}/\ell$. We assume w.l.o.g. that $x_1 \geq d_{u,t}/\ell$. Let $t' = (x_1, 0, \dots, 0)$ and $r' = x_1 - \lambda$, where $\lambda = d_{u,t} - r$. (Observe that $r' > 0$, because $\lambda < \varepsilon d_{u,t}$ as $r > (1-\varepsilon) \cdot d_{u,t}$, and $\varepsilon d_{u,t} \leq d_{u,t}/(4\ell^2) < d_{u,t}/\ell \leq x_1$.) If we consider a shortest path from u to t in the grid, which goes from u to t' and from there to t , it is easy to see that ball $B_{t'}(r')$ is a subset of $B_t(r)$: for each node $v \in B_{t'}(r')$ we have

$$d(t, v) \leq d(t, t') + d(t', v) \leq (d_{u,t} - x_1) + r' = r,$$

and thus $B_{t'}(r') \subseteq B_t(r)$. Therefore,

$$\Pr(u \rightarrow B_t(r)) \geq \Pr(u \rightarrow B_{t'}(r')).$$

We will now lower bound the probability on the right side. Let L_i , for $\lambda \leq i \leq x_1$, be the set of nodes $v = (y_1, \dots, y_\ell)$ from $B_{t'}(r')$ with $y_1 = i$. The size of L_i is then the number of possible ways to fix y_2, \dots, y_ℓ such that $d(v, t') \leq r'$, which is equivalent to

$$|y_2| + \dots + |y_\ell| \leq r' - (x_1 - i) = i - \lambda. \quad (3)$$

Thus if $\ell > 1$, by counting only non-negative combinations we obtain,

$$|L_i| \geq \sum_{s=0}^{i-\lambda} \binom{s+\ell-2}{\ell-2} = 1 + \sum_{s=1}^{i-\lambda} \Omega(s^{\ell-2}) = \Omega\left((i-\lambda)^{\ell-1}\right).$$

The above lower bound holds also when $\ell = 1$, as in this case $|L_i| = 1$ or 2 . Further, for each $v = (i, y_2, \dots, y_\ell) \in L_i$ we have $d(v, u) \leq i + |y_2| + \dots + |y_\ell| \leq 2i$, by (3). It follows

$$\begin{aligned} \Pr(u \rightarrow B_{t'}(r')) &\geq \sum_{i=\lambda}^{x_1} |L_i| \cdot p_{2i} \\ &= \Omega\left(\sum_{i=\lambda}^{x_1} \frac{(i-\lambda)^{\ell-1}}{\ln n \cdot (2i)^\ell}\right) \\ &= \Omega\left(\sum_{i=\lambda}^{x_1} \frac{(1-\lambda/i)^{\ell-1}}{\ln n \cdot i}\right) \\ &= \Omega\left(\sum_{i=2\lambda}^{x_1} \frac{(1/2)^{\ell-1}}{\ln n \cdot i}\right) \\ &= \Omega\left(\ln\left(\frac{x_1}{2\lambda}\right) / \ln n\right). \end{aligned}$$

To complete the proof we argue that $\ln\left(\frac{x_1}{2\lambda}\right) \geq \ln\left(\frac{d_{u,t}}{d_{u,t}-r}\right)/2$: We must show that $\left(\frac{x_1}{2\lambda}\right)^2 \geq \frac{d_{u,t}}{d_{u,t}-r}$, and since $x_1 \geq d_{u,t}/\ell$ and $\lambda = d_{u,t} - r$, it suffices to show that $\left(\frac{d_{u,t}/\ell}{2(d_{u,t}-r)}\right)^2 \geq \frac{d_{u,t}}{d_{u,t}-r}$, which is equivalent to $\frac{d_{u,t}-r}{d_{u,t}} \leq \frac{1}{4\ell^2}$. The last inequality holds because the left side is at most ε and we have assumed that $\varepsilon \leq 1/(4\ell^2)$.

(c), (d): Both results follow from (a) and (b).

(e): This result is a special case of [31, Fact 5].

(f): For a fixed set size $|U|$, the probability $\Pr(u \rightarrow U)$ is maximized when U consists of the $|U|$ nodes closest to u . From (1), the number of nodes at distance i from u is $|\mathcal{S}_u(i)| \leq c \cdot i^{\ell-1}$ for some constant c . It follows

$$\begin{aligned} \Pr(u \rightarrow U) &\leq \sum_{i=1}^{\lceil (|U|/c)^{1/(\ell-1)} \rceil} (c \cdot i^{\ell-1} \cdot p_i) \\ &= O\left(\sum_{i=1}^{\lceil (|U|/c)^{1/(\ell-1)} \rceil} \frac{1}{\ln n \cdot i}\right) \\ &= O\left(\frac{\ln\left(\lceil (|U|/c)^{1/(\ell-1)} \rceil\right)}{\ln n}\right) = O\left(\frac{\ln(|U|)}{\ln n}\right) \\ &= o(1), \end{aligned}$$

as $|U| = n^{o(1)}$.

(g): The upper bound follows from the union bound,

$$\begin{aligned} \Pr(U \rightarrow u) &\leq \sum_{v \in U} \Pr(v \rightarrow u) \\ &= \sum_{v \in U} \Pr(u \rightarrow v) = \Pr(u \rightarrow U). \end{aligned}$$

For the lower bound we have

$$\begin{aligned} \Pr(U \rightarrow u) &= 1 - \prod_{v \in U} (1 - \Pr(v \rightarrow u)) \\ &\geq 1 - \prod_{v \in U} e^{-\Pr(v \rightarrow u)} = 1 - e^{-\sum_{v \in U} \Pr(v \rightarrow u)} \\ &= 1 - e^{-\sum_{v \in U} \Pr(u \rightarrow v)} = 1 - e^{-\Pr(u \rightarrow U)} \\ &\geq \Pr(u \rightarrow U)/2, \end{aligned}$$

where for the last relation we used the fact that $e^{-x} \leq 1 - x/2$ for $0 \leq x \leq 1$.

(h): For each $v \in U$ we have $\Pr(v \not\rightarrow u) = 1 - \Pr(v \rightarrow u) \geq 4^{-\Pr(v \rightarrow u)}$, because of the facts that $1 - x \geq 4^{-x}$ for $0 \leq x \leq 1/2$, and $\Pr(v \rightarrow u) < 1/2$. It follows that

$$\begin{aligned} \Pr(U \not\rightarrow u) &= \prod_{v \in U} \Pr(v \not\rightarrow u) \geq \prod_{v \in U} 4^{-\Pr(v \rightarrow u)} \\ &= 4^{-\sum_{v \in U} \Pr(v \rightarrow u)} = 4^{-\sum_{v \in U} \Pr(u \rightarrow v)} \\ &= 4^{-\Pr(u \rightarrow U)} \geq 4^{-1}. \end{aligned}$$

This completes the proof of Observation 1. \square

The next claim gives bounds on the probability that a given node u in $\mathcal{SW}(\ell, n, \alpha)$ has out-contacts in some set U , conditionally on the event that u has (or has not) long-range edges to nodes in another given set. The bounds are in terms of simpler probability expressions that do not involve conditioning on u having long-range edges to other sets. This type of conditioning will be very common throughout our analysis of GreedyBiDir, as, e.g., nodes that are closer to the target than the node holding the message have no edges to previous nodes in the routing path.

Observation 2 Let u be a node in $\mathcal{SW}(\ell, n, \alpha)$, and U, W, W' be disjoint sets of nodes not containing u . Sets U and W' are nonempty, but W may be empty.

- (a) $\Pr(u \rightarrow U) \leq \Pr(u \rightarrow U \mid C_u = 2) \leq 2\Pr(u \rightarrow U \mid C_u = 1)$.
- (b) For any event \mathcal{E} , $\Pr(\mathcal{E} \mid u \not\rightarrow W) \leq (1 + o(1)) \cdot \Pr(\mathcal{E})$, if $|W| = n^{o(1)}$; in particular, $\Pr(u \rightarrow U \mid u \not\rightarrow W) \leq (1 + o(1)) \cdot \Pr(u \rightarrow U \mid C_u = 2)$, if $|W| = n^{o(1)}$.
- (c) $\Pr(u \rightarrow U \mid u \not\rightarrow W) \leq \Pr(u \rightarrow U \mid C_u = 3)$, if $|W| = n^{o(1)}$ and $|U| = n^{o(1)}$.
- (d) $\Pr(u \rightarrow U \mid u \not\rightarrow W) \geq \Pr(u \rightarrow U \mid C_u = 1)$.
- (e) $\Pr(u \rightarrow U \mid C_u = k, u \rightarrow W') \geq \Pr(u \rightarrow U \mid C_u = k - 1)$.
- (f) $\Pr(u \rightarrow U \mid C_u = k, u \not\rightarrow W) \leq \frac{\Pr(u \rightarrow U \mid C_u = k)}{\Pr(u \not\rightarrow W \mid C_u = 1)}$.

Proof (a): Let $a := \Pr(u \not\rightarrow U \mid C_u = 1)$; then $\Pr(u \not\rightarrow U \mid C_u = k) = a^k$. We have

$$\begin{aligned} \Pr(u \not\rightarrow U) &= \sum_k \Pr(u \not\rightarrow U \mid C_u = k) \cdot \Pr(C_u = k) \\ &= \sum_k a^k \cdot \Pr(C_u = k) = \mathbf{E}[a^{C_u}]. \end{aligned}$$

Since a^x is a convex function, it follows from Jensen's Inequality that

$$\mathbf{E}[a^{C_u}] \geq a^{\mathbf{E}[C_u]} = a^2 = \Pr(u \not\rightarrow U \mid C_u = 2).$$

Therefore, $\Pr(u \not\rightarrow U) \geq \Pr(u \not\rightarrow U \mid C_u = 2)$, and thus $\Pr(u \rightarrow U) \leq \Pr(u \rightarrow U \mid C_u = 2)$, as desired. Also, from the union bound it follows $\Pr(u \rightarrow U \mid C_u = 2) \leq 2\Pr(u \rightarrow U \mid C_u = 1)$.

(b): We have

$$\Pr(\mathcal{E} \mid u \not\rightarrow W) = \frac{\Pr(\mathcal{E} \wedge u \not\rightarrow W)}{\Pr(u \not\rightarrow W)} \leq \frac{\Pr(\mathcal{E})}{\Pr(u \not\rightarrow W)},$$

and $\Pr(u \not\rightarrow W) = 1 - \Pr(u \rightarrow W) \stackrel{(a)}{\geq} 1 - 2\Pr(u \rightarrow W \mid C_u = 1) \stackrel{\text{Obs.1(f)}}{=} 1 - o(1)$. Thus,

$$\Pr(\mathcal{E} \mid u \not\rightarrow W) \leq \Pr(\mathcal{E}) / (1 - o(1)) = (1 + o(1)) \cdot \Pr(\mathcal{E}).$$

The second part of (b) follows from the first and (a).

(c): Let $a := \Pr(u \rightarrow U \mid C_u = 1)$, and note that $a = o(1)$ from Observation 1(f). We have

$$\begin{aligned} \Pr(u \rightarrow U \mid u \not\rightarrow W) &\stackrel{(b)}{\leq} (1 + o(1)) \cdot \Pr(u \rightarrow U \mid C_u = 2) \\ &\leq 2a + o(a). \end{aligned}$$

Further,

$$\begin{aligned} \Pr(u \rightarrow U \mid C_u = 3) &= 1 - (1 - a)^3 \\ &= 1 - (1 - 3a + o(a)) = 3a - o(a), \end{aligned}$$

where the second-to-last relation holds because $a = o(1)$. From the two expressions above it follows $\Pr(u \rightarrow U \mid u \not\rightarrow W) \leq \Pr(u \rightarrow U \mid C_u = 3)$.

(d): Since u has at least one out-contact, we have

$$\begin{aligned} \Pr(u \rightarrow U \mid u \not\rightarrow W) &\geq \Pr(u \rightarrow U \mid u \not\rightarrow W, C_u = 1) \\ &= \frac{\Pr(u \rightarrow U \wedge u \not\rightarrow W \mid C_u = 1)}{\Pr(u \not\rightarrow W \mid C_u = 1)} \\ &= \frac{\Pr(u \rightarrow U \mid C_u = 1)}{\Pr(u \not\rightarrow W \mid C_u = 1)} \\ &\geq \frac{\Pr(u \rightarrow U \mid C_u = 1)}{1}. \end{aligned}$$

(e): Given that $C_u = k$ and $u \rightarrow W'$, we can assume that the out-contacts of u are chosen as follows. We start by choosing (from the right probability distribution) the smallest j such that the j -th out-contact of u is the first one that belongs to W' . Then we choose the first $j - 1$ out-contacts conditionally on the event that $u \not\rightarrow W'$, and the last $k - j$ out-contacts without any conditioning. We note that the first $j - 1$ out-contacts are more likely to belong to U than if they were chosen unconditionally. The claim then follows.

(f): From Bayes' Rule we have

$$\begin{aligned} \Pr(u \rightarrow U \mid C_u = k, u \not\rightarrow W) \\ &= \frac{\Pr(u \not\rightarrow W \mid C_u = k, u \rightarrow U)}{\Pr(u \not\rightarrow W \mid C_u = k)} \cdot \Pr(u \rightarrow U \mid C_u = k). \end{aligned}$$

Also from (e) it follows $\Pr(u \not\rightarrow W \mid C_u = k, u \rightarrow U) \leq \Pr(u \not\rightarrow W \mid C_u = k - 1)$, and thus,

$$\begin{aligned} \frac{\Pr(u \not\rightarrow W \mid C_u = k, u \rightarrow U)}{\Pr(u \not\rightarrow W \mid C_u = k)} \\ &\leq \frac{\Pr(u \not\rightarrow W \mid C_u = k - 1)}{\Pr(u \not\rightarrow W \mid C_u = k)} \\ &= \frac{(\Pr(u \not\rightarrow W \mid C_u = 1))^{k-1}}{(\Pr(u \not\rightarrow W \mid C_u = 1))^k} = \frac{1}{\Pr(u \not\rightarrow W \mid C_u = 1)}. \end{aligned}$$

Combining the above yields the claim.

