

On the Searchability of Small-World Networks with Arbitrary Underlying Structure*

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

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

ABSTRACT

Revisiting the “small-world” experiments of the ’60s, Kleinberg observed that individuals are very effective at constructing short chains of acquaintances between any two people, and he proposed a mathematical model of this phenomenon. In this model, individuals are the nodes of a *base graph*, the square grid, capturing the *underlying structure* of the social network; and this base graph is augmented with additional edges from each node to a few *long-range contacts* of this node, chosen according to some natural distance-based distribution. In this augmented graph, a greedy search algorithm takes only a polylogarithmic number of steps in the graph size. Following this work, several papers investigated the correlations between underlying structure and long-range connections that yield efficient decentralized search, generalizing Kleinberg’s results to broad classes of underlying structures, such as metrics of bounded doubling dimension, and minor-excluding graphs.

We focus on the case of *arbitrary* base graphs. We show that for a simple long-range contact distribution consistent with empirical observations on social networks, a slight variation of greedy search, where the next hop is to a distant node only if it yields sufficient progress towards the target, requires $n^{o(1)}$ steps, where n is the number of nodes. Precisely, the expected number of steps for any source–target pair is at most $2^{(\log n)^{1/2+o(1)}}$. This bound almost matches the best known lower bound of $\Omega(2^{\sqrt{\log n}})$ steps, which applies to a general class of search algorithms. In the context of social networks, our result could be interpreted as: individuals may well be able to construct short chains between people regardless of the underlying structure of the social network.

*Research supported in part by the ANR projects ALADDIN and PROSE, and by the INRIA project GANG.

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

STOC’10, June 5–8, 2010, Cambridge, Massachusetts, USA
Copyright 2010 ACM 978-1-4503-0050-6/10/06 ...\$10.00.

Categories and Subject Descriptors

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

General Terms

Algorithms, Performance, Theory

Keywords

small worlds, decentralized search, social networks

1. INTRODUCTION

Small worlds have been an active research topic in many fields, including sociology, physics, and computer science. The scientific interest in small worlds began with the famous experiments conducted in the 1960s by sociologist S. Milgram [17]. These experiments quantified the *small-world phenomenon*, that is, the principle that all people are linked by short chains of acquaintances. Milgram’s findings were subsequently confirmed by others (see, e.g., [3]). Revisiting Milgram’s experiments, J. Kleinberg [9] observed another striking aspect of the small-world phenomenon: not only do short chains between people *exist*, but individuals are collectively very effective at *finding* them, using only local information. He then proposed a mathematical graph model for studying this aspect of small worlds (extending a model proposed in [19]). In the small-world model of [9], individuals are the nodes of a two-dimensional $n \times n$ grid graph, which is augmented with some additional *long-range* edges as follows. From each node u , one directed edge is added to some other node chosen at random, such that each node v is chosen with probability proportional to $d_{u,v}^{-\alpha}$, where $d_{u,v}$ is the grid distance between u and v , and $\alpha \geq 0$ is a parameter of the model. The edges of the augmented grid represent acquaintance relationships between people. Intuitively, the grid captures the *underlying structure* of the social network, defining a notion of distance between people (modeling, e.g., geographical distance). In this model, Kleinberg showed that if $\alpha \neq 2$ then *any* search algorithm that does not know the long-range edges of the nodes that have not been visited yet, requires an expected number of $\Omega(n^\gamma)$ steps, for most source–target pairs, where $\gamma > 0$ depends on α but not on n . However, if $\alpha = 2$ then the simple *greedy search* algorithm (where greediness is with respect to the grid distance) per-

forms in $n^{o(1)}$ steps. Precisely, for any source–target pair, greedy search takes an expected number of $O(\log^2 n)$ steps.¹

Following [9], several papers explored further the correlations between underlying structure and long-range connections that yield efficiently-searchable small-world networks. Specifically, general models in which the results of [9] carry over have been proposed [1, 4, 5, 10, 15, 16, 18], and various non-greedy decentralized search algorithms have been studied [7, 12, 14, 15]. We consider a general framework for studying search in small worlds, which naturally extends the model in [9], and captures many of the models proposed in the follow-up work. In this framework, we show that for *any* underlying structure, we can construct a small-world network where search takes $n^{o(1)}$ steps. Specifically, this framework models a small world by a random graph, where nodes are individuals, and edges denote acquaintance relationships, as in [9]. It consists of three components:

- *Base graph*: It is a connected undirected graph G on the set of individuals. For any nodes u, v of G , their shortest-path distance in G , denoted by $d_{u,v}$, represents the distance between individuals u and v in the social network. This distance may capture various notions of proximity, such as geographic proximity, and/or similarity in occupation, culture, hobbies, etc. The neighbors of u in G , i.e., all v such that $d_{u,v} = 1$, are u ’s acquaintances with a profile most similar to that of u , and they are called the *local contacts* of u .² In this paper we will assume that G can be arbitrary; so, no extra assumptions about the underlying structure of the social network are imposed.
- *Augmentation scheme*: It is a collection $\{\varphi_u : u \in V(G)\}$ of probability distributions over the set of nodes $V(G)$ of G . The graph G is *augmented* with additional edges as follows: from each node u , one directed edge is added to a node chosen independently at random according to φ_u . This node is an acquaintance of u with, possibly, a very different profile than u ’s, and it is called the *long-range contact* of u .³ We will use a simple and natural augmentation scheme, where φ_u depends on the “density” of G around u .
- *Search algorithm*: It is a decentralized algorithm that computes a path in the augmented graph from any source to any target, using *only* local information. More precisely, the decisions of the algorithm may depend on G and on the long-range contacts of the nodes visited so far, but not on the long-range contacts of nodes that have not been visited yet. We will consider a simple search algorithm, where the next hop is chosen based only on the current node’s distance to its neighbors (in G), and on *their* distance to the target. A prominent example of such an algorithm is *greedy search*, where the next hop is to the neighbor of the current node that is closest to the target (in G).

¹Note that paths of polylogarithmic length exist between nodes for a wide range of exponents α [16]; however, they can be efficiently discovered by a decentralized algorithm only when $\alpha = 2$.

²We could define the local contacts of u as all nodes with $d_{u,v} \leq k$, for some constant $k \geq 1$; however, this would not affect the asymptotic behavior of the model.

³Again, having a constant number $k \geq 1$ of long-range contacts per node would not essentially affect the model’s behavior.

As mentioned earlier, a significant amount of work has focused on extending the results of [9] to underlying structures other than the grid. In fact, there has been an effort to identify the most general setting in which the results of [9] carry over. One motivation for that has been the observation that certain measures of social proximity, such as occupation, are not successfully captured by a grid-based model. Moreover, the underlying structure of a social network often results from the combination of several different measures of proximity. Therefore, it is highly unclear how this underlying structure should look.

The results of [9] were extended by [10] into hierarchical models, where the nodes are the leaves of a complete b -ary tree, and the distance between them is the length of the path between them in the tree. Also, [10] proposed a model generalizing both the grid-based and the hierarchical models, where the distance function is induced by certain families of node sets. Small-world networks with grid-like underlying structure were studied in [15, 16], where the diameter of these networks was computed. In [4], small-world networks where the base graph has a “bounded growth rate” property were considered. In [5], base graphs of bounded tree-width were studied. The case where the nodes are embedded in a metric space of bounded doubling dimension (instead of a base graph) was considered in [18]. Finally, weighted base graphs that exclude a fixed minor were studied in [1]. In all these cases, for suitable augmentation schemes, greedy search performs in a polylogarithmic expected number of steps, as a function of the number of nodes. (In [1] and [18], the expected number of steps is also a polylogarithmic function of the aspect ratio, i.e., the ratio of the largest to the smallest distance between nodes.)

For the case of *arbitrary* base graphs, the following results are known. An infinite family of graphs was described in [8], such that for any n -node base graph from this family, and any augmentation scheme, greedy search requires an expected number of $\Omega(2^{\sqrt{\log n}})$ steps, for some source–target pair. On the upper-bound side, a universal augmentation scheme was proposed in [6], such that for any n -node base graph augmented according to this scheme, the expected number of steps for greedy search is $\tilde{O}(n^{1/3})$.

Several non-greedy decentralized search algorithms have been suggested for grid-based small-world networks [7, 12, 14, 15]. Roughly speaking, all these algorithms construct shorter paths than greedy search, but they visit (or “consult”) an additional, small number of nodes before the next hop is decided. Finally, there is work on lower-bounds on the performance of greedy search in grid-based small-world networks for arbitrary augmentation schemes (see, e.g., [2]).

