

Decentralized Search with Random Costs

Oskar Sandberg *

March 31, 2008

Abstract

A decentralized search algorithm is a method of routing on a random graph that uses only limited, local, information about the realization of the graph. In some random graph models it is possible to define such algorithms which produce short paths when routing from any vertex to any other, while for others it is not.

We consider random graphs with random costs assigned to the edges. In this situation, we use the methods of stochastic dynamic programming to create a decentralized search method which attempts to minimize the total cost, rather than the number of steps, of each path. We show that it succeeds in doing so among all decentralized search algorithms which monotonically approach the destination. Our algorithm depends on knowing the expected cost of routing from every vertex to any other, but we show that this may be calculated iteratively, and in practice can be easily estimated from the cost of previous routes and compressed into a small routing table. The methods applied here can also be applied directly in other situations, such as efficient searching in graphs with varying vertex degrees.

1 Introduction

Jon Kleinberg introduced the concept of decentralized search algorithms in his celebrated 2000 paper on the Small-World phenomenon [1]. In particular, he showed that in certain random graphs it is possible to find paths between vertices of poly-logarithmic length even when limited to using only local knowledge at each step, while in others it is not.

Most of the by now large canon of work in the area (see [2] for a recent survey) has been dedicated to finding and analyzing algorithms that route between two given vertices in a small number of steps. Typically, the best

*The Department of Mathematical Sciences, Chalmers University of Technology and Göteborg University. ossa@math.chalmers.se




Figure 1: If the goal is to reach vertex 9 in the fewest steps, then certainly vertex 1 should choose the long edge to vertex 7. However, if the goal instead is to minimize the total cost, routing to 2 might be a better choice.

method in these situation is greedy: progress to the neighbor which is closest to the destination. In this paper, we consider a generalized situation where the cost of passing down an edge is not fixed, but may be a random variable whose value is known when the choice of where to proceed to next is made. The goal is then the minimize the cost of reaching the destination, which may lead to different priorities when routing (one may not wish to route to a vertex very close to the destination if the cost of passing down that edge is very high). This is a situation similar to that addressed by the field of stochastic dynamic programming [3], and we use similar methods.

The basic idea is this: If a vertex knows the expected cost of each of his neighbors routing to the destination, and also the cost of him routing to each of his neighbors, it makes sense for him to choose as the next step that neighbor which minimizes the sum of these two costs. While the expected costs are difficult to calculate analytically, we find that there is little reason to do so. When many queries are performed, one may start by using any guess as to these values, and then update these guesses based on past experience.

We will show analytically that this form of search is well defined, and that it is optimal among algorithms that monotonically approach the destination, as well as presenting some results on the order of the total cost as well as approximation trade-offs. We continue with a discussion and some experimentation on the practicality of the approach. Finally, we apply the algorithm with greater generality, and see that it performs well also in these cases. In particular, we note that it can be applied to routing with non-homogeneous degree distributions.

1.1 Previous Work

The original work on decentralized search was done by Kleinberg in [1] and [4]. Much work has since been done on related problems, in particular further generalization of the results and improvements on the bounds (see [5] [6] [7] [8] [9] for some examples).

In a bid to improve the performance of Freenet [10], a decentralized peer-to-peer network, Clarke [11] proposed an algorithm in some ways similar to that presented here under the name “NG Routing”. The method was implemented in Freenet at the time, but has not been used since the network was re-engineered to route according to the method described in [12]. We believe that the technical and architectural problems experienced by previous versions of Freenet were unrelated to what is discussed below, and, in light of our results, that Clarke’s ideas were fundamentally sound.

Şimşek and Jensen also proposed an algorithm based on the same principles in [13], intended for routing based on both vertex degree and similarity. Their “expected-value navigation” is based on the same idea as our cost-greedy search, however they use a rough estimate of the expected routing time, which cannot be applied to our problem, to make the decisions. They present no analytic results. For our take on search in graphs with variable node degree, see Section 6.2 below.

For an introduction to the field of stochastic dynamic programming, see the monograph by Ross [3].

1.1.1 A note on terminology

The terminology regarding algorithms for decentralized path-finding in random graphs has not yet settled, and different authors have used different terms. We have chosen the term “decentralized search” (following recent work by Kleinberg [2]) but others terms have been used to describe the same thing. “Navigation”, used in e.g. [4] [14], is quite common, but perhaps not as descriptive. We avoid formally calling our algorithms “routings” as this name has previously been used in computer science literature to describe flow assignments through a graph with limited edge capacities (for example [15]), a different problem from that currently studied. However, we do use the terms “route” and “to route” following their dictionary definitions.