This completes the proof of Observation 2. \square

4 Proof of the upper bounds

We prove now the upper bounds of Theorem 1. In Section 4.1, we show an $O(\log^2 n)$ bound that applies to both algorithms for all exponents α . In fact, we prove a more general lemma which states that $O((\log d + \log \beta^{-1}) \cdot \log n)$ steps suffice to reach the target from a node at distance d with probability $1 - \beta$. The main claim we use is that the distance to the target halves with probability $\Omega(1/\ln n)$ in each step. This claim holds even if nodes have only one out-contact each.

In Sections 4.2 and 4.3 we prove stronger bounds for GreedyBiDir for the cases of $2 < \alpha < 3$ and $\alpha = 2$, respectively. For the case of $2 < \alpha < 3$, the key claim is that the distance to the target halves with probability $\Omega(\log d / \log^{\alpha-1} n)$ within $O(1)$ steps, where d is the current distance. (We note that this bound is larger than $\Omega(1/\ln n)$ for d sufficiently large.) The proof of that claim lower bounds the probability of the event that the next node in the path: (1) is an in-contact of the previous node, (2) has $\Theta(\log n)$ out-contacts, and (3) has at least one out-contact within distance $d/2$ from the target. For the case of $\alpha = 2$, we have a similar claim saying that the distance to the target decreases to $d^{1-\varepsilon}$ for some constant $\varepsilon > 0$, with probability $\Omega(\log^2 d / \log^2 n)$ in the next $O(1)$ steps. The proof lower bounds the probability of the event that the next node in the path (1) is an in-contact of the previous node, (2) has between d^{ε_1} and d^{ε_2} out-contacts, for some constants $\varepsilon_1 < \varepsilon_2 < 1 - \varepsilon$, and (3) has some out-contact within distance d^ε from the target. In both cases, $2 < \alpha < 3$ and $\alpha = 2$, we use the lemma from Section 4.1 to bound the number of remaining steps when the distance to the target becomes sufficiently small.

4.1 A general upper bound

In this section we prove a bound of $O((\log d + \log \beta^{-1}) \cdot \log n)$ steps for routing between two given nodes at distance d , that holds with probability $1 - \beta$. This bound applies to both routing algorithms and all exponents α .

Lemma 1 *For either protocol and any $\alpha \geq 0$, the routing time for a pair of nodes s, t is $O((\log d_{s,t} + \log \beta^{-1}) \cdot \log n)$ with probability $1 - \beta$, for any $0 < \beta < 1$ (which may depend on n and $d_{s,t}$). Further, for GreedyBiDir we have that if s is an intermediate node in the routing path from some node s' to t , then the same probability bound holds for the time until the message reaches t from s , conditionally on the path up to s .*

This result implies the $O(\log^2 n)$ bound of Theorem 1 for GreedyUniDir, and for GreedyBiDir when $0 \leq \alpha < 2$ or $\alpha \geq 3$: by choosing $\beta = 1/d_{s,t}$, Lemma 1 gives that the routing time is $O(\log d_{s,t} \cdot \log n)$ with probability at least $1 - 1/d_{s,t}$, and thus the expected routing time is at most

$$\begin{aligned} (1 - 1/d_{s,t}) \cdot O(\log d_{s,t} \cdot \log n) + (1/d_{s,t}) \cdot d_{s,t} \\ = O(\log d_{s,t} \cdot \log n) = O(\log^2 n). \end{aligned}$$

Proof of Lemma 1 The proof structure follows that of [24, Theorem 2]. Denote by d_i the distance between the node that has the message after i steps and the target t . We will prove the following claim below.

Claim 1 $\Pr[d_{i+1} \leq d_i/2 \mid d_i] = \Omega(1/\ln n)$.

We use Claim 1 to bound the expectation of d_{i+1} . We have

$$\begin{aligned} \mathbf{E}[d_{i+1} \mid d_i] &\leq \Pr(d_{i+1} \leq d_i/2 \mid d_i) \cdot d_i/2 \\ &\quad + (1 - \Pr(d_{i+1} \leq d_i/2 \mid d_i)) \cdot d_i \\ &= (1 - \Pr(d_{i+1} \leq d_i/2 \mid d_i)/2) \cdot d_i \\ &\leq (1 - c/\ln n) \cdot d_i, \end{aligned}$$

for some constant $c > 0$, by Claim 1. Taking now the unconditional expectation we obtain that $\mathbf{E}[d_{i+1}] \leq (1 - c/\ln n) \cdot \mathbf{E}[d_i]$. Applying this inequality repeatedly yields $\mathbf{E}[d_i] \leq (1 - c/\ln n)^i \cdot d_0$, and thus for $t^* = (\ln d_0 + \ln \beta^{-1}) \cdot \ln n / c$,

$$\begin{aligned} \mathbf{E}[d_{t^*}] &\leq (1 - c/\ln n)^{(\ln d_0 + \ln \beta^{-1}) \cdot \ln n / c} d_0 \\ &\leq e^{-(\ln d_0 + \ln \beta^{-1})} d_0 = \beta. \end{aligned}$$

Markov's Inequality then gives $\Pr(d_{t^*} \geq 1) \leq \mathbf{E}[d_{t^*}] / 1 \leq \beta$. Hence, the probability that the target is reached in at most t^* steps is $\Pr(d_{t^*} = 0) = 1 - \Pr(d_{t^*} \geq 1) \geq 1 - \beta$.

It remains to prove Claim 1. Fix the i -th node u_i in the routing path, and thus its distance d_i from target t . For the case of GreedyUniDir, the out-contacts of u_i do not depend on the routing path up to node u_i . Since u_i has at least one out-contact, the probability that $u_i \rightarrow B_t(d_i/2)$ is lower bounded by the probability of the same event in $\mathcal{K}(\ell, n, 1)$:

$$\begin{aligned} \Pr(u_i \rightarrow B_t(d_i/2)) &\geq \Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow B_t(d_i/2)) \\ &\stackrel{\text{Obs.1(c)}}{=} \Theta(1/\ln n), \end{aligned}$$

and thus $\Pr(d_{i+1} \leq d_i/2) = \Omega(1/\ln n)$.

For the case of GreedyBiDir, we rely on the in-contacts of u_i to prove the claim rather than its out-contacts. Fix the path $u_0 \dots u_i$ up to node u_i . For each node $v \in B_t(d_i/2)$, we have $v \not\rightarrow \{u_0, \dots, u_{i-1}\}$ and thus from Observation 2(d), $\Pr(v \rightarrow u_i) \geq \Pr_{\mathcal{K}(\ell, n, 1)}(v \rightarrow u_i)$. It follows

$$\begin{aligned} \Pr(B_t(d_i/2) \rightarrow u_i) &\geq \Pr_{\mathcal{K}(\ell, n, 1)}(B_t(d_i/2) \rightarrow u_i) \\ &\stackrel{\text{Obs.1(g)}}{\geq} \Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow B_t(d_i/2))/2 \\ &\stackrel{\text{Obs.1(c)}}{=} \Theta(1/\ln n), \end{aligned}$$

and thus $\Pr(d_{i+1} \leq d_i/2) = \Omega(1/\ln n)$. This completes the proof of Claim 1 and Lemma 1. \square

4.2 Upper bound for GreedyBiDir for $2 < \alpha < 3$

In this section we prove an upper bound of $O(\log^{\alpha-1} n \cdot \log \log n)$ on the expected routing time of GreedyBiDir, for any pair of nodes s, t in $\mathcal{SW}(\ell, n, \alpha)$, when $2 < \alpha < 3$.

For this range of values for α , the normalizing factor v of distribution $q_k = \Pr(C_u = k)$ for choosing the number of out-contacts of a node is $v = \sum_{i \geq 2} (i^{1-\alpha} - i^{-\alpha}) = \Theta(1)$. Thus $q_k = \Theta(k^{-\alpha})$ for $k \geq 2$. Further, for $k = 1$ we have

$q_1 = \Theta(1)$, because since the expectation is 2, it follows that $q_1 \geq q_3 = \Theta(1)$. Therefore,

$$q_k = \Theta(k^{-\alpha}), \quad \text{for all } k \geq 1.$$

We observe that as soon as the message reaches a node within distance $d^* = 2^{\log^{\alpha-2} n}$ from target t , Lemma 1 yields an upper bound of $O(\log d^* \cdot \log n) = O(\log^{\alpha-1} n)$ on the remaining steps until t is reached with probability $1 - 1/d^*$. From this, it follows that the *expected* number of remaining steps is also $O(\log^{\alpha-1} n)$. Thus, we just need to bound the expected number of steps from source s until some node v with $d(v, t) \leq d^*$ is reached.

Let d_i denote the distance between the node that has the message after i steps and target t . We will prove the following key claim which lower bounds the probability that the distance to the target is halved in the next three steps.

Claim 2 For any $i \leq \log^3 n$, if $d_i \geq d^*/2$ then $\Pr(d_{i+3} \leq d_i/2 \mid d_i \dots d_0) = \Omega(\log d_i / \log^{\alpha-1} n)$.

Note that this bound requires that at most $\log^3 n$ steps have been taken.⁵ Because of that, we will bound first an auxiliary quantity we define now. Let T' be the minimum between $\log^3 n$ and the earliest step i for which $d_i \leq d^*$, i.e., $T' = \min\{\log^3 n, \min\{i: d_i \leq d^*\}\}$. We bound $\mathbf{E}[T']$ as follows. For $0 \leq k \leq \lfloor \log(n/d^*) \rfloor$, let $i_k = \min\{\log^3 n, \min\{i: d_i \leq n/2^k\}\}$. Then $T' \leq i_{\lfloor \log(n/d^*) \rfloor}$. From Claim 2 it follows that

$$\begin{aligned} \mathbf{E}[i_{k+1} - i_k] &= 3/\Omega(\log(n/2^{k+1})/\log^{\alpha-1} n) \\ &= O\left(\frac{\log^{\alpha-1} n}{\log n - (k+1)}\right). \end{aligned}$$

Applying the above result repeatedly yields

$$\mathbf{E}[i_k] = O\left(\sum_{1 \leq j \leq k} \frac{\log^{\alpha-1} n}{\log n - j}\right) = O(\log^{\alpha-1} n \cdot \ln \log n),$$

and thus

$$\mathbf{E}[T'] = O(\log^{\alpha-1} n \cdot \ln \log n).$$

If now $T = \min\{i: d_i \leq d^*\}$, then T and T' are identical except for when $T > \log^3 n$. Hence, $T' \geq T \cdot \mathbf{1}_{T \leq \log^3 n}$, and we have

$$\begin{aligned} \mathbf{E}[T] &= \mathbf{E}[T \cdot \mathbf{1}_{T \leq \log^3 n}] + \mathbf{E}[T \cdot \mathbf{1}_{T > \log^3 n}] \\ &\leq \mathbf{E}[T'] + n \cdot \Pr(T > \log^3 n). \end{aligned}$$

From Lemma 1 for $\beta = 1/n$, it follows $T = O(\log^2 n) < \log^3 n$ with probability $1 - 1/n$, and thus $\Pr(T > \log^3 n) < 1/n$. From this and the inequality above, $\mathbf{E}[T] \leq O(\log^{\alpha-1} n \cdot \ln \log n) + n \cdot (1/n) = O(\log^{\alpha-1} n \cdot \ln \log n)$.

It remains to prove Claim 2.

⁵ In fact, it holds for any i that is at most a polylogarithmic function of n , but for our purposes it suffices to assume that $i \leq \log^3 n$.

Proof of Claim 2 We fix the path $u_0 \dots u_i$ in the first i steps. Further, we reveal whether or not u_i is an out-contact of u_{i-1} . If u_i is not an out-contact of u_{i-1} and the two nodes are not grid-neighbors, then u_i must be an in-contact of u_{i-1} , and this changes significantly the a priori probability distribution of C_{u_i} . As we explain later, we avoid handling this case directly; instead we reduce this case to the case we describe next by consider one extra step.

We study now the case in which either u_i is an out-contact of u_{i-1} , or the two nodes are grid-neighbors. We describe four events such that if all of them occur, then it follows that $d_{i+2} \leq d_i/2$. Next we show that the probability of the intersection of these events is lower bounded by $\Omega(\log d_i / \log^{\alpha-1} n)$, and thus, the same bound holds for $\Pr(d_{i+2} \leq d_i/2)$.

Roughly speaking, the events we consider say that the next node u_{i+1} is an in-contact of u_i with approximately $\log n$ out-contacts, which is at distance between $d_i/2$ and $d_i - \sqrt{d_i}$ from target t , and has an out-contact within distance $d_i/2$ from t . Formally, let $R := B_t(d_i - \sqrt{d_i}) \setminus B_t(d_i/2)$ and $D := \{\ln n, \dots, 2 \ln n\}$. We define the following events.

- \mathcal{E}_1 : there is some node $v \in R$ with $C_v \in D$ for which we have $v \rightarrow u_i$; let z be the node that is closest to t among the nodes v that satisfy this event;
- \mathcal{E}_2 : $z \rightarrow B_t(d_i/2)$;
- \mathcal{E}_3 : for all $v \in R$ with $C_v \notin D$ we have $v \not\rightarrow u_i$;
- \mathcal{E}_4 : $u_i \not\rightarrow R$.

It is immediate that the intersection of the above four events implies $d_{i+2} \leq d_i/2$: If u_i has some in- or out-contact in $B_t(d_i/2)$ then we have $d_{i+1} \leq d_i/2$. Otherwise, from events \mathcal{E}_1 , \mathcal{E}_3 , and \mathcal{E}_4 , it follows that u_i will forward the message to node z . And by \mathcal{E}_2 , node z will subsequently forward the message to a node in $B_t(d_i/2)$.