For more background on search in small-world networks, see the survey [11].

1.1 Our contribution

In this paper, we show that for *any* base graph, there exists a simple, natural augmentation scheme, and a simple, natural search algorithm such that search in the augmented base graph takes $n^{o(1)}$ steps. Precisely, the expected number of steps for any source–target pair is at most $2^{(\log n)^{1/2 + o(1)}}$. This bound almost matches the lower bound of $\Omega(2^{\sqrt{\log n}})$ steps of [8], which applies to any search algorithm that decreases the distance to the target in each step (see Theorem 2, in Section 2.3).

The augmentation scheme we consider is the one where

the probability $\varphi_u(v)$ that node $v \neq u$ is the long-range contact of node u is inversely proportional to the size of the smallest ball in the base graph that is centered at u and contains v . This scheme is inspired by the experimental observation of [13] that two-thirds of the friendships are geographically distributed as follows: the probability of befriending a particular person is inversely proportional to the number of closer people. This augmentation scheme is also a natural generalization to arbitrary graphs of the augmentation scheme with $\alpha = 2$ proposed in [9] for the grid.

Our search algorithm is a slight variation of greedy search, where the next hop may be to a distant neighbor of the current node only if this hop results in sufficient progress towards the target. Quantitatively, from the current node u , search moves to the node that is closest to the target t , among all the neighbors v of u for which the ratio of the “benefit” $\beta(u, v, t) = d_{u,t} - d_{v,t}$ of moving to v , over the distance $d_{u,v}$, is larger than some small threshold:

$$\beta(u, v, t)/d_{u,v} \geq 1/\text{polylog } \beta(u, v, t).$$

(If this condition is replaced by $\beta(u, v, t)/d_{u,v} > 0$ then we obtain the standard greedy search.) This algorithm has a natural interpretation in the context of Milgram’s experiments [17]: people make greedy choices, but are reluctant to forward the letter to an acquaintance located far away, unless they are sure this acquaintance is sufficiently closer to the target. This reluctance may result from the fact that people can only *approximate* the distances between their acquaintances and the target, and, typically, the further the acquaintance resides, the worse this approximation is.

In the context of social networks, our result could be interpreted as follows: individuals may well be able to construct short chains between people *regardless of the underlying structure of the social network*.

1.2 Comparison with previous work

All previous work except [6, 8] has focused on achieving polylogarithmic search performance, but for restricted underlying structures. In contrast, we assume that the underlying structure is a graph, but we impose no restrictions on the structure of this graph. In this general setting, the lower bound of [8] precludes achieving polylogarithmic search. Still, search in an n -node network in $n^{o(1)}$ steps could also be regarded as “efficient search.” In [1], a *weighted* graph is used as a base graph. Our result can be extended to this case as well (see Section 3.3.5). In [18], a *metric space* of bounded doubling dimension is considered instead of a base graph. The concept of underlying metric is more general than that of base graph, and our result *cannot* be extended to *arbitrary* underlying metrics. Indeed, although a metric space can be seen as a weighted graph that induces a shortest-path metric, the edges of this graph are not necessarily present in the small-world network, unlike in the base graph approach. For our analysis to apply to an underlying metric, a condition that ensures that search can always make progress towards the target is necessary. In [18], this is achieved by combining the bounding doubling dimension condition, with a sufficiently large number of long-range contacts per node.

Our proof diverges significantly from the approach typically used in previous work. The standard approach has been to bound the number of steps until either the distance to the target is halved, or the size of the ball around the

target containing the current node is halved. The target is then reached in a logarithmic number of such halvings. For the underlying structures that have been studied in previous work, this halving can be achieved in a polylogarithmic number of steps, because, intuitively, for any target, the long-range links that facilitate progress towards this target are frequent enough. This approach, however, does not work for arbitrary base graphs, essentially, because the “dimensionality” of the graph may be so large that for some destinations the useful long-range links are very rare (see [8]). In our approach, instead of evaluating progress in terms of big jumps, we focus on making progress “locally.” Essentially, we recursively divide the search into smaller search subproblems, where in each subproblem either the distance between the source and the target is smaller, or the size of the part of the graph we focus on is smaller. To make this approach work we had to modify greedy search such that no long jumps with very small benefit are allowed, because such jumps could waste effort devoted to making progress locally, e.g., to overcoming a very dense part of the graph.

2. MODEL AND STATEMENT OF MAIN RESULT

In this section, we describe our augmentation scheme and search algorithm, and we state our main result.

2.1 Augmentation scheme

Recall that $d_{u,v}$ is the distance between the nodes u and v in the base graph G . We consider the augmentation scheme where the probability $\varphi_u(v)$ that v is the long-range contact of u is proportional to the *rank* of $d_{u,v}$ among the distances from u to all the other nodes; this scheme is called *rank-based augmentation*. Precisely, let $\mathbf{B}_u(r) = \{v \in V(G) : d_{u,v} \leq r\}$ be the ball of radius $\lfloor r \rfloor$ centered at u . ($\mathbf{B}_u(r) = \emptyset$, if $r < 0$.) In the *rank-based augmentation scheme*, for any nodes u, v with $v \neq u$,

$$\varphi_u(v) = \frac{1}{\nu_u \cdot |\mathbf{B}_u(d_{u,v})|}, \quad \text{where } \nu_u = \sum_{w \neq u} \frac{1}{|\mathbf{B}_u(d_{u,w})|}.$$

The normalizing factor ν_u is within the range:

$$1 - \frac{1}{n} \leq \nu_u \leq \sum_{i=2}^n \frac{1}{i} \leq \ln n. \quad (2.1)$$

The lower bound for ν_u follows from the fact that, for all $v \neq u$, $|\mathbf{B}_u(d_{u,v})| \leq n$. The upper bound follows from the fact that if $v_1, v_2, \dots, v_{n-1}$ is an ordering of all the nodes excluding u , such that $d_{u,v_1} \leq d_{u,v_2} \leq \dots \leq d_{u,v_{n-1}}$, then $|\mathbf{B}_u(d_{u,v_i})| \geq i + 1$.

2.2 Search algorithm

Let $\mathbf{N}_u$ denote the set of all neighbors of node u in the *augmented* graph, i.e., $\mathbf{N}_u$ consists of the neighbors of u in G , plus the long-range contact of u . In greedy search for target t , the next hop from node $u \neq t$ is to a node $v \in \mathbf{N}_u$ that *minimizes* the remaining distance $d_{v,t}$ to t , or, equivalently, that *maximizes* the *benefit*

$$\beta(u, v, t) = d_{u,t} - d_{v,t}$$

of moving to v . The search algorithm we consider in this paper, called *focused greedy*, is identical to greedy search, except that it ignores all nodes v for which the *relative benefit*

$\beta(u, v, t)/d_{u,v}$ is too small. Precisely, we define the *next-hop space* $\mathbf{D}_{u,t}$ of u for target t as

$$\mathbf{D}_{u,t} = \left\{ v \in V(G) : \frac{\beta(u, v, t)}{d_{u,v}} \geq \frac{1}{1 + \log \beta(u, v, t)} \right\}. \quad (2.2)$$

In *focused greedy search* for target t , if $u \neq t$ is the current node then the next hop is to a node $v \in \mathbf{N}_u \cap \mathbf{D}_{u,t}$ that minimizes the remaining distance $d_{v,t}$ to t . Note that $\mathbf{N}_u \cap \mathbf{D}_{u,t} \neq \emptyset$, because the first node after u in the shortest path in G from u to t is both in $\mathbf{N}_u$ and in $\mathbf{D}_{u,t}$. Thus, focused greedy search is well defined.

2.3 Main result

The main result of this paper is the following:

THEOREM 1 (UPPER BOUND). *For any base graph on n nodes augmented according to the rank-based augmentation scheme, the expected number of steps of focused greedy search from any source node to any target node is $\mathcal{O}(2^{(\log n)^{1/2+\epsilon}})$, for any fixed $\epsilon > 0$.*

The upper bound of Theorem 1 almost matches the lower bound of [8] repeated below.⁴

THEOREM 2 (LOWER BOUND [8]). *There exists an infinite family of graphs such that, for any n -node base graph in this family, any augmentation scheme, and any search algorithm that reduces the distance to the target in each step, there exists a source–target pair for which search requires an expected number of $\Omega(2^{\sqrt{\log n}})$ steps.*

A family of graphs as described in Theorem 2 is the one consisting of the $\sqrt{\log n}$ -dimensional grids where there are additional edges from each node to all its $2^{\sqrt{\log n}}$ diagonal neighbors. The source–target pairs for which search is hard consist of nodes that lie on the same diagonal.