1.2 Organization

In Section 2 we set out the basic definitions of decentralized search, as well as rigorously define our new “cost-greedy” search algorithm. Following this, in Section 3 we prove the main results regarding cost-greedy search in networks with sufficient independence. In Sections 4 and 5 we set out the methods for applying the results in practice, and, finally, in Section 6 we perform simulated experiments to look at the actual performance of the algorithm.

2 Definitions

A cost graph $G = (V, E, C)$ is a graph consisting of a vertex set V , a possibly directed set of edges E , and a collection of costs, C . For each element $(x, y) \in E$, $C(x, y) = C((x, y))$ is an i.i.d. positive random variable giving the cost of traveling down that edge (time taken).

G may be a random graph, by which we mean that for each $x, y \in V$, there exists a random indicator variable $E_{x,y}$ saying whether there is an edge from x to y . These may be dependent and differently distributed.

Definition 2.1. *For a given cost graph G , a z -search for a vertex $z \in V$, is a mapping $\mathcal{A} : V \mapsto V$ such that:*

1. $\mathcal{A}(x) = y$ only if $(x, y) \in E$.
2. $\mathcal{A}(z) = z$.
3. For all $x \in V$ there exists is $k < \infty$ such that $\mathcal{A}^k(x) = z$.

A search of G is a collection of z -routing algorithms for all $z \in V$.

We call d a *distance* on a set V if $d : V \times V \mapsto \mathbb{R}^+$, if for $x, y \in V$, $d(x, y) = 0$ implies that $x = y$, and if for $z, y, z \in V$,

$$d(x, z) \leq d(x, y) + d(y, z).$$

A distance is thus a metric without the symmetry requirement. In particular, any connected digraph G implies a distance d_G . For $x \in V$, $N(x) = \{y \in V : d_G(x, y) = 1\}$ is the set of neighbors of x in G .

Definition 2.2. *A distance d is adapted for search in a connected graph G if for every $x, z \in V$, where $x \neq z$, there exists a $y \in N(x)$ such that $d(y, z) < d(x, z)$.*

A distance function d is thus adapted for search if it, in some sense, reflects the structure of G . The most obvious example is of course d_G itself, but d may also be, for instance, graph distance on any connected spanning sub-graph H of G . Another important case is that if V is a set of points in a metric space, then the space's metric is adapted for search in V 's Delaunay triangulation (see [9]).

Definition 2.3. *Given a cost graph G , and a vertex z , a decentralized search is a z -search \mathcal{A} , such that for any $x \in V$ the random variable $\mathcal{A}(x)$ is measurable with respect to:*

1. $E_{x,y}$ for all $y \in V$.

2. $C(x, y)$ for all $y \in N(x)$.

Intuitively, this means that as well as any information about the graph model, routing at x may use information about which vertices x does (and does not) have edges to, as well as the costs of passing down those edges. The definition of decentralized search as originally given by Kleinberg was slightly broader than this, allowing a route started at a vertex x to use all the information from $\mathcal{A}^0(x), \mathcal{A}^1(x), \dots, \mathcal{A}^{k-1}(x)$ when taking its k -th step. Because our analysis will be restricted to algorithms meeting the following criteria, excluding this information makes little difference:

Definition 2.4. Given a graph G , a distance d adapted for search in G , and a vertex z , a forward search is a z -search \mathcal{F} such that for all $x \in V \setminus \{z\}$

$$d(\mathcal{F}(x), z) < d(x, z).$$

For a given search \mathcal{A} and vertex z

$$S^z(x; \mathcal{A}) = \inf \left(k : \mathcal{A}^k(x) = z \right)$$

is the number of *steps* it takes to reach z from x using \mathcal{A} . Let

$$T^z(x; \mathcal{A}) = \sum_{i=1}^{S^z(x; \mathcal{A})} C(\mathcal{A}^{i-1}(x), \mathcal{A}^i(x))$$

which is the *cost*.

2.1 Greedy and Weighted Greedy Search

In the following, we fix z . *Greedy search* is given by

$$\mathcal{A}_G(x) = \operatorname{argmin}_{y \in N(x)} (d(y, z)). \quad (1)$$

This is always a decentralized search, and if d is adapted, then it is a forward search (if d is not adapted, it may not be well defined to begin with).