We now bound the probability of $\mathcal{E}_1 \wedge \mathcal{E}_2 \wedge \mathcal{E}_3 \wedge \mathcal{E}_4$. Since \mathcal{E}_4 is independent of the other events,

$$\begin{aligned} \Pr(\mathcal{E}_1 \wedge \mathcal{E}_2 \wedge \mathcal{E}_3 \wedge \mathcal{E}_4) \\ = \Pr(\mathcal{E}_4) \cdot \Pr(\mathcal{E}_3) \cdot \Pr(\mathcal{E}_1 \mid \mathcal{E}_3) \cdot \Pr(\mathcal{E}_2 \mid \mathcal{E}_1, \mathcal{E}_3). \end{aligned}$$

We chose the order of conditioning to simplify the calculations below. We show next that each of the probabilities $\Pr(\mathcal{E}_4)$, $\Pr(\mathcal{E}_3)$, and $\Pr(\mathcal{E}_2 \mid \mathcal{E}_1, \mathcal{E}_3)$ is lower bounded by a positive constant, and $\Pr(\mathcal{E}_1 \mid \mathcal{E}_3) = \Omega(\log d_i / \log^{\alpha-1} n)$. It follows that $\Pr(\mathcal{E}_1 \wedge \mathcal{E}_2 \wedge \mathcal{E}_3 \wedge \mathcal{E}_4) = \Omega(\log d_i / \log^{\alpha-1} n)$, and hence the same lower bound applies to $\Pr(d_{i+2} \leq d_i/2)$.

Proof that $\Pr(\mathcal{E}_4) = \Omega(1)$: Recall that we have fixed the path $u_0 \dots u_i$ so far. For node u_i we have $u_i \not\rightarrow \{u_0, \dots, u_{i-2}\}$.

We remove this conditioning using Observation 2(b),

$$\begin{aligned} \Pr(\mathcal{E}_4) &= 1 - \Pr(u_i \rightarrow R) \\ &\stackrel{\text{Obs.2(b)}}{\geq} 1 - (1 + o(1)) \cdot \Pr_{\mathcal{K}(\ell, n, 2)}(u_i \rightarrow R) \\ &= \Pr_{\mathcal{K}(\ell, n, 2)}(u_i \not\rightarrow R) - o(1) \\ &= (\Pr_{\mathcal{K}(\ell, n, 1)}(u_i \not\rightarrow R))^2 - o(1). \end{aligned}$$

We will argue now that $\Pr_{\mathcal{K}(\ell, n, 1)}(u_i \not\rightarrow R) = \Omega(1)$, and thus $\Pr(\mathcal{E}_4) = \Omega(1)$, as desired. Let t' be a node at distance $\sqrt{d_i}/2$ from u_i . The ball $B := B_{t'}(\sqrt{d_i}/2 - 1)$ does not overlap with R , hence

$$\Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow B) + \Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow R) \leq 1. \quad (4)$$

We have $\Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow B) = \Omega(\ln d_i / \ln n)$ from Observation 1(d), and

$$\begin{aligned} \Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow R) &\leq \Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow B_t(d_i - 1)) \\ &= O(\ln d_i / \ln n), \end{aligned}$$

from the same observation. Thus, $\Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow B) \geq c \cdot \Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow R)$, for some constant $c > 0$. From this and Eq. (4) above, it follows $\Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow R) \leq 1/(1+c)$, and thus $\Pr_{\mathcal{K}(\ell, n, 1)}(u_i \not\rightarrow R) \geq c/(1+c) = \Omega(1)$.

Proof that $\Pr(\mathcal{E}_3) = \Omega(1)$: We will show a stronger result instead, that $\Pr(U \not\rightarrow u_i) = \Omega(1)$ for $U := B_t(d_{u_i, t} - 1)$. For each node $v \in U$ we have $v \not\rightarrow \{u_0, \dots, u_{i-1}\}$, and from Observation 2(c),

$$\Pr(v \rightarrow u_i) \leq \Pr_{\mathcal{K}(\ell, n, 3)}(v \rightarrow u_i).$$

From this it follows that $\Pr(U \rightarrow u_i) \leq \Pr_{\mathcal{K}(\ell, n, 3)}(U \rightarrow u_i)$, and thus

$$\begin{aligned} \Pr(U \not\rightarrow u_i) &\geq \Pr_{\mathcal{K}(\ell, n, 3)}(U \not\rightarrow u_i) \\ &= (\Pr_{\mathcal{K}(\ell, n, 1)}(U \not\rightarrow u_i))^3 \stackrel{\text{Obs.1(h)}}{=} \Omega(1). \end{aligned}$$

Proof that $\Pr(\mathcal{E}_2 \mid \mathcal{E}_1, \mathcal{E}_3) = \Omega(1)$: Fix z and the number $C_z \geq \ln n$ of its out-contacts. We must lower bound by $\Omega(1)$ the probability that $z \rightarrow B_t(d_i/2)$, given that $z \rightarrow u_i$ and $z \not\rightarrow \{u_0, \dots, u_{i-1}\}$. First we observe that we can drop the assumption that $z \not\rightarrow \{u_0, \dots, u_{i-1}\}$, since it can only increase the probability for each of the C_z long-range edges of u to point to a given node $v \notin \{u_0, \dots, u_{i-1}\}$, and thus to $B_t(d_i/2)$. Then, from Observation 2(e) we can also drop assumption $z \rightarrow u_i$, and assume instead that z has one out-contact fewer, i.e., at least $\ln n - 1$ out-contacts. From Observation 1(c), each of these out-contacts belongs to $B_t(d_i/2) \supseteq B_t(d_{z, t}/2)$ with probability at least $\Theta(1/\log n)$. Hence the probability that at least one of them belongs to $B_t(d_i/2)$ is lower bounded by

$$1 - (1 - \Omega(1/\ln n))^{\ln n - 1} \geq 1 - e^{-\Omega(1)} = \Omega(1).$$

Proof that $\Pr(\mathcal{E}_1 \mid \mathcal{E}_3) = \Omega(\log d_i / \log^{\alpha-1} n)$: Roughly speaking, we show that for each $v \in R$, the probability of the event that $v \rightarrow u_i \wedge C_v \in D$ equals $p_{d(v,u_i)}$ times $\Omega(1/\log^{\alpha-2} n)$, and combine this with the result that in $\mathcal{K}(\ell, n, 1)$ we have $R \rightarrow u_i$ with probability $\Omega(\ln d_i / \ln n)$.

We start by observing that conditioning on event \mathcal{E}_3 (i.e., u_i has no in-contacts $v \in R$ with $C_v \notin D$) increases the probability of event \mathcal{E}_1 (i.e., u_i has some in-contact $v \in R$ with $C_v \in D$), as it increases the probability that $C_v \in D$ for a given $v \in R$. Thus, we have $\Pr(\mathcal{E}_1 \mid \mathcal{E}_3) \geq \Pr(\mathcal{E}_1)$, and it suffices to lower bound $\Pr(\mathcal{E}_1)$. Let $H = \{u_0, \dots, u_{i-1}\}$ be the set of the first i nodes in the routing path. For each $v \in R$ we have $v \not\rightarrow H$, and from Bayes' Rule,

$$\begin{aligned} & \Pr(v \rightarrow u_i \wedge C_v \in D \mid v \not\rightarrow H) \\ &= \frac{\Pr(v \not\rightarrow H \mid v \rightarrow u_i, C_v \in D)}{\Pr(v \not\rightarrow H)} \\ & \quad \cdot \Pr(v \rightarrow u_i \mid C_v \in D) \cdot \Pr(C_v \in D) \\ & \geq \frac{\Pr(v \not\rightarrow H \mid v \rightarrow u_i, C_v = 2 \log n)}{1} \\ & \quad \cdot \Pr(v \rightarrow u_i \mid C_v = \log n) \cdot (|D| \cdot q_{2 \log n}), \end{aligned}$$

where for the last relation we used that $D = \{\ln n, \dots, 2 \ln n\}$, and for each probability in the numerator we chose the value from D that minimizes this probability. Next we bound the two probabilities in the last two lines above. We have

$$\begin{aligned} \Pr(v \rightarrow H \mid v \rightarrow u_i, C_v = 2 \log n) & \leq \Pr(v \rightarrow H \mid C_v = 2 \log n) \\ & \leq 2 \log n \cdot |H| \cdot p_{\sqrt{d_i}} \\ & \leq 2 \log n \cdot |H| \cdot p_{\sqrt{d^*/2}} \\ & = o(1), \end{aligned}$$

where for the inequality in the second line we used that the distance of v from any node in H is at least $\sqrt{d_i}$, by R 's definition. Also

$$\begin{aligned} \Pr(v \not\rightarrow u_i \mid C_v = \log n) &= (1 - p_{d(v,u_i)})^{\log n} \\ & \leq e^{-\log n \cdot p_{d(v,u_i)}} \\ & \leq 1 - (1 - o(1)) \cdot \log n \cdot p_{d(v,u_i)}, \end{aligned}$$

where for the last equality we used the facts that $e^{-x} \leq 1 - x + x^2/2$ and $\log n \cdot p_{d(v,u_i)} = o(1)$. Combining the above yields

$$\begin{aligned} & \Pr(v \rightarrow u_i \wedge C_v \in D \mid v \not\rightarrow H) \\ & \geq (1 - o(1)) \cdot ((1 - o(1)) \cdot \log n \cdot p_{d(v,u_i)}) \cdot (|D| \cdot q_{2 \log n}) \\ & = \rho \cdot p_{d(v,u_i)}, \end{aligned}$$

where $\rho := (1 - o(1))^2 \cdot \log n \cdot |D| \cdot q_{2 \log n} = \Omega(1/\log^{\alpha-2} n)$. From this it follows that the probability of the event $\neg \mathcal{E}_1$ that

u_i has no in-contacts $v \in R$ with $C_v \in D$ is

$$\begin{aligned} \Pr(\neg \mathcal{E}_1) & \leq \prod_{v \in R} (1 - \rho \cdot p_{d(v,u_i)}) \\ & \leq \prod_{v \in R} e^{-\rho \cdot p_{d(v,u_i)}} = e^{-\rho \sum_{v \in R} p_{d(v,u_i)}}. \end{aligned}$$

We compare this with the probability in $\mathcal{K}(\ell, n, 1)$ that u_i has no in-contacts $v \in R$,

$$\begin{aligned} \Pr_{\mathcal{K}(\ell, n, 1)}(R \not\rightarrow u_i) &= \prod_{v \in R} (1 - p_{d(v,u_i)}) \\ & \geq \prod_{v \in R} e^{-p_{d(v,u_i)}/2} = e^{-(1/2) \cdot \sum_{v \in R} p_{d(v,u_i)}}, \end{aligned}$$

where for the inequality we used the fact that $1 - x \geq e^{-x-x^2}$ when $0 \leq x \leq 1/2$, and $p_{d(v,u_i)} = o(1)$. Further, we have

$$\begin{aligned} \Pr_{\mathcal{K}(\ell, n, 1)}(R \rightarrow u_i) & \stackrel{\text{Obs.1(g)}}{\geq} \Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow R)/2 \\ & = \Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow B_t(d_i - \sqrt{d_i}))/2 \\ & \quad - \Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow B_t(d_i/2))/2 \\ & \stackrel{\text{Obs.1(b)\&(c)}}{=} \Omega(\ln d_i / \ln n) - \Theta(1/\ln n) \\ & = \Omega(\ln d_i / \ln n). \end{aligned}$$

Combining the last three results above yields

$$\begin{aligned} \Pr(\neg \mathcal{E}_1) & \leq \left(\Pr_{\mathcal{K}(\ell, n, 1)}(R \not\rightarrow u_i) \right)^{2\rho} \\ & \leq (1 - \Omega(\ln d_i / \ln n))^{2\rho} \leq e^{-\Omega(\rho \ln d_i / \ln n)}, \end{aligned}$$

and substituting the value of ρ gives

$$\Pr(\neg \mathcal{E}_1) \leq e^{-\Omega(\ln d_i / \ln^{\alpha-1} n)} = 1 - \Omega(\ln d_i / \ln^{\alpha-1} n).$$

Thus, $\Pr(\mathcal{E}_1) = \Omega(\ln d_i / \ln^{\alpha-1} n)$, as desired.

We have now completed the proof for the case in which either u_i is an out-contact of u_{i-1} , or the two nodes are grid-neighbors. We have shown for this case that

$$\Pr(d_{i+2} \leq d_i/2) = \Omega(\log d_i / \log^{\alpha-1} n).$$

We consider now the case in which u_i is neither an out-contact nor a grid-neighbor of u_{i-1} , and thus it is an in-contact of u_{i-1} . Fix the out-contacts of u_i , and let w be the out-contact or grid-neighbor of u_i that is closest to t . If $d_{w,t} \leq d_i/2$ then $d_{i+1} \leq d_i/2$ as desired; so, we assume that $d_{w,t} > d_i/2$. The probability that u_i has no in-contacts in $B_t(d_{w,t} - 1)$ is $\Pr(B_t(d_{w,t} - 1) \not\rightarrow u_i) = \Omega(1)$; the proof of this result is the same as that for $\Pr(\mathcal{E}_3) = \Theta(1)$, described earlier. Given now that $B_t(d_{w,t} - 1) \not\rightarrow u_i$, the next node u_{i+1} in the path is node w , and we can apply the results of the previous case to obtain that $\Pr(d_{i+3} \leq d_{i+1}/2) = \Omega(\log d_{i+1} / \log^{\alpha-1} n) = \Omega(\log d_i / \log^{\alpha-1} n)$, as $d_{i+1} = d_{w,t} > d_i/2$. It follows that $\Pr(d_{i+3} \leq d_i/2) \geq \Theta(1) \cdot \Omega(\log d_i / \log^{\alpha-1} n)$. This completes the proof of this case, and the proof of Claim 2. \square

4.3 Upper bound for GreedyBiDir for $\alpha = 2$

In this section we prove an upper bound of $O(\log^{4/3} n)$ on the expected routing time of GreedyBiDir, for any pair of nodes s, t in $\mathcal{SW}(\ell, n, 2)$. The analysis is similar to that for $\mathcal{SW}(\ell, n, \alpha)$ with $2 < \alpha < 3$, in Section 4.2. Arguments that are the same in both proofs are not repeated here.