3. ANALYSIS

We now present the proof of Theorem 1. We start with a sketch of this proof, in Section 3.1. In Section 3.2, we establish some properties of the next-hop space $\mathbf{D}_{u,t}$. The main proof is described in Section 3.3.

3.1 Proof sketch of Theorem 1

We begin with some properties of the next-hop space. $\mathbf{D}_{u,t}$ has a teardrop shape and is included in the ball $\mathbf{B}_t(d_{u,t} - 1)$, the analogue of $\mathbf{D}_{u,t}$ for greedy search (see Figure 1). Also, for any $v \in \mathbf{D}_{u,t}$, $\mathbf{D}_{v,t}$ is included in $\mathbf{D}_{u,t}$. Let $\varrho_u(v)$, for $v \in \mathbf{D}_{u,t}$, be the radius of the largest ball centered at v and included in $\mathbf{D}_{u,t}$; and let $r_u(v)$ be the radius of the smallest ball centered at t containing $\mathbf{D}_{v,t} \setminus \mathbf{B}_v(\varrho_u(v))$ (see Figure 3). We use ϱ_u to measure the progress of the search process. We have that for any $w \in \mathbf{D}_{v,t}$, $\varrho_u(w) - \varrho_u(v)$ is at least equal to some positive constant. In particular, if $w \in \mathbf{D}_{v,t} \cap \mathbf{B}_t(r_u(v))$ then $\varrho_u(w) - \varrho_u(v) \geq \varrho_u(v) / \text{polylog } \varrho_u(v)$. This last property is at the core of our analysis, and it depends critically on the definition of the next-hop space.

The proof of the theorem proceeds by decomposing the search path into subpaths of certain types, and bounding the number and lengths of these subpaths. The decomposition is

⁴This is what is actually proved in [8], although the authors only claim their result for greedy search.




Figure 1: Illustration of Lemma 1(a). The next-hop space $\mathbf{D}_{u,t}$ is included in the ball $\mathbf{B}_t(d_{u,t} - 1)$, and includes the ball $\mathbf{B}_t(d_u^*)$, where $d_u^* = d_{u,t} - \Theta(d_{u,t} / \log d_{u,t})$.

recursive, and it employs three types of subpaths: σ -paths, δ -paths, and π -paths. Subpaths of the last two types are also decomposed into smaller subpaths of the three types, while σ -paths are not decomposed further; so, the search path is eventually decomposed into a collection of σ -paths. Each σ -path P is associated with a *reference node* u such that if v is P 's start node then $v \in \mathbf{D}_{u,t}$, and with a *target-ball radius* $r \geq r_u(v)$. P ends when a node $\hat{v}$ is reached such that (1) $\hat{v} \in \mathbf{B}_t(r)$, or (2) for some node $\hat{u}$ with $\hat{v} \in \mathbf{D}_{\hat{u},t}$, the set $\mathbf{D}_{\hat{u},t} \setminus \mathbf{B}_t(r)$ is at least $\kappa \approx n^{1/\sqrt{\log n}}$ times smaller than $\mathbf{D}_{u,t} \setminus \mathbf{B}_t(r)$ — $\hat{u}$ will be the reference node of the next σ -path. Intuitively, if (1) occurs, search makes progress because ϱ_u increases, and if (2) occurs, because the size of the part of the base graph we focus on decreases. We show that for the rank-based augmentation scheme, the expected length of a σ -path is at most $\kappa \cdot \text{polylog } n$. Also, the decomposition of the search path essentially yields a total number of $(\log n)^{\mathcal{O}(\log \kappa)}$ σ -path—this result is based on the last property of the next-hop space mentioned above. The desired bound on the length of the search path then follows.

3.2 Properties of next-hop space

Throughout this section, we assume that $u \neq t$. Also, the shorthand notations d_v and $\mathbf{D}_v$ are used for $d_{v,t}$ and $\mathbf{D}_{v,t}$, respectively. In the same spirit, in the notation we introduce we do not make explicit the dependence on t .

Let $f : [1, \infty) \rightarrow [1, \infty)$, with

$$f(x) = x(1 + \log x).$$

We can rewrite the definition in Equation (2.2) as

$$\mathbf{D}_u = \{v \in V(G) : d_{u,v} \leq f(d_u - d_v)\}.$$

Lemma 1 below describes the shape of $\mathbf{D}_u$ (see also Figures 1 and 2). As mentioned earlier, $\mathbf{D}_u$ has the shape of a teardrop. The distance from t to the “top” and to the “bottom” of this teardrop is $d_u - 1$ and d_u^* , respectively, where d_u^* is the solution of the following equation in x :

$$d_u + x = f(d_u - x).$$




Figure 2: Illustration of Lemma 1(b). For any d with $d_u^* \leq d \leq d_{u,t} - 1$, $\mathbf{D}_{u,t}$ is included in the union of $\mathbf{B}_t(d)$ and $\mathbf{B}_u(f(d_{u,t} - d))$, and includes their intersection.

Solving this equation, yields

$$d_u^* = d_u - \Theta(d_u / \log d_u).$$

LEMMA 1.

- (a) $\mathbf{B}_t(d_u^*) \subseteq \mathbf{D}_u \subseteq \mathbf{B}_t(d_u - 1)$.
- (b) For all reals d with $d_u^* \leq d \leq d_u - 1$,
 $\mathbf{B}_t(d) \cap \mathbf{B}_u(f(d_u - d)) \subseteq \mathbf{D}_u \subseteq \mathbf{B}_t(d) \cup \mathbf{B}_u(f(d_u - d))$.
- (c) For all $v \in \mathbf{D}_u$, $\mathbf{D}_v \subseteq \mathbf{D}_u$.

PROOF. Part (a) is obtained as follows. If $v \in \mathbf{B}_t(d_u^*)$ then $d_v \leq d_u^*$, so,

$$f(d_u - d_v) \geq f(d_u - d_u^*) = d_u + d_u^* \geq d_u + d_v \geq d_{u,v},$$

where the first relation holds because f is non decreasing, and the second follows from the definition of d_u^* . Thus, $v \in \mathbf{D}_u$. So, $\mathbf{B}_t(d_u^*) \subseteq \mathbf{D}_u$. If $v \in \mathbf{D}_u$ then, by the definition of f , $d_u - d_v \geq 1$, so, $v \in \mathbf{B}_t(d_u - 1)$. Hence, $\mathbf{D}_u \subseteq \mathbf{B}_t(d_u - 1)$.

We now prove (b). If $v \in \mathbf{B}_t(d) \cap \mathbf{B}_u(f(d_u - d))$ then $d_v \leq d$ and $d_{u,v} \leq f(d_u - d)$. So, $f(d_u - d_v) \geq f(d_u - d) \geq d_{u,v}$, and, thus, $v \in \mathbf{D}_u$. Therefore, $\mathbf{B}_t(d) \cap \mathbf{B}_u(f(d_u - d)) \subseteq \mathbf{D}_u$. If $v \in \mathbf{D}_u \setminus \mathbf{B}_t(d)$ then $d_{u,v} \leq f(d_u - d_v)$, and $d_v > d$. So, $d_{u,v} \leq f(d_u - d)$, and, thus, $v \in \mathbf{B}_u(f(d_u - d))$. Therefore, $\mathbf{D}_u \subseteq \mathbf{B}_t(d) \cup \mathbf{B}_u(f(d_u - d))$.

Finally, for (c), we have that if $w \in \mathbf{D}_v$ then

$$d_{u,w} \leq d_{u,v} + d_{v,w} \leq f(d_u - d_v) + f(d_v - d_w) \leq f(d_u - d_w),$$

where the last relation holds because $f(x) + f(y) \leq f(x + y)$. So, $w \in \mathbf{D}_u$. Thus, $\mathbf{D}_v \subseteq \mathbf{D}_u$. $\square$

We now introduce a quantity that plays a key role in our analysis. For any $v \in \mathbf{D}_u$, we define $\varrho_u(v)$ to be the solution of the following equation in x :

$$d_{u,v} + x = f(d_u - d_v - x),$$

Below, we will write ϱ_v instead of $\varrho_u(v)$. Intuitively, ϱ_v is a conservative estimate of the distance of v from the “border” of $\mathbf{D}_u$ (see Figure 3). The next lemma formalizes this intuition.