For $\alpha = 2$, the normalizing factor ν of the probability distribution $q_k = \Pr(C_u = k)$ is $\nu = \sum_{2 \leq i \leq k_{\max}} (i^{-1} - i^{-2}) = \Theta(\ln k_{\max}) = \Theta(\ln n)$, and thus

$$q_k = \begin{cases} \Theta\left(\frac{1}{k^2 \ln n}\right), & \text{if } k \geq 2; \\ 1 - \Theta\left(\frac{1}{\ln n}\right), & \text{if } k = 1. \end{cases}$$

As in Section 4.2, we use Lemma 1 to argue that the expected number of remaining steps when the message has reached a node within distance $d^* = 2^{\log^{1/3} n}$ of t , is bounded by $O(\log d^* \cdot \log n) = O(\log^{4/3} n)$. Next we bound the expected number of steps from source s until some node v with $d(v, t) \leq d^*$ is reached.

Led d_i be the distance between the node that has the message after i steps and target t . The next result is the analogue to Claim 2; we give its proof later.

Claim 3 *Let ε be an arbitrary constant such that $1 - 1/(3\ell) < \varepsilon < 1$. If $i \leq \log^3 n$ and $d_i \geq (d^*)^\varepsilon$, then $\Pr(d_{i+3} \leq d_i^\varepsilon \mid d_i \dots d_0) = \Omega(\log^2 d_i / \log^2 n)$.*

As before we let $T' = \min\{\log^3 n, \min\{i: d_i \leq d^*\}\}$, and we bound $\mathbf{E}[T']$ as follows. For each $0 \leq k \leq k^* := \lceil \log_{1/\varepsilon} \left(\frac{\log n}{\log d^*}\right) \rceil$, let $i_k = \min\{\log^3 n, \min\{i: d_i \leq n^{\varepsilon^k}\}\}$. Hence, $T' \leq i_{k^*}$. From Claim 3 then it follows that

$$\mathbf{E}[i_{k+1} - i_k] = 3/\Omega\left(\log^2(n^{\varepsilon^{k+1}}) / \log^2 n\right) = O(\varepsilon^{-2(k+1)}).$$

Applying the above result repeatedly yields

$$\mathbf{E}[i_k] = O\left(\sum_{1 \leq j \leq k} \varepsilon^{-2j}\right) = O(\varepsilon^{-2k}),$$

and thus

$$\begin{aligned} \mathbf{E}[T'] &\leq \mathbf{E}[i_{k^*}] = O\left(\varepsilon^{-2 \log_{1/\varepsilon} \left(\frac{\log n}{\log d^*}\right)}\right) \\ &= O\left(\left(\frac{\log n}{\log d^*}\right)^2\right) \\ &= O\left(\left(\frac{\log n}{\log^{1/3} n}\right)^2\right) = O(\log^{4/3} n). \end{aligned}$$

Letting now $T = \min\{i: d_i \leq d^*\}$, we argue that $\mathbf{E}[T] \leq \mathbf{E}[T'] + O(1)$ in exactly the same way as in Section 4.2. Hence, we obtain $\mathbf{E}[T] = O(\log^{4/3} n)$.

It remains to prove Claim 3.

Proof of Claim 3 The proof is very similar to that of Claim 2. As before, we fix the path $u_0 \dots u_i$ in the first i steps, and we consider first the case in which either u_i is an out-contact of u_{i-1} , or the two nodes are grid-neighbors. We consider the same four events we used in the proof of Claim 2, but for different sets R and D . We let $R := B_t(d_i - \sqrt{d_i}) \setminus B_t(d_i^\varepsilon)$ and $D := \{d_i^{\varepsilon_1}, \dots, d_i^{\varepsilon_2}\}$, where ε_1 and ε_2 are arbitrary constants with $1/3 < \varepsilon_1 < \varepsilon_2 < 1/2$. Then,

\mathcal{E}_1 : there is some node $v \in R$ with $C_v \in D$ for which we have $v \rightarrow u_i$; let z be the node that is closest to t among the nodes v that satisfy this event;

\mathcal{E}_2 : $z \rightarrow B_t(d_i^\varepsilon)$;

\mathcal{E}_3 : for all $v \in R$ with $C_v \notin D$ we have $v \not\rightarrow u_i$;

\mathcal{E}_4 : $u_i \not\rightarrow R$.

By the same reasoning as in the proof of Claim 2, we have that the probability of $d_{i+2} \leq d_i^\varepsilon$ is lower bounded by the probability of the intersection of the four events above, and

$$\begin{aligned} \Pr(\mathcal{E}_1 \wedge \mathcal{E}_2 \wedge \mathcal{E}_3 \wedge \mathcal{E}_4) \\ = \Pr(\mathcal{E}_4) \cdot \Pr(\mathcal{E}_3) \cdot \Pr(\mathcal{E}_1 \mid \mathcal{E}_3) \cdot \Pr(\mathcal{E}_2 \mid \mathcal{E}_1, \mathcal{E}_3). \end{aligned}$$

We show that $\Pr(\mathcal{E}_4)$, $\Pr(\mathcal{E}_3)$, and $\Pr(\mathcal{E}_2 \mid \mathcal{E}_1, \mathcal{E}_3)$ are $\Omega(1)$, as before, and that $\Pr(\mathcal{E}_1 \mid \mathcal{E}_3) = \Omega(\log^2 d_i / \log^2 n)$. It follows that $\Pr(\mathcal{E}_1 \wedge \mathcal{E}_2 \wedge \mathcal{E}_3 \wedge \mathcal{E}_4) = \Omega(\log^2 d_i / \log^2 n)$, and hence the same bound applies to $\Pr(d_{i+2} \leq d_i^\varepsilon)$.

Proof that $\Pr(\mathcal{E}_4) = \Omega(1)$ and $\Pr(\mathcal{E}_3) = \Omega(1)$: The derivations of these two bounds described in the proof for Claim 2 apply also in the current setting without changes.

Proof that $\Pr(\mathcal{E}_2 \mid \mathcal{E}_1, \mathcal{E}_3) = \Omega(1)$: Fix z and $C_z \geq d_i^{\varepsilon_1}$. From the same argument as in the proof for Claim 2, it suffices to lower bound by $\Omega(1)$ the probability that $z \rightarrow B_t(d_i^\varepsilon)$ assuming that z has at least $C_z - 1 \geq d_i^{\varepsilon_1} - 1$ out-contacts (without the assumptions that $z \rightarrow u_i$ and $z \not\rightarrow \{u_0, \dots, u_{i-1}\}$). From Observation 1(b) it follows that each of these out-contacts belongs to $B_t(d_i^\varepsilon)$ with probability

$$\begin{aligned} \Omega\left(\left(\frac{d_i^\varepsilon}{d_{z,t}}\right)^\ell / \ln n\right) &\geq \Omega\left(\left(\frac{d_i^\varepsilon}{d_i}\right)^\ell / \ln n\right) \\ &= \Omega\left(d_i^{-(1-\varepsilon)\ell} / \ln n\right). \end{aligned}$$

Hence the probability that at least one of the out-contacts belongs to $B_t(d_i^\varepsilon)$ is lower bounded by

$$\begin{aligned} 1 - \left(1 - \Omega\left(d_i^{-(1-\varepsilon)\ell} / \ln n\right)\right)^{d_i^{\varepsilon_1} - 1} &\geq 1 - e^{-\Omega\left(d_i^{\varepsilon_1} d_i^{-(1-\varepsilon)\ell} / \ln n\right)} \\ &= 1 - o(1), \end{aligned}$$

where the last relation holds because $\varepsilon_1 > 1/3 > (1-\varepsilon)\ell$.

Proof that $\Pr(\mathcal{E}_1 \mid \mathcal{E}_3) = \Omega(\log^2 d_i / \log^2 n)$: As explained in the proof for Claim 2, we have $\Pr(\mathcal{E}_1 \mid \mathcal{E}_3) \geq$

$\Pr(\mathcal{E}_1)$ and thus it suffices to lower bound $\Pr(\mathcal{E}_1)$. For each node $v \in R$ and for $H = \{u_0, \dots, u_{i-1}\}$ we have

$$\begin{aligned} & \Pr(v \rightarrow u_i \wedge C_v \in D \mid v \not\rightarrow H) \\ &= \sum_{k \in D} \Pr(v \rightarrow u_i \wedge C_v = k \mid v \not\rightarrow H) \\ &= \sum_{k \in D} \left(\frac{\Pr(v \not\rightarrow H \mid v \rightarrow u_i, C_v = k)}{\Pr(v \not\rightarrow H)} \right. \\ & \quad \left. \cdot \Pr(v \rightarrow u_i \mid C_v = k) \cdot \Pr(C_v = k) \right). \end{aligned}$$

For $k \in D$,

$$\begin{aligned} & \Pr(v \rightarrow H \mid v \rightarrow u_i, C_v = k) \\ & \leq \Pr(v \rightarrow H \mid C_v = k) \leq |H| \cdot kp_{\sqrt{d_i}} \\ & = O\left(\log^3 n \cdot d_i^{\varepsilon_2} / \left(\ln n \cdot \sqrt{d_i}^\ell\right)\right) = o(1), \end{aligned}$$

where the last relation holds because $\varepsilon_2 < 1/2 \leq \ell/2$. Also

$$\begin{aligned} & \Pr(v \not\rightarrow u_i \mid C_v = k) = (1 - p_{d(v, u_i)})^k \\ & \leq e^{-kp_{d(v, u_i)}} = 1 - (1 - o(1)) \cdot kp_{d(v, u_i)}, \end{aligned}$$

where for the last equality we used the facts that $e^{-x} \leq 1 - x + x^2/2$ and $kp_{d(v, u_i)} \leq kp_{\sqrt{d_i}} = o(1)$. Finally, we have $\Pr(C_v = k) = q_k = \Theta(1/(\ln n \cdot k^2))$. Combining the above yields

$$\begin{aligned} & \Pr(v \rightarrow u_i \wedge C_v = k \mid v \not\rightarrow H) \\ & \geq \sum_{k \in D} \frac{(1 - o(1)) \cdot (1 - o(1)) \cdot kp_{d(v, u_i)} \cdot q_k}{1} \\ & = \rho \cdot p_{d(v, u_i)}, \end{aligned}$$

where

$$\begin{aligned} \rho & := (1 - o(1)) \sum_{k \in D} kq_k = \Theta\left(\sum_{k \in D} \frac{1}{\ln n \cdot k}\right) \\ & = \Theta\left(\frac{\ln d_i^{\varepsilon_2} - \ln d_i^{\varepsilon_1}}{\ln n}\right) = \Theta\left(\frac{\ln d_i}{\ln n}\right). \end{aligned}$$

In the proof of Claim 2, we showed that $\Pr(\neg \mathcal{E}_1) \leq e^{-\Omega(\rho \ln d_i / \ln n)}$. Substituting ρ 's value gives

$$\Pr(\neg \mathcal{E}_1) \leq e^{-\Omega(\ln^2(d_i) / \ln^2 n)} = 1 - \Omega(\ln^2(d_i) / \ln^2 n).$$

Thus, $\Pr(\mathcal{E}_1) = \Omega(\ln^2(d_i) / \ln^2 n)$.

This completes the proof for the case in which either u_i is an out-contact of u_{i-1} , or the two nodes are grid-neighbors. The proof of the complementary case, in which u_i is neither an out-contact nor a grid-neighbor of u_{i-1} , follows from the previous case as explained in the proof of Claim 2.

We have thus finished the proof of Claim 3. \square

5 Proof of the lower bounds

We prove now the lower bounds of Theorem 2. For GreedyUniDir we consider the logarithm of the distances d_0, d_1, \dots , where d_i is the distance to the target after i steps. We show that the expected decrease of this quantity in a step is $\mu = O(1/\log n)$, and then apply Wald's Theorem to obtain a lower bound of $\Omega(\log d_0 / \mu)$ on the expected routing time.

For GreedyBiDir we consider a slightly difference sequence of distances, in order to decrease the dependency on the past. Roughly speaking, if the node u_i that has the message after i rounds is *not* an in-contact of u_{i-1} then, as before, d_i is the distance between u_i and the target. But if u_i is an in-contact of u_{i-1} then d_i is instead the minimum distance of u_i and its out-contacts from the target. Similar to GreedyUniDir, we bound the expected decrease in the logarithm of this distance in a step, and then apply Wald's Theorem. The bound we establish on this decrease is in terms of the distribution of d_i , i.e., $\Pr(d_i \leq r \mid d_{i-1} = \delta)$. To bound this probability we must bound the probability of the event that u_i is an in-contact of u_{i-1} and has an out-contact at distance at most r from the target. Bounding this ‘‘two-hop’’ probability is a main component of our analysis.

In Section 5.1 we describe our proof technique. In Section 5.2 we show the lower bound for GreedyUniDir. In Section 5.3 we define the sequence of distances to be used in the analysis of GreedyBiDir, and then derive the lower bounds for GreedyBiDir in Sections 5.4–5.8.

5.1 The method

In this section we describe the steps that we follow to prove the lower bounds. These steps, or slight variations of them, are used for all the cases of the theorem, except if stated otherwise.

Let s be the source and t the target nodes, and let u_i be the node that has the message after the first i steps of routing.

Step 1 We describe a sequence d_0, d_1, \dots of distances to target t , such that $d_0 = d(s, t)$, and $d_i \leq d(u_i, t)$ for each $i \geq 1$. Hence, the routing time for s, t is lower bounded by $\min\{i : d_i = 0\}$. For the case of GreedyUniDir we define d_i simply as $d_i = d(u_i, t)$. For GreedyBiDir the definition of d_i is a bit more involved and is given in Section 5.3.

Step 2 We bound the distribution of d_i . Specifically, we compute a function $f(r, \delta)$ such that for any $\delta > \ln n$ and $1 \leq r \leq \delta - 1$ we have

$$\Pr(d_i \leq r \mid d_{i-1} = \delta) \leq f(r, \delta).$$

Step 3 We bound the expected decrease in a step, of the logarithm of distance d_i . Let

$$l_i = \begin{cases} \ln d_i, & \text{if } d_i \geq 1; \\ 0, & \text{if } d_i = 0. \end{cases}$$

Further, for $i \geq 1$, let ∂l_i be the decrease of this quantity in step i , i.e.,

$$\partial l_i = l_{i-1} - l_i = \begin{cases} \ln(d_{i-1}/d_i), & \text{if } d_i \geq 1; \\ \ln(d_{i-1}), & \text{if } d_{i-1} > d_i = 0. \end{cases}$$

We compute some $\mu > 0$ such that if $d_{i-1} > \ln n$ then

$$\mathbf{E}[\partial l_i \mid d_{i-1}] \leq \mu.$$

We use the next claim, which bounds $\mathbf{E}[\partial l_i \mid d_{i-1}]$ in terms of function f from Step 2.