Figure 3: $\varrho_u(v)$ is a lower bound on the radius of the largest ball centered at u and included in $\mathbf{D}_{u,t}$ (Lemma 2); $r_u(v)$ is such that $\mathbf{D}_{v,t}$ is included in the union of the balls $\mathbf{B}_v(\varrho_u(v))$ and $\mathbf{B}_t(r_u(v))$, and includes their intersection (Lemma 3).

LEMMA 2. For all $v \in \mathbf{D}_u$, (a) $0 \leq \varrho_v \leq d_u - d_v - 1$; and (b) $\mathbf{B}_v(\varrho_v) \subseteq \mathbf{D}_u$.

PROOF. Part (a) is obtained as follows. The function $g(x) = f(d_u - d_v - x) - x$ is strictly decreasing. Also $g(0) = f(d_u - d_v) \geq d_{u,v}$, and $g(d_u - d_v - 1) = f(1) - (d_u - d_v - 1) = 2 - (d_u - d_v) \leq 1 \leq d_{u,v}$. Therefore, the equation $g(x) = d_{u,v}$ has a unique solution $x = \varrho_v$, and $0 \leq \varrho_v \leq d_u - d_v - 1$.

For (b), we have that if $w \in \mathbf{B}_v(\varrho_v)$ then

$$d_{u,w} \leq d_{u,v} + \varrho_v = f(d_u - d_v - \varrho_v) \leq f(d_u - d_w),$$

so, $w \in \mathbf{D}_u$. Thus, $\mathbf{B}_v(\varrho_v) \subseteq \mathbf{D}_u$. $\square$

We also introduce another related quantity. For any $v \in \mathbf{D}_u$, we define

$$r_u(v) = \begin{cases} d_v - f^{-1}(\max\{1, \varrho_v\}), & \text{if } \varrho_v \leq d_v + d_v^*; \\ -1, & \text{if } \varrho_v > d_v + d_v^*, \end{cases}$$

where f^{-1} is the inverse function of f . We will write r_v to denote $r_u(v)$. Informally, r_v is the radius of the smallest ball centered at t containing all the nodes in $\mathbf{D}_u$ that are not in $\mathbf{B}_v(\varrho_v)$ (see Figure 3). Precisely, we have the following lemma.

LEMMA 3. For all $v \in \mathbf{D}_u$,

$$\mathbf{B}_t(r_v) \cap \mathbf{B}_v(\varrho_v) \subseteq \mathbf{D}_v \subseteq \mathbf{B}_t(r_v) \cup \mathbf{B}_v(\varrho_v).$$

PROOF. If $\varrho_v < 1$ then $r_v = d_v - 1$. So,

$$\mathbf{B}_t(r_v) \cap \mathbf{B}_v(\varrho_v) = \mathbf{B}_t(d_v - 1) \cap \{v\} = \emptyset \subseteq \mathbf{D}_v.$$

Also, $\mathbf{B}_t(r_v) \cup \mathbf{B}_v(\varrho_v) \supseteq \mathbf{B}_v(d_v - 1) \supseteq \mathbf{D}_v$, by Lemma 1(a).

If $\varrho_v > d_v + d_v^*$ then $r_v = -1$ and $\mathbf{B}_t(r_v) = \emptyset$. So, $\mathbf{B}_t(r_v) \cap \mathbf{B}_v(\varrho_v) = \emptyset \subseteq \mathbf{D}_v$. Also, by Lemma 1(b),

$$\begin{aligned} \mathbf{D}_v &\subseteq \mathbf{B}_t(d_v^*) \cup \mathbf{B}_v(f(d_v - d_v^*)) = \mathbf{B}_t(d_v^*) \cup \mathbf{B}_v(d_v + d_v^*) \\ &= \mathbf{B}_v(d_v + d_v^*) \subseteq \mathbf{B}_v(\varrho_v) = \mathbf{B}_t(r_v) \cup \mathbf{B}_v(\varrho_v). \end{aligned}$$

Finally, if $1 \leq \varrho_v \leq d_v + d_v^*$ then $d_v^* \leq r_v \leq d_v - 1$. So, by Lemma 1(b),

$$\mathbf{D}_v \supseteq \mathbf{B}_t(r_v) \cap \mathbf{B}_v(f(d_v - r_v)) = \mathbf{B}_t(r_v) \cap \mathbf{B}_v(\varrho_v),$$

and $\mathbf{D}_v \subseteq \mathbf{B}_t(r_v) \cup \mathbf{B}_v(f(d_v - r_v)) = \mathbf{B}_t(r_v) \cup \mathbf{B}_v(\varrho_v)$. $\square$

The next lemma, Lemma 4, is the main result of this section. Lemma 4(a) says that if $v \in \mathbf{D}_u$ and $w \in \mathbf{D}_v$ then the difference between ϱ_w and ϱ_v is lower-bounded by a positive constant. This implies that when searching for t from a node in $\mathbf{D}_u$, the value of ϱ is increased by at least that constant in each step. Lemma 4(b) improves this bound to $\varrho_w - \varrho_v \geq \varrho_v / \text{polylog } \varrho_v$, for the case where $w \in \mathbf{B}_t(r_v)$. Lemma 4(b) lies at the core of our analysis, and it depends critically on the fact that the threshold for $\beta(u, v, t)/d_{u,v}$ in Equation (2.2) is $1/\text{polylog } \beta(u, v, t)$.⁵

LEMMA 4. For all $v \in \mathbf{D}_u$ and $w \in \mathbf{D}_v$,

- (a) $\varrho_w \geq \varrho_v + 0.46$;
- (b) if $w \in \mathbf{B}_t(r_v)$ then $\varrho_w \geq \varrho_v + h(\varrho_v)$, for some function h with $h(x) = \Omega(x/\log^2 x)$, and $h(x) \geq 0.46$ for all $x \geq 0$.

PROOF. Let $\Delta = \varrho_w - \varrho_v$. For (a) we must show that $\Delta \geq 0.46$. If $\varrho_w \geq d_u - d_v$ then $\Delta \geq d_u - d_v - \varrho_v \geq 1$, by Lemma 2(a). So, below we assume that $\varrho_w < d_u - d_v$.

By the definitions of ϱ_w and ϱ_v ,

$$\Delta = f(d_u - d_w - \varrho_w) - f(d_u - d_v - \varrho_v) - d_{u,w} + d_{u,v}.$$

Also, $d_{u,w} \leq d_{u,v} + d_{v,w} \leq d_{u,v} + f(d_v - d_w)$, since $w \in \mathbf{D}_v$. So,

$$\Delta \geq f(d_u - d_w - \varrho_w) - f(d_u - d_v - \varrho_v) - f(d_v - d_w). \quad (3.1)$$

The difference $f(d_u - x - \varrho_w) - f(d_v - x)$ decreases when x increases, because $\varrho_w < d_u - d_v$ and f is convex. And, since $d_w \leq d_v - 1$, $f(d_u - d_w - \varrho_w) - f(d_v - d_w)$ is at least

$$f(d_u - d_v + 1 - \varrho_w) - f(1) = f(d_u - d_v - \varrho_v + 1 - \Delta) - 1.$$

Thus,

$$\Delta \geq f(d_u - d_v - \varrho_v + 1 - \Delta) - f(d_u - d_v - \varrho_v) - 1.$$

Suppose that $\Delta \leq 1$ (otherwise, $\Delta \geq 0.46$ holds). By a similar argument as above, we have that

$$f(d_u - d_v - \varrho_v + 1 - \Delta) - f(d_u - d_v - \varrho_v) \geq f(2 - \Delta) - 1,$$

since, by Lemma 2(a), $d_u - d_v - \varrho_v \geq 1$. Therefore, $\Delta \geq f(2 - \Delta) - 2$. Let $\alpha = 0.4696\dots$ be the solution of the equation $\alpha = f(2 - \alpha) - 2$. Then, we have $\Delta \geq \alpha$, because by assuming the opposite we have a contradiction:

$$\Delta < \alpha = f(2 - \alpha) - 2 < f(2 - \Delta) - 2 = \Delta.$$

We now proceed to prove (b). From (a), we have that $\Delta \geq 0.46$. So, we just have to show that $\Delta = \Omega(\varrho_v / \log^2 \varrho_v)$. Since this result is asymptotic in ϱ_v , we can assume that ϱ_v is larger than some constant. We will assume that $\varrho_v \geq 8$.

Let $z = d_u - d_v - \varrho_v$. By the definition of ϱ_v ,

$$z = f^{-1}(d_{u,v} + \varrho_v) \geq f^{-1}(\varrho_v) = \Omega(\varrho_v / \log \varrho_v) \quad (3.2)$$

Hence, if $\varrho_w \geq d_u - d_v$ then $\Delta \geq z = \Omega(\varrho_v / \log^2 \varrho_v)$. So, below we assume that $\varrho_w < d_u - d_v$.

Inequality (3.1) still applies in the context of (b). Again we observe that $f(d_u - x - \varrho_w) - f(d_v - x)$ decreases as x increases. Also, since $w \in \mathbf{B}_t(r_v)$, $d_w \leq r_v = d_v - f^{-1}(\varrho_v)$ (since $\varrho_v \geq 8 > 1$). Thus, $f(d_u - d_w - \varrho_w) - f(d_v - d_w) \geq$

⁵In particular, a threshold equal to $1/(\beta(u, v, t))^c$, for any constant $c > 0$, would yield $\varrho_w - \varrho_v \geq \varrho_v^{1-\Theta(1)}$; and a constant threshold would yield $\varrho_w - \varrho_v = \Omega(1)$.

$f(d_u - d_w + f^{-1}(\varrho_v) - \varrho_w) - \varrho_v$. Applying this to (3.1), we obtain

$$\Delta \geq f(z + f^{-1}(\varrho_v) - \Delta) - f(z) - \varrho_v. \quad (3.3)$$

Next we bound $f(z + f^{-1}(\varrho_v) - \Delta)$. We will use the following simple fact.

FACT 1.

- (a) For all $x, y \geq 1$, $f(x + y) \geq f(x) + f(y) + \min\{x, y\}$.
- (b) For all x, y with $x \geq 2y$ and $y \geq 1$, $f(x - y) \geq f(x) - f(y) - y(\log \frac{x}{y} + 2.2)$.

PROOF. For (a), suppose, w.l.o.g., that $x \geq y$. Then,

$$\begin{aligned} f(x + y) &= (x + y)(1 + \log(x + y)) \\ &\geq x(1 + \log x) + y(1 + \log(2y)) = f(x) + f(y) + y. \end{aligned}$$

For (b), we have

$$f(x) - f(y) - f(x - y) = x \log \frac{x}{x - y} + y \log \frac{x - y}{y}.$$

Also, $\log \frac{x - y}{y} \leq \log \frac{x}{y}$, and

$$\begin{aligned} \log \frac{x}{x - y} &= \log \frac{1}{1 - y/x} \leq \log \frac{1}{e^{-y/x - (y/x)^2}} \\ &= (y/x + (y/x)^2) \log e \leq \frac{3y}{2x} \log e \leq 2.2 \frac{y}{x}, \end{aligned}$$

where for the second and the second-to-last relations we used the fact that $x \geq 2y$. Combining the above yields the desired inequality. $\square$ {of Fact 1}

To apply the above result, we will need the condition

$$\Delta \leq \min\{f^{-1}(\varrho_v) - 1, f^{-1}(\varrho_v)/2\}. \quad (3.4)$$

Note, however, that if the above inequality does not hold then $\Delta = \Omega(f^{-1}(\varrho_v)) \subseteq \Omega(\varrho_v / \log^2 \varrho_v)$. Also, we will have to make sure that $\Delta \geq 1$. This inequality, however, follows from our assumption that $\varrho_v \geq 8$, as we explain now. If $\Delta < 1$ then (3.3) yields

$$1 \geq f(z + f^{-1}(\varrho_v) - 1) - f(z) - \varrho_v.$$

Also, by (3.2), $z \geq f^{-1}(\varrho_v)$. Thus,

$$1 \geq f(2f^{-1}(\varrho_v) - 1) - 2f(f^{-1}(\varrho_v)).$$

Solving the above inequality numerically, yields $f^{-1}(\varrho_v) < 3$, which gives $\varrho_v < 8$, contradicting the assumption that $\varrho_v \geq 8$. Thus, $\Delta \geq 1$.

By Fact 1(a), we have

$$f(z + f^{-1}(\varrho_v) - \Delta) \geq f(z) + f(f^{-1}(\varrho_v) - \Delta) + f^{-1}(\varrho_v) - \Delta,$$

because, by (3.2), $z \geq f^{-1}(\varrho_v) - \Delta$, and, by (3.4), $f^{-1}(\varrho_v) - \Delta \geq 1$. Also, by Fact 1(b),

$$f(f^{-1}(\varrho_v) - \Delta) \geq \varrho_v - f(\Delta) - \Delta \left(\log \frac{f^{-1}(\varrho_v)}{\Delta} + 2.2 \right),$$

because, by (3.4), $f^{-1}(\varrho_v) \geq 2\Delta$, and $\Delta \geq 1$. Combining the two results above with (3.3), yields

$$\Delta \left(\log \frac{f^{-1}(\varrho_v)}{\Delta} + 4.2 \right) \geq f^{-1}(\varrho_v) - f(\Delta),$$

and solving for Δ we obtain that $\Delta \geq \frac{f^{-1}(\varrho_v)}{\log f^{-1}(\varrho_v) + 5.2} = \Theta\left(\frac{\varrho_v}{\log^2 \varrho_v}\right)$. $\square$

3.3 Proof of Theorem 1

The proof describes a recursive decomposition of the search path into subpaths of certain types, and it establishes upper bounds on the number of those paths and on their lengths. In Section 3.3.1, we define the types of subpaths employed in the decomposition. The actual decomposition is described in Section 3.3.2. In Section 3.3.3, we bound the expected length of the base subpath type. In Section 3.3.4, we combine all the pieces to derive the theorem. We finish with some remarks, in Section 3.3.5.

3.3.1 Subpath types

We define three types of subpaths of a search path: σ -paths, δ -paths, and π -paths. A path P of any of the three types is identified by:

- a *starting node* v ;
- a *reference node* u such that $\mathbf{D}_u$ contains all the nodes of P ; $\mathbf{D}_u$ is called the *reference set* of P ;⁶ and
- an integer $r \geq 0$; the ball $\mathbf{B}_t(r)$ is the *target ball* of P .

The triple (v, u, r) satisfies some *pre-condition* that depends on the type of P . There is also a *post-condition*, depending on the type of P , that describes the last node of P . P is the *shortest* subpath of the search path from v to t such that the first node of P is v , and its last node satisfies the post-condition. Below, we described the pre- and post-condition for each path type.

The σ -path with parameters v, u, r is denoted $\sigma(v, u, r)$, and the pre-condition that v, u, r must satisfy is

$$\text{PRE}_\sigma(v, u, r) : (v \in \mathbf{D}_u) \wedge (r_u(v) \leq r).$$

This condition implies that every node in $\mathbf{D}_v$ is in $\mathbf{B}_v(\varrho_u(v))$ or in the target ball $\mathbf{B}_t(r)$ (see Lemma 3). The following post-condition determines the last node $\hat{v}$ of $\sigma(v, u, r)$. For any node w , and any real d , define $N(w, d) = |\mathbf{D}_w \setminus \mathbf{B}_t(d)|$.

$$\text{POST}_\sigma(v, u, r, \hat{v}) : \exists \hat{u} : (\hat{v} \in \mathbf{D}_{\hat{u}}) \wedge (N(\hat{u}, r) \leq N(u, r)/\kappa),$$

where $\kappa = 2^{\sqrt{\log n \log \log n}}$. Informally, $\text{POST}_\sigma(v, u, r, \hat{v})$ says that there is a smaller reference set $\mathbf{D}_{\hat{u}}$ containing $\hat{v}$, such that the size of its fraction that is outside the target ball is at least κ times smaller than the size of the corresponding fraction of the original reference set $\mathbf{D}_u$. Note that it is not required that $\hat{u}$ belong to the search path from v to t . We will denote by $\partial_\sigma(v, u, r)$ the node $\hat{u}$ described in $\text{POST}_\sigma(v, u, r, \hat{v})$.⁷

The δ -path with parameters v, u, r is denoted $\delta(v, u, r)$, and its pre-condition is

$$\text{PRE}_\delta(v, u, r) : (v \in \mathbf{D}_u) \wedge (r_u(v) > r).$$

The post-condition specifying the last node $\hat{v}$ of $\delta(v, u, r)$ is

$$\text{POST}_\delta(v, u, r, \hat{v}) : (\varrho_u(\hat{v}) \geq \varrho_u(v) + h(\varrho_u(v)) \vee (r_u(\hat{v}) \leq r),$$

where h is the function defined in the statement of Lemma 4. Informally, $\text{POST}_\delta(v, u, r, \hat{v})$ says that either the distance ϱ_u from the border of the reference set is increased sufficiently, or $\text{PRE}_\sigma(\hat{v}, u, r)$ is satisfied.