Claim 4 *If $\delta > \ln n$ then $\mathbf{E}[\partial l_i \mid d_{i-1} = \delta] \leq \sum_{1 \leq r \leq \delta-1} f(r, \delta)/r$.*

Proof Suppose $d_{i-1} = \delta$. We have

$$\begin{aligned} \mathbf{E}[\partial l_i] &= \sum_{1 \leq r \leq \delta-1} \Pr(d_i = r) \cdot \ln(\delta/r) + \Pr(d_i = 0) \cdot \ln(\delta) \\ &= \sum_{1 \leq r \leq \delta-1} (\Pr(d_i \leq r) - \Pr(d_i \leq r-1)) \cdot \ln(\delta/r) \\ &\quad + \Pr(d_i = 0) \cdot \ln(\delta) \\ &= \sum_{1 \leq r \leq \delta-1} \Pr(d_i \leq r) \cdot \ln(\delta/r) \\ &\quad - \sum_{1 \leq r \leq \delta-1} \Pr(d_i \leq r-1) \cdot \ln(\delta/r) \\ &\quad + \Pr(d_i = 0) \cdot \ln(\delta) \\ &= \sum_{1 \leq r \leq \delta-1} \Pr(d_i \leq r) \cdot \ln(\delta/r) \\ &\quad - \sum_{2 \leq r \leq \delta-1} \Pr(d_i \leq r-1) \cdot \ln(\delta/r) \\ &= \sum_{1 \leq r \leq \delta-1} \Pr(d_i \leq r) \cdot \ln(\delta/r) \\ &\quad - \sum_{1 \leq r \leq \delta-2} \Pr(d_i \leq r) \cdot \ln(\delta/(r+1)) \\ &= \sum_{1 \leq r \leq \delta-1} \Pr(d_i \leq r) \cdot \ln(\delta/r) \\ &\quad - \sum_{1 \leq r \leq \delta-1} \Pr(d_i \leq r) \cdot \ln(\delta/(r+1)) \\ &= \sum_{1 \leq r \leq \delta-1} \Pr(d_i \leq r) \cdot \ln((r+1)/r) \\ &\leq \sum_{1 \leq r \leq \delta-1} \Pr(d_i \leq r) \cdot (1/r). \quad \square \end{aligned}$$

Step 4 We lower bound the expected number of steps until we have $d_i \leq \ln n$, by using the bound $\mathbf{E}[\partial l_i \mid d_{i-1}] \leq \mu$ obtained in Step 3 and Wald's Theorem. Let $T_{s,t} = \min\{i: d_i \leq \ln n\}$, and consider the sum $\sum_{1 \leq i \leq T_{s,t}} \partial l_i = l_0 - l_{T_{s,t}}$. By definition, $d_{T_{s,t}} \leq \ln n$. Thus $l_{T_{s,t}} \leq \ln \ln n$, and

$$\sum_{1 \leq i \leq T_{s,t}} \partial l_i \geq \ln(d_0/\ln n).$$

Further, since $\mathbf{E}[\partial l_i \mid i \leq T_{s,t}] = \mathbf{E}[\mathbf{E}[\partial l_i \mid d_{i-1}] \mid i \leq T_{s,t}] \leq \mu$ from Step 3, Wald's Theorem gives

$$\mathbf{E} \left[\sum_{1 \leq i \leq T_{s,t}} \partial l_i \right] \leq \mathbf{E}[T_{s,t}] \cdot \mu.$$

From the two inequalities above it follows

$$\mathbf{E}[T_{s,t}] \geq \ln(d_0/\ln n)/\mu. \quad (5)$$

Step 5 We observe that for a uniformly random pair s, t , it holds that $d(s, t) = n^{\Omega(1)}$ with probability $1 - o(1)$. Then from the result of Step 4 it follows that the expected routing time T for a random pair s, t is

$$\mathbf{E}[T] = \Omega(\ln n)/\mu. \quad (6)$$

This step is not used for the case of GreedyBiDir when $\alpha \leq 2$, as in this case we bound the expected routing time for the worst-case pair only.

5.2 Lower bound for GreedyUniDir

We prove a lower bound of $\Omega(\log^2 n)$ on the expected routing time of GreedyUniDir for a random pair s, t in $\mathcal{SW}(\ell, n, \alpha)$. This bound holds for all $\alpha \geq 0$.

We follow the steps listed in Section 5.1. For $\delta > \ln n$ and $1 \leq r \leq \delta - 1$, we define

$$f(r, \delta) = \begin{cases} c \cdot \ln\left(\frac{\delta}{\delta-r}\right) / \ln n, & \text{if } 1 \leq r \leq \delta - 2; \\ 1, & \text{if } r = \delta - 1, \end{cases}$$

for a constant $c > 0$. We argue now that if c is sufficiently large, then $\Pr(d_i \leq r \mid d_{i-1} = \delta) \leq f(r, \delta)$ as required (cf. Step 2): Fix node u_{i-1} and assume that $d_{i-1} = \delta$. Suppose also that $r \neq \delta - 1$ (because $f(\delta - 1, \delta) = 1$, and thus the desired inequality holds for $r = \delta - 1$). We have

$$\begin{aligned} \Pr(d_i \leq r) &= \Pr(u_i \rightarrow B_i(r)) \\ &\stackrel{\text{Obs.2(a)}}{\leq} 2\Pr_{\mathcal{K}(\ell, n, 1)}(u_i \rightarrow B_i(r)) \\ &\stackrel{\text{Obs.1(a)}}{=} O\left(\ln\left(\frac{\delta}{\delta-r}\right) / \ln n\right). \end{aligned} \quad (7)$$

Therefore, $\Pr(d_i \leq r) \leq f(r, \delta)$ for a large enough constant c .

Next we use Claim 4 to bound $\mathbf{E}[\partial l_i \mid d_{i-1} = \delta]$ (cf. Step 3).

$$\begin{aligned} \mathbf{E}[\partial l_i \mid d_{i-1} = \delta] &\leq \sum_{1 \leq r \leq \delta-1} f(r, \delta)/r \\ &= \sum_{1 \leq r \leq \delta-2} \frac{c}{r} \cdot \ln\left(\frac{\delta}{\delta-r}\right) / \ln n + \frac{1}{\delta-1}. \end{aligned}$$

To bound the last sum we observe that

$$\begin{aligned} \sum_{1 \leq r \leq \delta/2} \frac{1}{r} \cdot \ln\left(\frac{\delta}{\delta-r}\right) &= \sum_{1 \leq r \leq \delta/2} \frac{1}{r} \cdot \ln\left(1 + \frac{r}{\delta-r}\right) \\ &\leq \sum_{1 \leq r \leq \delta/2} \frac{1}{r} \cdot \frac{r}{\delta-r} \\ &\leq \sum_{1 \leq r \leq \delta/2} \frac{1}{\delta - \delta/2} = O(1); \end{aligned} \quad (8)$$

$$\begin{aligned} \sum_{\delta/2 < r \leq \delta-2} \frac{1}{r} \cdot \ln\left(\frac{\delta}{\delta-r}\right) &\leq \frac{1}{\delta/2} \cdot \sum_{\delta/2 < r \leq \delta-2} \ln\left(\frac{\delta}{\delta-r}\right) \\ &\leq \frac{1}{\delta/2} \cdot \ln\left(\frac{\delta^{\lceil \delta/2 \rceil}}{\lceil \delta/2 \rceil!}\right) \\ &\leq \frac{1}{\delta/2} \cdot \ln\left(\frac{\delta^{\lceil \delta/2 \rceil}}{(\delta/2e)^{\lceil \delta/2 \rceil}}\right) \\ &\quad \text{(by Stirling's Approximation)} \\ &\leq \frac{1}{\delta/2} \cdot O(\delta) = O(1). \end{aligned} \quad (9)$$

It follows that $\mathbf{E}[\partial l_i \mid d_{i-1} = \delta] = O(c/\ln n) + 1/(\delta - 1) = O(1/\ln n)$.

We have thus shown that if $d_{i-1} > \ln n$ then $\mathbf{E}[\partial l_i \mid d_{i-1}] \leq \mu$ for some $\mu = O(1/\ln n)$. We can now apply Eq. (6) (cf. Step 5) to conclude that the expected routing time for a random pair is $\mathbf{E}[T] = \Omega(\ln n)/\mu = \Omega(\ln^2 n)$.

5.3 Distance sequence for GreedyBiDir

In the analysis of GreedyBiDir we will use a different definition for the distances d_i than in the analysis of GreedyUniDir, where d_i was just the distance between node u_i that has the message after i steps and target t .

Below we will write $d(U, t)$ for a node set U to denote the minimum distance between some node from U and t , i.e., $d(U, t) = \min_{v \in U} d(v, t)$. Further, by $N_{out}(v)$ we denote the set of out-contacts of v . For each $i \geq 0$, let

$$r_i = \begin{cases} d(\{u_i\} \cup N_{out}(u_i), t), & \text{if } u_i \rightarrow u_{i-1}; \\ d(u_i, t), & \text{otherwise;} \end{cases}$$

i.e., if u_i is an in-contact of the previous node in the routing path, then r_i is the minimum distance of u_i and its out-contacts from the target; otherwise, r_i is just the distance of u_i from the target. Then we define $d_0 = r_0 = d(s, t)$, and for $i \geq 1$,

$$d_i = \min \left\{ r_i, \left\lfloor \left(1 - \frac{1}{\ln n}\right) \cdot d_{i-1} \right\rfloor \right\}.$$

The second quantity inside $\min\{\}$ ensures that $d_i \leq d_{i-1} - 1$ when $d_{i-1} > 0$, and also that $d_i = 0$ after at most $i =$

$O(\log^2 n)$ steps. The latter is useful for the analysis because when the set of nodes visited grows very large it affects significantly the edge distribution of nodes not visited yet.

From the definitions above it is immediate that $d_i \leq r_i \leq d(u_i, t)$. Hence, for any node pair s, t the routing time of GreedyBiDir is lower bounded by $\min\{i: d_i = 0\}$. We have thus completed Step 1 of the analysis for GreedyBiDir (cf. Section 5.1). For Step 2 we must bound the distribution of d_i ; we use the next result.

Claim 5 For any $b \geq \delta > \ln n$ and $1 \leq r < \lfloor (1 - \frac{1}{\ln n})\delta \rfloor$ we have

$$\begin{aligned} \Pr(d_i \leq r \mid d_{i-1} = \delta, d(u_{i-1}, t) = b) \\ = O\left(\ln\left(\frac{\delta}{\delta-r}\right) / \ln n + f_{2-hop}(r, b)\right), \end{aligned}$$

where

$$\begin{aligned} f_{2-hop}(r, b) := \sum_{j=r+1}^{b-1} \sum_{k \geq 2} \left(\min \left\{ q_k j^{\ell-1}, \frac{kq_k}{\ln n \cdot (b-j)} \right\} \right. \\ \left. \cdot \min \left\{ 1, k \ln\left(\frac{j}{j-r}\right) / \ln n \right\} \right). \end{aligned}$$

In the above upper bound on the probability of $d_i \leq r$, the $O(\ln(\frac{\delta}{\delta-r})/\ln n)$ term bounds the probability that u_i belongs to $B_r(r)$; this term is the same as the bound we used for the case of GreedyUniDir (cf. Eq. (7)). The term $O(f_{2-hop}(r, b))$ bounds the probability that $u_i \rightarrow u_{i-1}$ and some out-contact of u_i belongs to $B_r(r)$. Intuitively, in the definition of $f_{2-hop}(r, b)$, the quantities $q_k j^{\ell-1}$ and $\frac{kq_k}{\ln n \cdot (b-j)}$ are upper bounds on the probability of the event that u_{i-1} has some in-contact v with $C_v = k$ on the sphere $S_r(j)$. The quantity $k \ln(\frac{j}{j-r})/\ln n$ then bounds the probability that a given node $v \in S_r(j)$ with $C_v = k$ has an out-contact in $B_r(r)$. In the following sections we will bound f_{2-hop} for the different values of α . Bounds for f_{2-hop} that are larger than $O(\ln(\frac{\delta}{\delta-r})/\ln n)$ will yield lower bounds for the routing time that are smaller than $\Omega(\ln^2 n)$.

Proof of Claim 5 We fix the path $u_0 \dots u_{i-1}$, and also the out-contacts and in-contacts of each of the nodes u_0, \dots, u_{i-2} . Further, we reveal whether $u_{i-1} \rightarrow u_{i-2}$ holds, and if it does, we fix the out-contacts of u_{i-1} . From these we can also compute d_0, \dots, d_{i-1} . Suppose that $d(u_{i-1}, t) = b$ and $d_{i-1} = \delta$.

First we bound the probability that u_{i-1} has some out-contact in $B_r(r)$. If $u_{i-1} \rightarrow u_{i-2}$ then this probability is zero, as by definition in this case we have $d(N_{out}(u_{i-1}), t) \geq r_{i-1} \geq d_{i-1} = \delta > r$. Hence, we assume that $u_{i-1} \not\rightarrow u_{i-2}$ (or $i = 1$). We have $u_{i-1} \not\rightarrow H := \{u_0, \dots, u_{i-2}\}$, and thus

$$\begin{aligned} \Pr(u_{i-1} \rightarrow B_r(r)) \\ \stackrel{\text{Obs.2(b)}}{\leq} (1 + o(1)) \cdot \Pr_{\mathcal{K}(\ell, n, 2)}(u_{i-1} \rightarrow B_r(r)) \\ \stackrel{\text{Obs.1(a)}}{=} O\left(\ln\left(\frac{b}{b-r}\right) / \ln n\right) = O\left(\ln\left(\frac{\delta}{\delta-r}\right) / \ln n\right), \end{aligned}$$

where the last relation holds because $\delta \leq b$ and the function $x/(x-r)$ is decreasing for $x > r$.