⁶As before, $\mathbf{D}_v$ and d_v are shorthand notations for $\mathbf{D}_{v,t}$ and $d_{v,t}$, respectively.

⁷If there are more such $\hat{u}$, one of them is selected deterministically.

The π -path with parameters v, u, r is denoted $\pi(v, u, r)$. Its pre-condition $\text{PRE}_\pi(v, u, r)$ is the same as $\text{PRE}_\sigma(v, u, r)$, and the post-condition specifying its last node $\hat{v}$ is

$$\text{POST}_\pi(v, u, r, \hat{v}) :$$

$$\exists \hat{u} : (\hat{v} \in \mathbf{D}_{\hat{u}}) \wedge (N(\hat{u}, r) \leq N(u, r)/\kappa) \wedge (r_{\hat{u}}(\hat{v}) \leq r),$$

that is, the conjunction of $\text{POST}_\sigma(v, u, r, \hat{v})$ and $\text{PRE}_\sigma(\hat{v}, \hat{u}, r)$. So, the triple $(\hat{v}, \hat{u}, r)$ satisfies the same conditions as (v, u, r) , but $N(\hat{u}, r) \leq N(u, r)/\kappa$. We will denote by $\partial_\pi(v, u, r)$ the node $\hat{u}$ described in $\text{POST}_\pi(v, u, r, \hat{v})$.

The case of $u = \perp$. We extend the definitions of σ -paths and π -paths to the case where the reference set consists of all the nodes. Let $\mathbf{D}_\perp$ denote the set of all nodes, and let $\rho_\perp(v) = \infty$ and $r_\perp(v) = -1$, for any v . We define the paths $\sigma(v, \perp, r)$ and $\pi(v, \perp, r)$ using exactly the same definitions as for $\sigma(v, u, r)$ and $\pi(v, u, r)$, respectively, except that all occurrences of u in these definitions are replaced by $\perp$. E.g., the pre-condition for both $\sigma(v, \perp, r)$ and $\pi(v, \perp, r)$ is $(v \in \mathbf{D}_\perp) \wedge (r \geq r_\perp(v))$, which holds for any v and r .

3.3.2 Search-path decomposition

We now describe a decomposition of the search path into subpaths of the types above. The decomposition is defined recursively in Lemmata 5–7. Lemma 5 decomposes a π -path into a σ -path and a set of δ -paths; Lemma 6 decomposes a δ -path into a set of π -paths; and Lemma 7 decomposes a search path into a set of π -paths as well. In the first two lemmata, the union of the paths to which the initial path is decomposed may yield a larger path than the initial one. This is not a problem, however, since these results are used to obtain upper bounds on the length of the paths decomposed. Unlike δ -paths and π -paths, σ -paths are not decomposed into smaller paths; we bound the expected length of σ -paths in the next section.

Below, we denote the last node of path P by $\text{last}(P)$. Also, we assume that the starting node v is different than the target t of the search.

LEMMA 5. $\pi(v, u, r)$ is a subpath of $\bigcup_{i=0}^I P_i$, where $P_0 = \sigma(v, u, r)$, for $i > 0$, $P_i = \delta(v_i, \hat{u}, r)$ with $v_i = \text{last}(P_{i-1})$ and $\hat{u} = \partial_\sigma(v, u, r)$, and $I = \min\{i : r_{\hat{u}}(v_{i+1}) \leq r\}$. Also, $I \leq c \log^3 d_v$, for some constant c .

PROOF. First we show by induction that for all $i \leq I$, the pre-condition of P_i holds, thus, P_i is well defined. Clearly, the pre-condition $\text{PRE}_\sigma(v, u, r)$ of P_0 holds, since it is the same as $\text{PRE}_\pi(v, u, r)$. Let $0 < i \leq I$, and assume that the pre-condition of P_j holds, for all $j < i$. By the post-condition $\text{POST}_\sigma(v, u, r, v_1)$ of P_0 , $v_1 \in \mathbf{D}_{\hat{u}}$, so, by Lemma 1(c), $v_i \in \mathbf{D}_{\hat{u}}$ as well. Also, $r_{\hat{u}}(v_i) > r$, since $I \geq i$. Therefore, the pre-condition $\text{PRE}_\delta(v_i, \hat{u}, r)$ of P_i holds.

Next, we prove the bound on I . If $0 < i < I$, by combining the post-condition $\text{POST}_\delta(v_i, \hat{u}, r, v_{i+1})$ of P_i and the fact that $r_{\hat{u}}(v_{i+1}) > r$ (since $I > i$), we obtain that $\varrho_{\hat{u}}(v_{i+1}) \geq \varrho_{\hat{u}}(v_i) + h(\varrho_{\hat{u}}(v_i))$. From this and the definition of h , it follows that for appropriate constants i_0 and $c' > 0$, if $i_0 \leq i < I$ then

$$\varrho_{\hat{u}}(v_{i+1}) \geq (1 + c' / \log^2 \varrho_{\hat{u}}(v_i))^{i-i_0}.$$

Also, since $r_{\hat{u}}(v_i) > r \geq 0$,

$$\varrho_{\hat{u}}(v_i) \leq d_{v_i} + d_{v_i}^* \leq 2d_{v_i} \leq 2d_v.$$

From the two inequalities above, it follows easily that $I \leq c \log^3 d_v$, for some constant c .⁸

It remains to show that $\pi(v, u, r)$ is a subpath of $\bigcup_{i=0}^I P_i$. By the post-condition $\text{POST}_\sigma(v, u, r, v_1)$ of P_0 , $v_1 \in \mathbf{D}_{\hat{u}}$, so, $v_{I+1} \in \mathbf{D}_{\hat{u}}$. From the same post-condition, we have that $N(\hat{u}, r) \leq N(u, r)/\kappa$. Also, $r_{\hat{u}}(v_{I+1}) \leq r$, by the definition of I . Therefore, $\text{POST}_\pi(v, u, r, v_{I+1})$ holds, which implies that $\pi(v, u, r)$ is a subpath of $\bigcup_{i=0}^I P_i$.⁹ $\square$

LEMMA 6. $\delta(v, u, r)$ is a subpath of $\bigcup_{i=0}^I \pi(v_i, u_i, \hat{r})$, where $v_0 = v$, $u_0 = u$, $\hat{r} = r_u(v)$, for $i > 0$, $v_i = \text{last}(\pi(v_{i-1}, u_{i-1}, \hat{r}))$ and $u_i = \partial_\pi(v_{i-1}, u_{i-1}, \hat{r})$, and $I = \min\{i : v_{i+1} \in \mathbf{B}_t(\hat{r})\}$. Also, $I \leq \log_\kappa \max\{1, N(u, \hat{r})\}$.

PROOF. It is similar to the proof of Lemma 5. First we show that for all $i \leq I$, the pre-condition of $P_i = \pi(v_i, u_i, \hat{r})$ holds. The pre-condition $\text{PRE}_\pi(v, u, \hat{r})$ of P_0 holds because, by $\text{PRE}_\delta(v, u, r)$, $v \in \mathbf{D}_u$, and $\hat{r} = r_u(v)$. Also, if $0 < i \leq I$, the pre-condition $\text{PRE}_\pi(v_i, u_i, \hat{r})$ of P_i follows from the post-condition $\text{POST}_\pi(v_{i-1}, u_{i-1}, \hat{r}, v_i)$ of P_{i-1} .

Next, we derive the bound on I . Combining the post-conditions $\text{POST}_\pi(v_i, u_i, \hat{r}, v_{i+1})$ of P_i , for all $i < I$, yields $N(u_I, \hat{r}) \leq N(u, \hat{r})/\kappa^I$. Also, $N(u_I, \hat{r}) \geq 1$, since, otherwise, $v_I \in \mathbf{B}_t(\hat{r})$, which contradicts the definition of I . Therefore, $N(u, \hat{r})/\kappa^I \geq 1$, and, so, $I \leq \log_\kappa \max\{1, N(u, \hat{r})\}$.

Finally, by the definition of I , $v_{I+1} \in \mathbf{B}_t(r_u(v))$, and, by applying Lemma 4(b), we obtain $\varrho_u(v_{I+1}) \geq \varrho_u(v) + h(\varrho_u(v))$. Hence, $\text{POST}_\delta(v, u, r, v_{I+1})$ holds, and, so, $\delta(v, u, r)$ is a subpath of $\bigcup_{i=0}^I P_i$. $\square$

The proof of the next lemma is completely analogous to that of Lemma 6 and is omitted.

LEMMA 7. The search path from node v to t is equal to $\bigcup_{i=0}^I \pi(v_i, u_i, 0)$, where $v_0 = v$, $u_0 = \perp$, for $i > 0$, $v_i = \text{last}(\pi(v_{i-1}, u_{i-1}, 0))$ and $u_i = \partial_\pi(v_{i-1}, u_{i-1}, 0)$, and $I = \min\{i : v_{i+1} = t\}$. Also, $I \leq \log_\kappa n$.

3.3.3 Expected length of σ -path

In this section, we establish the following upper bound on the expected length of a σ -path. We denote the length of a path P by $\|P\|$.

LEMMA 8. $\mathbb{E}[\|\sigma(v, u, r)\|] = O(\kappa \log n)$.

PROOF. We show that if $\|\sigma(v, u, r)\|$ is larger than i then, with probability at least $1/\kappa \log n$, it is no larger than $i + 2$. Formally, let $L = \|\sigma(v, u, r)\|$. We show that for all $i \geq 0$,

$$\mathbb{P}\mathbb{r}[L \leq i + 2 \mid L > i] \geq 1/\kappa \log n.$$

From this it is immediate that $\mathbb{E}[L] = O(\kappa \log n)$. We now prove the above inequality.

Let the path $\sigma(v, u, r)$ be $v_0 v_1 \cdots v_L$. Suppose that $i < L$, and fix the value of v_i , say $v_i = w$. Let $\mathbf{A} = \mathbf{B}_w(x)$ for the *smallest* radius x such that $\mathbf{D}_w \subseteq \mathbf{B}_t(r) \cup \mathbf{B}_w(x)$. We describe two properties of the set $\mathbf{A}$ that we need for the proof. The first property is

$$\mathbf{A} \subseteq \mathbf{D}_u. \quad (3.5)$$

⁸Note that this bound depends critically on function h , which is determined by the shape of the next-hop space; see also the footnote before the statement of Lemma 4.

⁹Note that, in general, $\pi(v, u, r) \neq \bigcup_{i=0}^I P_i$, since it is possible that $\partial_\pi(v, u, r) \neq \hat{u} = \partial_\sigma(v, u, r)$.

We can derive this inclusion as follows. Using the definitions of $r_u(w)$ and $r_u(v)$, and the inequalities $d_w \leq d_v$, $d_w^* \leq d_v^*$, and $\varrho_u(w) \geq \varrho_u(v)$, where the last one follows from Lemma 4(a), it is easy to show that $r_u(w) \leq r_u(v)$. So, by $\text{PRE}_\sigma(v, u, r)$, $r_u(w) \leq r$. And since, by Lemma 3, $\mathbf{D}_w \subseteq \mathbf{B}_t(r_u(w)) \cup \mathbf{B}_w(\varrho_u(w))$, $\mathbf{D}_w \subseteq \mathbf{B}_t(r) \cup \mathbf{B}_w(\varrho_u(w))$. Combining this and the definition of $\mathbf{A}$, yields $\mathbf{A} \subseteq \mathbf{B}_w(\varrho_u(w))$. And since, by Lemma 2(b), $\mathbf{B}_w(\varrho_u(w)) \subseteq \mathbf{D}_u$, (3.5) follows. The second property of $\mathbf{A}$ that we need is

$$\mathbf{B}_t(r) \cap \mathbf{A} \subseteq \mathbf{D}_w. \quad (3.6)$$

This can be shown by applying Lemmata 1(a) and (b) (for $\mathbf{D}_w$ and $d = r$), and using the definition of $\mathbf{A}$.

For the rest of the proof, we distinguish two cases. First we consider the case of $N(w, r) \leq |\mathbf{A} \setminus \mathbf{B}_t(r)|/\kappa$. By (3.5),

$$N(w, r) \leq |\mathbf{D}_u \setminus \mathbf{B}_t(r)|/\kappa = N(u, r)/\kappa,$$

and since $v_{i+1} \in \mathbf{D}_w$, $\text{POST}_\sigma(v, u, r, v_{i+1})$ holds. Therefore,

$$\mathbb{P}\mathbb{r}[L = i + 1 \mid (L > i) \wedge (v_i = w)] = 1.$$

We consider now the complementary case of $N(w, r) > |\mathbf{A} \setminus \mathbf{B}_t(r)|/\kappa$. Let r' be the *smallest* radius so that

$$|\mathbf{D}_w \cap \mathbf{B}_t(r') \setminus \mathbf{B}_t(r)| > |\mathbf{A} \setminus \mathbf{B}_t(r)|/\kappa.$$

Clearly, $r' > r$. Then,

$$\begin{aligned} \mathbb{P}\mathbb{r}[v_{i+1} \in \mathbf{B}_t(r') \mid (L > i) \wedge (v_i = w)] &\geq \mathbb{P}\mathbb{r}[v_{i+1} \in \mathbf{B}_t(r') \cap \mathbf{A} \mid (L > i) \wedge (v_i = w)] \\ &\geq \sum_{v' \in \mathbf{B}_t(r') \cap \mathbf{A} \cap \mathbf{D}_w} \varphi_w(v') \\ &\geq \frac{|\mathbf{B}_t(r') \cap \mathbf{A} \cap \mathbf{D}_w|}{\ln n \cdot |\mathbf{A}|} \\ &\geq \frac{|\mathbf{B}_t(r') \cap \mathbf{A} \cap \mathbf{D}_w| - |\mathbf{B}_t(r) \cap \mathbf{A} \cap \mathbf{D}_w|}{\ln n \cdot (|\mathbf{A}| - |\mathbf{B}_t(r) \cap \mathbf{A} \cap \mathbf{D}_w|)}, \end{aligned}$$

where the second-to-last line is obtained by applying the definition of φ and (2.1), and the last by using the fact that $\frac{a}{b} \geq \frac{a-c}{b-c}$, for $0 \leq c < a \leq b$. It is

$$\begin{aligned} |\mathbf{B}_t(r') \cap \mathbf{A} \cap \mathbf{D}_w| - |\mathbf{B}_t(r) \cap \mathbf{A} \cap \mathbf{D}_w| &= |\mathbf{B}_t(r') \cap \mathbf{A} \cap \mathbf{D}_w \setminus \mathbf{B}_t(r)| \\ &= |\mathbf{B}_t(r') \cap \mathbf{D}_w \setminus \mathbf{B}_t(r)|, \end{aligned}$$

where the first relation holds because $r' \geq r$, and the second because, by the definition of $\mathbf{A}$, $\mathbf{A} \supseteq \mathbf{D}_w \setminus \mathbf{B}_t(r)$. Also,

$$|\mathbf{A}| - |\mathbf{B}_t(r) \cap \mathbf{A} \cap \mathbf{D}_w| = |\mathbf{A}| - |\mathbf{B}_t(r) \cap \mathbf{A}| = |\mathbf{A} \setminus \mathbf{B}_t(r)|,$$

where the first relation holds because of (3.6). Combining the above gives

$$\begin{aligned} \mathbb{P}\mathbb{r}[v_{i+1} \in \mathbf{B}_t(r') \mid (L > i) \wedge (v_i = w)] &\geq \frac{|\mathbf{D}_w \cap \mathbf{B}_t(r') \setminus \mathbf{B}_t(r)|}{\ln n \cdot |\mathbf{A} \setminus \mathbf{B}_t(r)|} \\ &\geq \frac{1}{\kappa \ln n}, \end{aligned} \quad (3.7)$$

by the definition of r' . If $v_{i+1} \in \mathbf{B}_t(r')$ then

$$\begin{aligned} N(v_{i+1}, r) &\leq |\mathbf{D}_w \cap \mathbf{B}_t(r' - 1) \setminus \mathbf{B}_t(r)| \\ &\leq |\mathbf{A} \setminus \mathbf{B}_t(r)|/\kappa \\ &\leq N(u, r)/\kappa, \end{aligned}$$

where the first inequality holds because $\mathbf{D}_{v_{i+1}} \subseteq \mathbf{D}_w \cap \mathbf{B}_t(r'-1)$, the second because of the minimality of r' , and the last because of (3.5). So, if $v_{i+1} \in \mathbf{B}_t(r')$, $\text{POST}_\sigma(v, u, r, v_{i+2})$ holds, and, thus, $L \leq i+2$. Combining this and (3.7) yields

$$\Pr[L \leq i+2 \mid (L > i) \wedge (v_i = w)] \geq 1/\kappa \ln n. \quad \square$$