Next we establish the same bound for the probability that u_{i-1} has some in-contact in $B_t(r)$. For each $v \in B_t(r)$ we have $v \not\rightarrow H$, and from Observation 2(c), $\Pr(v \rightarrow u_{i-1}) \leq \Pr_{\mathcal{K}(\ell, n, 3)}(v \rightarrow u_{i-1})$. From this it follows that $\Pr(B_t(r) \rightarrow u_{i-1}) \leq \Pr_{\mathcal{K}(\ell, n, 3)}(B_t(r) \rightarrow u_{i-1})$, and thus

$$\begin{aligned} \Pr(B_t(r) \rightarrow u_{i-1}) &\leq 3 \Pr_{\mathcal{K}(\ell, n, 1)}(B_t(r) \rightarrow u_{i-1}) \\ &\stackrel{\text{Obs.1(g)}}{\leq} 3 \Pr_{\mathcal{K}(\ell, n, 1)}(u_{i-1} \rightarrow B_t(r)) \\ &\stackrel{\text{Obs.1(a)}}{=} O\left(\ln\left(\frac{b}{b-r}\right) / \ln n\right) \\ &= O\left(\ln\left(\frac{\delta}{\delta-r}\right) / \ln n\right), \end{aligned}$$

as before.

Last, we bound the probability that u_{i-1} has some in-contact v for which $r < d(v, t) < b$ and $v \rightarrow B_t(r)$. For each node v with $d(v, t) < b$ we have $v \not\rightarrow H$; we make this explicit as conditioning in the probability statements below, so that we can remove this condition. We have

$$\begin{aligned} \Pr(v \rightarrow u_{i-1} \wedge C_v = k \mid v \not\rightarrow H) &\stackrel{\text{Obs.2(b)}}{\leq} 2 \Pr(v \rightarrow u_{i-1} \wedge C_v = k) \\ &= 2q_k \cdot \Pr(v \rightarrow u_{i-1} \mid C_v = k) \\ &\leq 2q_k \cdot \min\{1, kp_{d(v, u_{i-1})}\}. \end{aligned}$$

Let $O_{j,k}$, for $r < j < b$ and $k \geq 2$, be the set of nodes $v \in S_t(j)$ with $C_v = k$ that are in-contacts of u_{i-1} , i.e., $O_{j,k} = \{v \in S_t(j) : C_v = k, v \rightarrow u_{i-1}\}$. We have

$$\begin{aligned} \mathbf{E}[|O_{j,k}|] &= \sum_{v \in S_t(j)} \Pr(v \rightarrow u_{i-1} \wedge C_v = k \mid v \not\rightarrow H) \\ &\leq 2q_k \sum_{v \in S_t(j)} \min\{1, kp_{d(v, u_{i-1})}\} \\ &\leq 2q_k \cdot \min\left\{|S_t(j)|, \sum_{v \in S_t(j)} kp_{d(v, u_{i-1})}\right\}. \end{aligned}$$

Also, $|S_t(j)| = O(j^{\ell-1})$ from Eq. (1) on page 7, and

$$\begin{aligned} \sum_{v \in S_t(j)} p_{d(v, u_{i-1})} &= \Pr_{\mathcal{K}(\ell, n, 1)}(u_{i-1} \rightarrow S_t(j)) \\ &\stackrel{\text{Obs.1(e)}}{=} O\left(\frac{1}{\ln n \cdot (b-j)}\right). \end{aligned}$$

Therefore,

$$\mathbf{E}[|O_{j,k}|] = O\left(q_k \cdot \min\left\{j^{\ell-1}, \frac{k}{\ln n \cdot (b-j)}\right\}\right).$$

We now bound the probability that a given node from $O_{j,t}$ has an out-contact in $B_t(r)$. For each $v \in S_t(j)$, we have

$$\begin{aligned} \Pr(v \rightarrow B_t(r) \mid C_v = k, v \not\rightarrow H) &\leq \Pr(v \rightarrow B_t(r) \mid C_v = k, v \not\rightarrow H) \\ &\stackrel{\text{Obs.2(f)}}{\leq} \frac{\Pr(v \rightarrow B_t(r) \mid C_v = k)}{\Pr(v \not\rightarrow H \mid C_v = 1)} \\ &\stackrel{\text{Obs.1(f)}}{=} \Pr(v \rightarrow B_t(r) \mid C_v = k) / (1 - o(1)) \\ &\leq k \cdot \Pr(v \rightarrow B_t(r) \mid C_v = 1) / (1 - o(1)) \\ &\stackrel{\text{Obs.1(e)}}{=} O\left(k \ln\left(\frac{j}{j-r}\right) / \ln n\right). \end{aligned}$$

From the two results above it follows that the expected number of in-contacts v of u_{i-1} for which $r < d(v, t) < b$ and $v \rightarrow B_t(r)$ is upper bounded by

$$\begin{aligned} &\sum_{j=r+1}^{b-1} \sum_{k \geq 2} \mathbf{E}[|O_{j,k}|] \cdot \min\left\{1, O\left(k \ln\left(\frac{j}{j-r}\right) / \ln n\right)\right\} \\ &= O\left(\sum_{j=r+1}^{b-1} \sum_{k \geq 2} \left(\min\left\{q_k j^{\ell-1}, \frac{kq_k}{\ln n \cdot (b-j)}\right\}\right) \cdot \min\left\{1, k \ln\left(\frac{j}{j-r}\right) / \ln n\right\}\right); \end{aligned}$$

and by Markov's Inequality the same upper bound applies to the probability that at least one such v exists.

Combining the last bound with the bounds for $\Pr(B_t(r) \rightarrow u_{i-1})$ and $\Pr(u_{i-1} \rightarrow B_t(r))$ shown earlier, and using the union bound completes the proof of Claim 5. \square

5.4 Lower bound for GreedyBiDir for $\alpha > 3$

We prove a lower bound of $\Omega(\log^2 n)$ on the expected routing time of GreedyBiDir for a random pair s, t in $\mathcal{SVW}(\ell, n, \alpha)$, when $\alpha > 3$.

First we bound the distribution of d_i (cf. Step 2 in Section 5.1). We will use Claim 5, and thus we must compute an upper bound for f_{2-hop} . From the definition of f_{2-hop} , it follows that for any integers δ, b, r for which $\ln n < \delta \leq b \leq d_0 = d(s, t)$ and $1 \leq r < \lfloor (1 - \frac{1}{\ln n}) \delta \rfloor$,

$$\begin{aligned} f_{2-hop}(r, b) &\leq \sum_{j=r+1}^{b-1} \sum_{k \geq 2} \frac{kq_k}{\ln n \cdot (b-j)} \cdot k \ln\left(\frac{j}{j-r}\right) / \ln n \\ &= \frac{1}{\nu \ln^2 n} \sum_{j=r+1}^{b-1} \frac{1}{b-j} \cdot \ln\left(\frac{j}{j-r}\right) \cdot \sum_{k \geq 2} \frac{1}{k^{\alpha-2}}. \end{aligned}$$

Since $\alpha > 3$, the normalizing factor ν is $\Theta(1)$ and $\sum_{k \geq 2} \frac{1}{k^{\alpha-2}} = \Theta(1)$. Further, we show below that

$$\sum_{j=r+1}^{b-1} \frac{1}{b-j} \cdot \ln\left(\frac{j}{j-r}\right) = O\left(\ln n \cdot \ln\left(\frac{\delta}{\delta-r}\right)\right). \quad (10)$$

It follows

$$f_{2-hop}(r, b) = O\left(\ln\left(\frac{\delta}{\delta-r}\right)/\ln n\right).$$

Claim 5 then yields that $\Pr(d_i \leq r \mid d_{i-1} = \delta) = O\left(\ln\left(\frac{\delta}{\delta-r}\right)/\ln n\right)$, for $1 \leq r < \lfloor(1 - \frac{1}{\ln n})\delta\rfloor$, which is the same bound as the one we had for GreedyUniDir (cf. Eq. (7) in Section 5.2). Thus the following definition for f satisfies Step 2 for a sufficiently large constant $c > 0$,

$$f(r, \delta) = \begin{cases} c \ln\left(\frac{\delta}{\delta-r}\right)/\ln n, & \text{if } 1 \leq r < \lfloor(1 - \frac{1}{\ln n})\delta\rfloor; \\ 1, & \text{if } \lfloor(1 - \frac{1}{\ln n})\delta\rfloor \leq r \leq \delta - 1. \end{cases} \quad (11)$$

Next, from Claim 4 (cf. Step 3) it follows

$$\begin{aligned} \mathbf{E}[\partial l_i \mid d_{i-1} = \delta] &\leq \sum_{1 \leq r \leq \delta-1} f(r, \delta)/r \\ &= \sum_{1 \leq r < \lfloor(1 - \frac{1}{\ln n})\delta\rfloor} \frac{c}{r} \ln\left(\frac{\delta}{\delta-r}\right)/\ln n + \sum_{\lfloor(1 - \frac{1}{\ln n})\delta\rfloor \leq r \leq \delta-1} \frac{1}{r}. \end{aligned}$$

Both sums in the last line are $O(1/\ln n)$: the first one because of Eq. (8) and (9) in Section 5.2, and the second because it is $O\left(\frac{\delta}{\ln n} \cdot \frac{1}{\delta}\right) = O(1/\ln n)$. Thus, $\mathbf{E}[\partial l_i \mid d_{i-1}] \leq \mu$ for some $\mu = O(1/\ln n)$, and from Eq. (6) (cf. Step 5) we conclude that the expected routing time for a random pair is $\mathbf{E}[T] = \Omega(\ln(n)/\mu) = \Omega(\ln^2 n)$.

Proof of Eq. (10) We distinguish two cases.

Case $r \leq \delta/2$:

$$\begin{aligned} \sum_{j=r+1}^{b-1} \frac{1}{b-j} \cdot \ln\left(\frac{j}{j-r}\right) &= \sum_{j=r+1}^{b-1} \frac{1}{b-j} \cdot \ln\left(1 + \frac{r}{j-r}\right) \\ &\leq \sum_{j=r+1}^{b-1} \frac{1}{b-j} \cdot \frac{r}{j-r} \\ &= 2 \sum_{j=r+1}^{(b+r)/2} \frac{1}{b-j} \cdot \frac{r}{j-r} \\ &\quad (\text{by symmetry}) \\ &\leq 2 \sum_{j=r+1}^{(b+r)/2} \frac{2}{b-r} \cdot \frac{r}{j-r} \\ &\quad (\text{we set } j = (b+r)/2 \text{ in the} \\ &\quad \text{first fraction}) \\ &= O\left(\frac{r}{b-r} \cdot \ln(b-r)\right) \\ &= O\left(\frac{r}{\delta-r} \cdot \ln n\right). \end{aligned}$$

Finally, to obtain (10) we observe that $\frac{r}{\delta-r} \leq \log\left(\frac{\delta}{\delta-r}\right)$: we have $\frac{\delta}{\delta-r} = 1 + \frac{r}{\delta-r} \geq 2^{r/(\delta-r)}$, because $1+x \geq 2^x$ for $0 \leq x \leq 1$, and $\frac{r}{\delta-r} \leq \frac{\delta/2}{\delta-\delta/2} = 1$.

Case $r > \delta/2$: We break the sum into two sums, for $i \leq (b+r)/2$ and $i > (b+r)/2$.

$$\sum_{j=r+1}^{(b+r)/2} \frac{1}{b-j} \cdot \ln\left(\frac{j}{j-r}\right) \leq \sum_{j=r+1}^{(b+r)/2} \frac{2}{b-r} \cdot \ln\left(\frac{j}{j-r}\right)$$

(we set $j = (b+r)/2$ in the first fraction)

$$\leq \sum_{j=r+1}^{(b+r)/2} \frac{2}{b-r} \cdot \ln b = O(\ln b) = O(\ln n).$$

$$\sum_{j=(b+r)/2}^{b-1} \frac{1}{b-j} \cdot \ln\left(\frac{j}{j-r}\right) \leq \sum_{j=(b+r)/2}^{b-1} \frac{1}{b-j} \cdot \ln\left(\frac{b+r}{b-r}\right)$$

(we set $j = (b+r)/2$ in the second fraction)

$$= O(\ln(b-r)) \cdot \ln\left(\frac{b+r}{b-r}\right) = O(\ln n \cdot \ln\left(\frac{b+r}{b-r}\right)).$$

To obtain (10) it suffices to prove that $\frac{b+r}{b-r} \leq \left(\frac{\delta}{\delta-r}\right)^2$: we have $\frac{b+r}{b-r} \leq \frac{\delta+r}{\delta-r}$ and $\frac{\delta+r}{\delta-r} \leq \left(\frac{\delta}{\delta-r}\right)^2$, where the latter holds because is equivalent to $(\delta+r) \cdot (\delta-r) \leq \delta^2$, which is equivalent to $\delta^2 - r^2 \leq \delta^2$. This completes the proof of Eq. (10). \square

5.5 Lower bound for GreedyBiDir for $\alpha = 3$

We prove a lower bound of $\Omega(\log^2 n / \log \log n)$ on the expected routing time of GreedyBiDir for a random pair s, t in $\mathcal{SW}(\ell, n, 3)$.