3.3.4 Expected length of search path

We are now ready to bound the expected length of the search path. Let $T_\sigma = \max_{v,u,r} \mathbb{E}[\|\sigma(v, u, r)\|]$, and, for $m \geq 1$, let

$$\begin{aligned} T_\delta(m) &= \max\{\mathbb{E}[\|\delta(v, u, r)\|] : N(u, r) \leq m\}; \\ T_\pi(m) &= \max\{\mathbb{E}[\|\pi(v, u, r)\|] : N(u, r) \leq m\}. \end{aligned}$$

Define also $T_\delta(0) = T_\pi(0) = 0$. By Lemma 5, we have

$$T_\pi(m) \leq T_\sigma + c \log^3 n \cdot T_\delta(\lfloor m/\kappa \rfloor).$$

Also, by Lemma 6 and the fact that $\hat{r} > r$,

$$T_\delta(m) \leq (1 + \log_\kappa m) \cdot T_\pi(m).$$

Combining yields

$$T_\pi(m) \leq T_\sigma + c \log^3 n \log_\kappa m \cdot T_\pi(\lfloor m/\kappa \rfloor),$$

and unfolding the recurrence we obtain

$$\begin{aligned} T_\pi(m) &\leq T_\sigma \cdot \sum_{0 \leq i \leq \log_\kappa m} (c \log^3 n \log_\kappa m)^i \\ &= O(T_\sigma \cdot (c \log^3 n \log_\kappa m)^{\log_\kappa m}). \end{aligned}$$

By Lemma 7, the expected length T of the search path from s to t is

$$T \leq (1 + \log_\kappa n) \cdot T_\pi(n),$$

and by Lemma 8, $T_\sigma = O(\kappa \log n)$. Combining gives

$$T = O(\log_\kappa n \cdot \kappa \log n \cdot (c \log^3 n \log_\kappa m)^{\log_\kappa n}).$$

Finally, substituting the value of $\kappa = 2^{\sqrt{\log n \log \log n}}$, we obtain $T = O(2^{c' \sqrt{\log n \log \log n}})$, for some constant $c' > 0$. Therefore, $T = o(2^{(\log n)^{1/2+\epsilon}})$, for any fixed $\epsilon > 0$. This completes the proof of Theorem 1.

3.3.5 Remarks

The bound we have shown above holds also with high probability (not only in expectation); i.e., with high probability, $2^{O(\sqrt{\log n \log \log n})}$ steps are required, for all source–target pairs. This follows from the fact that, with high probability, $\|\sigma(v, u, r)\| = O(\kappa \log n)$, which is immediate from the proof of Lemma 8. Our result can be extended to the case of arbitrary weighted base graphs. For this case, the bound is $2^{O(\sqrt{\log n \log \log(n+\Delta)})}$ steps, where Δ is the aspect ratio, i.e., the ratio of the largest to the smallest distance between nodes.

4. OPEN PROBLEMS

We have established an $n^{o(1)}$ bound on the expected number of steps of focused greedy search, a variant of greedy search, for any source–target pair, in any n -node base graph that has been augmented via some natural density-based distribution. This upper bound is essentially tight.

It would be interesting to determine if our bound holds for greedy search as well. In greedy search, the next-hop

space of node u for target t consists of all nodes v with $\beta(u, v, t) = d_{u,t} - d_{v,t} > 0$, while in focused greedy, the next-hop space consists of those v with $\beta(u, v, t)/d_{u,v} \geq 1/\text{polylog}(\beta(u, v, t))$. Our proof works only for this specific threshold for $\beta(u, v, t)/d_{u,v}$. We do not know if this is intrinsically required for achieving search paths of length $n^{o(1)}$, or if it is just an artifact of our proof.

The lower bound of Theorem 2 holds for the worst-case source–target pair, but not necessarily when the source–target pair is chosen at random. It may thus be the case that for a random pair, even greedy search takes a polylogarithmic expected number of steps, for any base graph (for a suitable augmentation). It is an open problem whether this is true.

Note also that Theorem 2 applies only to search algorithms that move closer to the target in each step. An interesting question is whether there is a combination of a decentralized search algorithm (which sometimes moves further away from the target) and a “natural” augmentation scheme, such as the rank-based augmentation, that achieves search paths of polylogarithmic expected length, for any base graph.

5. ACKNOWLEDGMENTS

We thank the anonymous referees of this paper for their helpful comments.

6. REFERENCES

- [1] I. Abraham and C. Gavoille. Object location using path separators. In *Proc. 25th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 188–197, 2006.
- [2] M. Dietzfelbinger and P. Woelfel. Tight lower bounds for greedy routing in uniform small world rings. In *Proc. 41st ACM Symp. on Theory of Computing (STOC)*, pages 591–600, 2009.
- [3] P. S. Dodds, R. Muhamad, and D. J. Watts. An experimental study of search in global social networks. *Science*, 301:827–829, 2003.
- [4] P. Duchon, N. Hanusse, E. Lebhar, and N. Schabanel. Could any graph be turned into a small-world? In *Proc. 19th Int. Symp. on Distributed Computing (DISC)*, pages 511–513, 2005.
- [5] P. Fraigniaud. Greedy routing in tree-decomposed graphs. In *Proc. 13th European Symp. on Algorithms (ESA)*, pages 791–802, 2005.
- [6] P. Fraigniaud, C. Gavoille, A. Kosowski, E. Lebhar, and Z. Lotker. Universal augmentation schemes for network navigability: Overcoming the $\sqrt{n}$ -barrier. In *Proc. 19th ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*, pages 1–7, 2007.
- [7] P. Fraigniaud, C. Gavoille, and C. Paul. Eclecticism shrinks even small worlds. In *Proc. 23rd ACM Symp. on Principles of Distributed Computing (PODC)*, pages 169–178, 2004.
- [8] P. Fraigniaud, E. Lebhar, and Z. Lotker. A lower bound for network navigability. *SIAM J. Discrete Math.*, to appear. (Preliminary version in *Proc. 13th ESA*, 2006.).
- [9] J. Kleinberg. The small-world phenomenon: An algorithm perspective. In *Proc. 32nd ACM Symp. on Theory of Computing (STOC)*, pages 163–170, 2000.

- [10] J. Kleinberg. Small-world phenomena and the dynamics of information. In *Proc. 15th Neural Information Processing Systems Conf. (NIPS)*, pages 431–438, 2001.
- [11] J. Kleinberg. Complex networks and decentralized search algorithms. In *Proc. Int. Congress of Mathematicians (ICM)*, 2006.
- [12] E. Lebar and N. Schabanel. Almost optimal decentralized routing in long-range contact networks. In *Proc. 31st Int. Colloq. on Automata, Languages and Programming (ICALP)*, pages 894–905, 2004.
- [13] D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, and A. Tomkins. Geographic routing in social networks. *Proc. National Academy of Sciences of the USA*, 102(33):11623–11628, 2005.
- [14] G. S. Manku, M. Naor, and U. Wieder. Know thy neighbor’s neighbor: The power of lookahead in randomized P2P networks. In *Proc. 36th ACM Symp. on Theory of Computing (STOC)*, pages 54–63, 2004.
- [15] C. Martel and V. Nguyen. Analyzing Kleinberg’s (and other) small-world models. In *Proc. 23rd ACM Symp. on Principles of Distributed Computing (PODC)*, pages 179–188, 2004.
- [16] C. Martel and V. Nguyen. Analyzing and characterizing small-world graphs. In *Proc. 16th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 311–320, 2005.
- [17] S. Milgram. The small world problem. *Psychology Today*, 67(1):60–67, 1967.
- [18] A. Slivkins. Distance estimation and object location via rings of neighbors. In *Proc. 24th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 41–50, 2005.
- [19] D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393:440–442, 1998.