The difference from the case of $\mathcal{SW}(\ell, n, \alpha)$ with $\alpha > 3$, is that the bound we will show for f_{2-hop} will be larger by a factor of $O(\log \log n)$, and consequently the bound on the routing time will be smaller by a factor of $1/O(\log \log n)$. To bound f_{2-hop} we distinguish between nodes with at most $\ln^2 n$ out-contacts, and nodes with more out-contacts. From the definition of f_{2-hop} in Claim 5, it follows that for $\ln n < \delta \leq b \leq d_0$ and $1 \leq r < \lfloor(1 - \frac{1}{\ln n})\delta\rfloor$,

$f_{2-hop}(r, b)$

$$\begin{aligned} &\leq \sum_{j=r+1}^{b-1} \sum_{k \geq 2} \frac{kq_k}{\ln n \cdot (b-j)} \cdot \min\left\{1, k \ln\left(\frac{j}{j-r}\right)/\ln n\right\} \\ &\leq \sum_{j=r+1}^{b-1} \sum_{k=2}^{\ln^2 n} \frac{kq_k}{\ln n \cdot (b-j)} \cdot k \ln\left(\frac{j}{j-r}\right)/\ln n \\ &\quad + \sum_{j=r+1}^{b-1} \sum_{k > \ln^2 n} \frac{kq_k}{\ln n \cdot (b-j)} \cdot 1 \\ &= \frac{1}{\nu \ln^2 n} \sum_{j=r+1}^{b-1} \frac{1}{b-j} \cdot \ln\left(\frac{j}{j-r}\right) \cdot \sum_{k=2}^{\ln^2 n} \frac{1}{k} \\ &\quad + \frac{1}{\nu \ln n} \sum_{j=r+1}^{b-1} \frac{1}{b-j} \cdot \sum_{k > \ln^2 n} \frac{1}{k^2}. \end{aligned}$$

We bound the quantities in the last two lines: $\nu = \Theta(1)$; $\sum_{j=r+1}^{b-1} \frac{1}{b-j} \ln\left(\frac{j}{j-r}\right) = O(\ln n \cdot \ln\left(\frac{\delta}{\delta-r}\right))$, from Eq.(10); $\sum_{k=2}^{\ln^2 n} \frac{1}{k} = O(\ln \ln n)$; $\sum_{j=r+1}^{b-1} \frac{1}{b-j} = O(\ln n)$; and $\sum_{k > \ln^2 n} \frac{1}{k^2} = O(1/\ln^2 n)$. Applying these above yields

$$f_{2\text{-hop}}(r, b) = O\left(\ln \ln n \cdot \ln\left(\frac{\delta}{\delta-r}\right) / \ln n\right).$$

From Claim 5 it follows that the same bound holds for $\Pr(d_i \leq r \mid d_{i-1} = \delta)$. Thus we can define $f(r, \delta)$ as

$$\begin{cases} c \ln \ln n \cdot \ln\left(\frac{\delta}{\delta-r}\right) / \ln n, & \text{if } 1 \leq r < \lfloor (1 - \frac{1}{\ln n}) \delta \rfloor; \\ 1, & \text{if } \lfloor (1 - \frac{1}{\ln n}) \delta \rfloor \leq r \leq \delta - 1. \end{cases}$$

Applying now Claim 4 as in Section 5.4, and observing that the above definition of $f(r, \delta)$ differs from the one in Eq. (11) by an $O(\ln \ln n)$ factor, we obtain that $\mathbf{E}[\partial l_i \mid d_{i-1}] \leq \mu$ for some $\mu = O(\ln \ln n / \ln n)$. Finally, from Eq. (6) it follows that the expected routing time for a random pair is $\mathbf{E}[T] = \Omega(\ln n / \mu) = \Omega(\ln^2 n / \ln \ln n)$.

5.6 Lower bound for GreedyBiDir for $2 < \alpha < 3$

We prove a lower bound of $\Omega(\log^{\alpha-1} n)$ on the expected routing time of GreedyBiDir for a random pair s, t in $\mathcal{SW}(\ell, n, 3)$, when $2 < \alpha < 3$.

In this case the computations needed to bound $f_{2\text{-hop}}$ are a bit more involved than in the previous two cases. From the definition of $f_{2\text{-hop}}$ in Claim 5, it follows that for $\ln n < \delta \leq b \leq d_0$ and $1 \leq r < \lfloor (1 - \frac{1}{\ln n}) \delta \rfloor$, and for $k_j := \ln n / \ln\left(\frac{j}{j-r}\right)$,

$$\begin{aligned} f_{2\text{-hop}}(r, b) &\leq \sum_{j=r+1}^{b-1} \sum_{k \geq 2} \frac{k q_k}{\ln n \cdot (b-j)} \cdot \min\left\{1, k \ln\left(\frac{j}{j-r}\right) / \ln n\right\} \\ &= \sum_{j=r+1}^{b-1} \sum_{2 \leq k \leq k_j} \frac{k q_k}{\ln n \cdot (b-j)} \cdot k \ln\left(\frac{j}{j-r}\right) / \ln n \\ &\quad + \sum_{j=r+1}^{b-1} \sum_{k > k_j} \frac{k q_k}{\ln n \cdot (b-j)} \cdot 1 \\ &= \frac{1}{\nu \ln^2 n} \sum_{j=r+1}^{b-1} \frac{\ln\left(\frac{j}{j-r}\right)}{b-j} \sum_{2 \leq k \leq k_j} \frac{1}{k^{\alpha-2}} \\ &\quad + \frac{1}{\nu \ln n} \sum_{j=r+1}^{b-1} \frac{1}{b-j} \sum_{k > k_j} \frac{1}{k^{\alpha-1}}. \end{aligned}$$

We have $\nu = \Theta(1)$, $\sum_{2 \leq k \leq k_j} \frac{1}{k^{\alpha-2}} = O(k_j^{3-\alpha})$, and $\sum_{k > k_j} \frac{1}{k^{\alpha-1}} = O(1/k_j^{\alpha-2})$. Applying these above and plug-

ging the value of $k_j = \ln n / \ln\left(\frac{j}{j-r}\right)$ yields

$$\begin{aligned} f_{2\text{-hop}}(r, b) &= O\left(\frac{1}{\ln^2 n} \sum_{j=r+1}^{b-1} \frac{\ln\left(\frac{j}{j-r}\right)}{b-j} \cdot \left(\frac{\ln n}{\ln\left(\frac{j}{j-r}\right)}\right)^{3-\alpha}\right. \\ &\quad \left. + \frac{1}{\ln n} \sum_{j=r+1}^{b-1} \frac{1}{b-j} \cdot \left(\frac{\ln\left(\frac{j}{j-r}\right)}{\ln n}\right)^{\alpha-2}\right) \\ &= O\left(\frac{1}{\ln^{\alpha-1} n} \sum_{j=r+1}^{b-1} \frac{1}{b-j} \cdot \ln^{\alpha-2}\left(\frac{j}{j-r}\right)\right). \end{aligned}$$

Further, we will show similar to Eq. (10) that

$$\begin{aligned} &\sum_{j=r+1}^{b-1} \frac{1}{b-j} \cdot \ln^{\alpha-2}\left(\frac{j}{j-r}\right) \\ &= \begin{cases} O\left(\ln n \cdot \left(\frac{r}{\delta}\right)^{\alpha-2}\right), & \text{if } 1 \leq r \leq \delta/2; \\ O\left(\ln n \cdot \ln\left(\frac{\delta}{\delta-r}\right)\right), & \text{if } \delta/2 < r \leq \delta - 1. \end{cases} \end{aligned} \quad (12)$$

Therefore, $f_{2\text{-hop}}(r, b)$ is

$$\begin{cases} O\left(\left(\frac{r}{\delta}\right)^{\alpha-2} / \ln^{\alpha-2} n\right), & \text{if } 1 \leq r \leq \delta/2; \\ O\left(\ln\left(\frac{\delta}{\delta-r}\right) / \ln^{\alpha-2} n\right), & \text{if } \delta/2 < r < \lfloor (1 - \frac{1}{\ln n}) \delta \rfloor. \end{cases}$$

From this and Claim 5, it follows $\Pr(d_i \leq r \mid d_{i-1} = \delta) \leq f(r, \delta)$ for $\delta > \ln n$ and $1 \leq r \leq \delta - 1$, if $f(r, \delta)$ is

$$\begin{cases} c \cdot \left(\frac{r}{\delta}\right)^{\alpha-2} / \ln^{\alpha-2} n, & \text{if } 1 \leq r \leq \delta/2; \\ c \cdot \ln\left(\frac{\delta}{\delta-r}\right) / \ln^{\alpha-2} n, & \text{if } \delta/2 < r < \lfloor (1 - \frac{1}{\ln n}) \delta \rfloor; \\ 1, & \text{if } \lfloor (1 - \frac{1}{\ln n}) \delta \rfloor \leq r \leq \delta - 1. \end{cases}$$

Then from Claim 4 we have

$$\begin{aligned} \mathbf{E}[\partial l_i \mid d_{i-1} = \delta] &\leq \sum_{1 \leq r \leq \delta-1} f(r, \delta) / r \\ &= \sum_{1 \leq r \leq \delta/2} \frac{c}{r} \cdot \left(\frac{r}{\delta}\right)^{\alpha-2} / \ln^{\alpha-2} n \\ &\quad + \sum_{\delta/2 < r < \lfloor (1 - \frac{1}{\ln n}) \delta \rfloor} \frac{c}{r} \cdot \ln\left(\frac{\delta}{\delta-r}\right) / \ln^{\alpha-2} n \\ &\quad + \sum_{\lfloor (1 - \frac{1}{\ln n}) \delta \rfloor \leq r \leq \delta-1} \frac{1}{r}. \end{aligned}$$

All the three sums above are $O(1/\ln^{\alpha-2} n)$: the first one because

$$\begin{aligned} \sum_{1 \leq r \leq \delta/2} \frac{1}{r} \cdot \left(\frac{r}{\delta}\right)^{\alpha-2} &= \frac{1}{\delta^{\alpha-2}} \sum_{1 \leq r \leq \delta/2} r^{\alpha-3} \\ &= \frac{1}{\delta^{\alpha-2}} O(\delta^{\alpha-2}) = O(1); \end{aligned}$$

the second one because of Eq. (8) and (9) in Section 5.2; and the third one is $O\left(\frac{\delta}{\ln n} \cdot \frac{1}{\delta}\right) = O(1/\ln n)$. Therefore, $\mathbf{E}[\partial l_i \mid$

$d_{i-1}] \leq \mu$ for some $\mu = O(1/\ln^{\alpha-2}n)$, and from Eq. (6) we conclude that the expected routing time for a random pair is $\mathbf{E}[T] = \Omega(\ln(n)/\mu) = \Omega(\ln^{\alpha-1}n)$.

Proof of Eq. (12) The proof is similar to that of Eq. (10).

Case $r \leq \delta/2$:

$$\begin{aligned} \sum_{j=r+1}^{b-1} \frac{1}{b-j} \ln^{\alpha-2} \left(\frac{j}{j-r} \right) &\leq \sum_{j=r+1}^{b-1} \frac{1}{b-j} \cdot \left(\frac{r}{j-r} \right)^{\alpha-2} \\ &\leq 2 \sum_{j=(b+r)/2}^{b-1} \frac{1}{b-j} \cdot \left(\frac{r}{j-r} \right)^{\alpha-2} \\ &\quad \text{(the sum of the first half terms} \\ &\quad \text{is smaller than the other half)} \\ &\leq 2 \sum_{j=(b+r)/2}^{b-1} \frac{1}{b-j} \cdot \left(\frac{2r}{b-r} \right)^{\alpha-2} \\ &\quad \text{(we set } j = (b+r)/2 \text{ in the} \\ &\quad \text{second fraction)} \\ &= O \left(\ln(b-r) \cdot \left(\frac{2r}{b-r} \right)^{\alpha-2} \right) \\ &= O \left(\ln n \cdot \left(\frac{r}{\delta} \right)^{\alpha-2} \right). \end{aligned}$$

Case $r > \delta/2$: We break the sum into two sums, for $i \leq (b+r)/2$ and $i > (b+r)/2$.

$$\begin{aligned} \sum_{j=r+1}^{(b+r)/2} \frac{1}{b-j} \cdot \ln^{\alpha-2} \left(\frac{j}{j-r} \right) &\leq \sum_{j=r+1}^{(b+r)/2} \frac{2}{b-r} \cdot \ln^{\alpha-2} b \\ &= O(\ln^{\alpha-2} b) = o(\ln n), \end{aligned}$$

as $\alpha < 3$.

$$\begin{aligned} \sum_{j=(b+r)/2}^{b-1} \frac{1}{b-j} \cdot \ln^{\alpha-2} \left(\frac{j}{j-r} \right) &\leq \sum_{j=(b+r)/2}^{b-1} \frac{1}{b-j} \cdot \ln^{\alpha-2} \left(\frac{b+r}{b-r} \right) \\ &= O(\ln(b-r)) \cdot \ln^{\alpha-2} \left(\frac{b+r}{b-r} \right) \\ &= O \left(\ln n \cdot \ln^{\alpha-2} \left(\frac{\delta}{\delta-r} \right) \right) \\ &= O \left(\ln n \cdot \ln \left(\frac{\delta}{\delta-r} \right) \right). \end{aligned}$$

This completes the proof of Eq. (12). \square

5.7 Lower bound for GreedyBiDir for $\alpha < 2$

We show that when $\alpha < 2$, the expected routing time of GreedyUniDir in $\mathcal{SW}(\ell, n, \alpha)$ is $\Omega(\log^2 n)$ for *some* pairs s, t .

The normalizing factor v of distribution q_k is a polynomial function of n in this case. Specifically, we have $v = \sum_{2 \leq i \leq k_{\max}} (i^{1-\alpha} - i^{-\alpha}) = \Theta((k_{\max})^{2-\alpha})$, where $k_{\max} = n^\gamma$ and $0 < \gamma \leq \ell$ is a constant. It follows that the probability

$1 - p_1$ for a node v to have more than one out-contacts is polynomially small. As a result, with probability $\Omega(1)$ the message visits only nodes v with $C_v = 1$, if $d(s, t)$ is a sufficiently small polynomial in n .

The above intuition is captured by a simple bound on f_{2-hop} we show next. From the definition of f_{2-hop} in Claim 5, it follows that for $\ln n < \delta \leq b \leq d_0$ and $1 \leq r < \lfloor (1 - \frac{1}{\ln n})\delta \rfloor$,

$$\begin{aligned} f_{2-hop}(r, b) &\leq \sum_{j=r+1}^{d_0-1} \sum_{k=2}^{k_{\max}} q_k j^{\ell-1} \cdot 1 \\ &= \sum_{j=r+1}^{d_0-1} j^{\ell-1} \sum_{k=2}^{k_{\max}} \frac{1}{v k^\alpha} = O \left(\frac{d_0^\ell}{(k_{\max})^{2-\alpha}} \sum_{k=2}^{k_{\max}} \frac{1}{k^\alpha} \right), \end{aligned}$$

and thus

$$f_{2-hop}(r, b) = \begin{cases} O(d_0^\ell / (k_{\max})^{2-\alpha}), & \text{if } 1 < \alpha < 2; \\ O(d_0^\ell \cdot \ln(k_{\max}) / k_{\max}), & \text{if } \alpha = 1; \\ O(d_0^\ell / k_{\max}), & \text{if } 0 \leq \alpha < 1. \end{cases}$$

Therefore, if⁶

$$d_0 \leq \lambda := \left(\frac{k_{\max}^{(2-\alpha)/2}}{\ln^2 n} \right)^{1/\ell} = n^{\Theta(1)},$$

then it follows that $f_{2-hop}(r, b) = O(1/\ln^2 n)$. From this bound and Claim 5 we obtain that $\Pr(d_i \leq r \mid d_{i-1} = \delta) \leq f(r, \delta)$, for $1 \leq r < \delta \leq d_0 \leq \lambda$, if $f(r, \delta)$ is

$$\begin{cases} c \ln \left(\frac{\delta}{\delta-r} \right) / \ln n + c / \ln^2 n, & \text{if } 1 \leq r < \lfloor (1 - \frac{1}{\ln n})\delta \rfloor; \\ 1, & \text{if } \lfloor (1 - \frac{1}{\ln n})\delta \rfloor \leq r \leq \delta - 1. \end{cases}$$

Claim 4 then yields

$$\begin{aligned} \mathbf{E}[\partial l_i \mid d_{i-1} = \delta] &\leq \sum_{1 \leq r \leq \delta-1} f(r, \delta) / r \\ &= \sum_{1 \leq r < \lfloor (1 - \frac{1}{\ln n})\delta \rfloor} \left(\frac{c}{r} \cdot \ln \left(\frac{\delta}{\delta-r} \right) / \ln n + \frac{c}{r \ln^2 n} \right) \\ &\quad + \sum_{\lfloor (1 - \frac{1}{\ln n})\delta \rfloor \leq r \leq \delta-1} \frac{1}{r} \\ &= O(1/\ln n), \end{aligned}$$

similar to case $\alpha > 3$, in Section 5.4. Thus, $\mathbf{E}[\partial l_i \mid d_{i-1}] \leq \mu$ for some $\mu = O(1/\ln n)$, and from Eq. (5) (cf. Step 4) we conclude that the expected routing time for a pair s, t with $d(s, t) = \lambda$ is $\Omega(\ln(\lambda/\ln n)/\mu) = \Omega(\ln^2 n)$.

⁶ The value of λ is not optimized; it was chosen so that the same simple expression works for all three cases.

5.8 Lower bound for GreedyBiDir for $\alpha = 2$

We show that for some pairs s, t in $\mathcal{SW}(\ell, n, 2)$ the expected routing time of GreedyBiDir is $\Omega(\log^{4/3} n)$.

Recall from Section 4.3 that $v = \Theta(\ln n)$ in this case, and thus the probability q_1 that a node has only one out-contact is $1 - \Theta(1/\ln n)$.

We will prove the lower bound for a pair s, t for which $d(s, t) = \lambda := 2^{\ln^{1/3} n}$. From the definition of f_{2-hop} in Claim 5, it follows that for $\ln n < \delta \leq b \leq d_0$ and $1 \leq r < \lfloor (1 - \frac{1}{\ln n})\delta \rfloor$,

$$\begin{aligned} f_{2-hop}(r, b) &\leq \sum_{j=r+1}^{b-1} \sum_{k \geq 2} \min \left\{ q_k j^{\ell-1}, \frac{kq_k}{\ln n \cdot (b-j)} \right\} \cdot 1 \\ &\leq \sum_{j=r+1}^{b-1} \sum_{k=2}^{b^\ell \ln n} \frac{kq_k}{\ln n \cdot (b-j)} + \sum_{j=r+1}^{b-1} \sum_{k > b^\ell \ln n} q_k j^{\ell-1} \\ &\leq \sum_{j=r+1}^{b-1} \sum_{k=2}^{b^\ell \ln n} \frac{1}{vk \ln n \cdot (b-j)} + \sum_{j=r+1}^{b-1} \sum_{k > b^\ell \ln n} \frac{j^{\ell-1}}{vk^2} \\ &= O\left(\frac{\ln^2 b}{v \ln n}\right) + O\left(\frac{b^\ell}{vb^\ell \ln n}\right) \\ &= O\left(\frac{\ln^2 \lambda}{\ln^2 n}\right), \end{aligned}$$

where for the last relation we used that $v = \Theta(\ln n)$ and $b \leq d_{s,t} = \lambda$. From this and Claim 5, it follows that $\Pr(d_i \leq r \mid d_{i-1} = \delta) \leq f(r, \delta)$ for $1 \leq r < \delta \leq d_0 \leq \lambda$, if $f(r, \delta)$ is

$$\begin{cases} c \ln\left(\frac{\delta}{\delta-r}\right) / \ln n + c \frac{\ln^2 \lambda}{\ln^2 n}, & \text{if } 1 \leq r < \lfloor (1 - \frac{1}{\ln n})\delta \rfloor; \\ 1, & \text{if } \lfloor (1 - \frac{1}{\ln n})\delta \rfloor \leq r \leq \delta - 1. \end{cases}$$

Claim 4 then yields

$$\begin{aligned} \mathbf{E}[\partial l_i \mid d_{i-1} = \delta] &\leq \sum_{1 \leq r \leq \delta-1} f(r, \delta) / r \\ &= \sum_{1 \leq r < \lfloor (1 - \frac{1}{\ln n})\delta \rfloor} \left(\frac{c}{r} \cdot \ln\left(\frac{\delta}{\delta-r}\right) / \ln n + \frac{c \ln^2 \lambda}{r \ln^2 n} \right) \\ &\quad + \sum_{\lfloor (1 - \frac{1}{\ln n})\delta \rfloor \leq r \leq \delta-1} \frac{1}{r}. \end{aligned}$$

Except for $\sum \frac{c \ln^2 \lambda}{r \ln^2 n}$, we have already seen that all the sums in the last two lines are $O(1/\ln n)$, e.g., in the analysis of case $\alpha > 3$. Further, we have

$$\sum_{1 \leq r \leq \delta} \frac{c \ln^2 \lambda}{r \ln^2 n} = O\left(\frac{\ln \delta \cdot \ln^2 \lambda}{\ln^2 n}\right) = O\left(\frac{\ln^3 \lambda}{\ln^2 n}\right) = O\left(\frac{1}{\ln n}\right),$$

because $\lambda = 2^{\ln^{1/3} n}$. Thus $\mathbf{E}[\partial l_i \mid d_{i-1}] \leq \mu := O(1/\ln n)$, and from Eq. (5) (cf. Step 4) we conclude that the routing time for pair s, t is $\Omega(\ln(\lambda/\ln n)/\mu) = \Omega(\ln^{3/4} n)$.

Acknowledgements We thank Vassos Hadzilacos and Philipp Woelfel for helpful discussions. We also thank the anonymous reviewers for their helpful feedback.

References

1. Abraham, I., Gavoille, C.: Object location using path separators. In: Proceedings of the 25th ACM Symposium on Principles of Distributed Computing (PODC), pp. 188–197 (2006)
2. Adamic, L.A., Lukose, R.M., Puniyani, A.R., Huberman, B.A.: Search in power-law networks. *Physical Review E* **64**, 46,135 (2001)
3. Albert, R., Barabási, A.L.: Statistical mechanics of complex networks. *Reviews of Modern Physics* **74**(1), 47–97 (2002)
4. Aspnes, J., Diamadi, Z., Shah, G.: Fault-tolerant routing in peer-to-peer systems. In: Proceedings of the 21st ACM Symposium on Principles of Distributed Computing (PODC), pp. 223–232 (2002)
5. Backstrom, L., Boldi, P., Rosa, M., Ugander, J., Vigna, S.: Four degrees of separation. In: Proceedings of the 3rd ACM Web Science Conference (WebSci), pp. 33–42 (2012)
6. Barrière, L., Fraigniaud, P., Kranakis, E., Krizanc, D.: Efficient routing in networks with long range contacts. In: Proceedings of the 15th International Symposium on Distributed Computing (DISC), pp. 270–284 (2001)
7. Bollobás, B., Riordan, O.: The diameter of a scale-free random graph. *Combinatorica* **24**(1), 5–34 (2004)
8. Chung, F.R.K., Lu, L.: The average distance in a random graph with given expected degrees. *Internet Mathematics* **1**(1), 91–113 (2003)
9. Coppersmith, D., Gamarnik, D., Sviridenko, M.: The diameter of a long range percolation graph. In: Proceedings of the 13th ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 329–337 (2002)
10. Dietzfelbinger, M., Rowe, J., Wegener, I., Woelfel, P.: Tight bounds for blind search on the integers. In: Proceedings of the 25th International Symposium on Theoretical Aspects of Computer Science (STACS), pp. 241–252 (2008)
11. Dietzfelbinger, M., Woelfel, P.: Tight lower bounds for greedy routing in uniform small world rings. In: Proceedings of the 41st ACM Symposium on Theory of Computing (STOC), pp. 591–600 (2009)
12. Dodds, P.S., Muhamad, R., Watts, D.J.: An experimental study of search in global social networks. *Science* **301**(5634), 827–829 (2003)
13. Dorogovtsev, S.N., Mendes, J.F.F.: Evolution of networks: From biological networks to the Internet and WWW. Oxford University Press (2003)

14. Duchon, P., Hanusse, N., Lebhar, E., Schabanel, N.: Could any graph be turned into a small-world? In: Proceedings of the 19th International Symposium on Distributed Computing (DISC), pp. 511–513 (2005)
15. Flammini, M., Moscardelli, L., Navarra, A., Pérennes, S.: Asymptotically optimal solutions for small world graphs. In: Proceedings of the 19th International Symposium on Distributed Computing (DISC), pp. 414–428 (2005)
16. Fraigniaud, P.: Greedy routing in tree-decomposed graphs. In: Proceedings of the 13th European Symposium on Algorithms (ESA), pp. 791–802 (2005)
17. Fraigniaud, P., Gavoille, C., Paul, C.: Eclecticism shrinks even small worlds. In: Proceedings of the 23rd ACM Symposium on Principles of Distributed Computing (PODC), pp. 169–178 (2004)
18. Fraigniaud, P., Giakkoupis, G.: The effect of power-law degrees on the navigability of small worlds. In: Proceedings of the 28th ACM Symposium on Principles of Distributed Computing (PODC), pp. 240–249 (2009)
19. Fraigniaud, P., Giakkoupis, G.: On the searchability of small-world networks with arbitrary underlying structure. In: Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC), pp. 389–398 (2010)
20. Giakkoupis, G., Hadzilacos, V.: On the complexity of greedy routing in ring-based peer-to-peer networks. In: Proceedings of the 26th ACM Symposium on Principles of Distributed Computing (PODC), pp. 99–108 (2007)
21. Giakkoupis, G., Schabanel, N.: Optimal path search in small worlds: Dimension matters. In: Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC), pp. 393–402 (2011)
22. Kim, B., Yoon, C., Han, S., Jeong, H.: Path finding strategies in scale-free networks. *Physical Review E* **65**, 027,103 (2002)
23. Kleinberg, J.: Navigation in a small world. *Nature* **406**, 845 (2000)
24. Kleinberg, J.: The small-world phenomenon: An algorithmic perspective. In: Proceedings of the 32nd ACM Symposium on Theory of Computing (STOC), pp. 163–170 (2000)
25. Kleinberg, J.: Small-world phenomena and the dynamics of information. In: Advances in Neural Information Processing Systems (NIPS) 14, pp. 431–438 (2001)
26. Kleinberg, J.: Complex networks and decentralized search algorithms. In: Proceedings of the International Congress of Mathematicians (ICM)(2006)
27. Lattanzi, S., Panconesi, A., Sivakumar, D.: Milgram-routing in social networks. In: Proceedings of the 20th ACM International Conference on World Wide Web (WWW), pp. 725–734 (2011)
28. Lebhar, E., Schabanel, N.: Almost optimal decentralized routing in long-range contact networks. In: Proceedings of the 31st International Colloquium on Automata, Languages, and Programming (ICALP), pp. 894–905 (2004)
29. Liben-Nowell, D., Novak, J., Kumar, R., Raghavan, P., Tomkins, A.: Geographic routing in social networks. *Proceedings of the National Academy of Sciences of the USA* **102**(33), 11,623–11,628 (2005)
30. Manku, G.S., Naor, M., Wieder, U.: Know thy neighbor’s neighbor: The power of lookahead in randomized P2P networks. In: Proceedings of the 36th ACM Symposium on Theory of Computing (STOC), pp. 54–63 (2004)
31. Martel, C., Nguyen, V.: Analyzing Kleinberg’s (and other) small-world models. In: Proceedings of the 23rd ACM Symposium on Principles of Distributed Computing (PODC), pp. 179–188 (2004)
32. Martel, C., Nguyen, V.: Analyzing and characterizing small-world graphs. In: Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 311–320 (2005)
33. Milgram, S.: The small world problem. *Psychology Today* **67**(1), 60–67 (1967)
34. Newman, M.E.J.: The structure and function of complex networks. *SIAM Review* **45**(2), 167–256 (2003)
35. Sarshar, N., Boykin, P.O., Roychowdhury, V.P.: Percolation search in power law networks: Making unstructured peer-to-peer networks scalable. In: Proceedings of the 4th IEEE International Conference on Peer-to-Peer Computing (P2P), pp. 2–9 (2004)
36. Simsek, Ö., Jensen, D.: Decentralized search in networks using homophily and degree disparity. In: Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI), pp. 304–310 (2005)
37. Slivkins, A.: Distance estimation and object location via rings of neighbors. In: Proceedings of the 24th ACM Symposium on Principles of Distributed Computing (PODC), pp. 41–50 (2005)
38. Watts, D.J., Strogatz, S.H.: Collective dynamics of ‘small-world’ networks. *Nature* **393**, 440–442 (1998)