Standard greedy search does not take the costs C into account. A variant that does, is *weighted greedy search*. For a given z , let $w_z(x)$ for $x \in V$ be a collection of weights, where $w_z(z) = 0$. These weights specify a search algorithm

$$\mathcal{A}_w(x) = \operatorname{argmin}_{y \in N(x)} (C(x, y) + w_z(x)) \quad (2)$$

which may or may not be well defined. If we restrict this to being a forward search, we get

$$\mathcal{F}_w(x) = \operatorname{argmin}_{y \in N(x) : d(y, z) < d(x, z)} (C(x, y) + w_z(x)) \quad (3)$$

which is always a well-defined if d is adapted for routing in G .

The behavior of weighted routing depends on w . If $w(x)$ is strictly increasing in $d(x, z)$ then \mathcal{A}_w and \mathcal{F}_w are \mathcal{A}_G when the costs are constant. If $w(x) = 0$ for all $x \in V$, then \mathcal{F}_w will simply choose the edge approaching z with the lowest cost. Since we are trying to minimize the cost of routing, it makes sense to see $w_z(x)$ as a guess of the cost of reaching z from x . Imagine that we are given a black-box function f , where

$$f(x, \mathcal{A}) = \mathbf{E}[T^z(x; \mathcal{A})]$$

that is, f tells us the expected cost of routing from x using any given algorithm \mathcal{A} . If a vertex x knows that all its neighbors will use the algorithm \mathcal{A} to route, it makes sense that it should want to use the y given by (2) (or (3) if restricted to forward search) with $w_z(y) = f(y, \mathcal{A})$. Extending this reasoning to any vertex motivates the following definition.

Definition 2.5. *A forward search \mathcal{F} is called cost-greedy search for z if $\mathcal{F} = \mathcal{F}_{\tilde{w}}$, where the weights \tilde{w} are given by the solution to the equations:*

$$w(x) = \mathbf{E}[T^z(x; \mathcal{F}_w)] \quad x \in V \tag{4}$$

Equation 4 is what is the routing equivalent of what is known in stochastic dynamic programming texts as the *optimality equation*. We will show below that a solution exists in this context, that this solution is a globally attractive fix-point, and that cost-greedy search is optimal with respect to expected routing cost for all forward searches.

3 Results for Independent Graphs

Let d be an distance, and $y \rightsquigarrow x$ denote the event that there is an edge from y to x (we allow multiple edges per pair). We let G be a random graph model on which d is adapted for routing (with probability 1), and for which $y, x, v, u \in V$ with $d(y, z) > d(x, z)$ and $d(v, z) > d(u, z)$, $y \rightsquigarrow x$ is independent of $v \rightsquigarrow u$ if $y \neq v$. Examples of such graphs are adding outgoing edges from each vertex with destinations chosen independently (as in Kleinberg's work [1]) or allowing each edge, either seen as directed or undirected, to exist independently of all others (like in classical random graphs and long-range percolation [16]). We call a graph constructed in this manner *edge independent*.

Theorem 3.1. *For an edge independent random graph G , there exists a solution \tilde{w} to equation (4) so that cost-greedy search is well defined. \tilde{w} is a globally attractive fix-point of the iteration given by*

$$w_{i+1} = \mathbf{E}[T^z(x; \mathcal{F}_{w_i})]. \tag{5}$$

Proof. Define the rank of x with respect to z , $r_z(x)$, as the position of x when all the elements in $V \setminus \{z\}$ are ordered by increasing distance from z , using some deterministic tie-breaking rule. Let $r_z(z) = 0$.

We will proceed by induction on $r_z(x)$.

Let w_0 be any weighting. Let $x \in V$ such that $r_z(x) = 1$. For any forward search \mathcal{F} , $\mathcal{F}(x) = z$, whence $\mathbf{E}[T^z(x; \mathcal{F})]$ does not depend on the algorithm. In particular $\mathbf{E}[T^z(x; \mathcal{F}_w)]$ does not depend on w , whence $w_i(x)$ is constant for $i \geq 1$.

Let $r_z(x) = k$, and assume that for all $y \in V$ such that $r_z(y) < r_z(x)$, $w_i(y)$ takes the same value for all $i \geq k - 1$. This means that $\mathbf{E}[T^z(x; \mathcal{F}_{w_i})]$ takes the same value for all $i \geq k$.

It follows that for all $x \in V$ $w_i(x)$ is fixed for $i \geq r_z(x)$. Hence $w_k = w_{k+1}$ for all $k \geq n - 1$, and $\tilde{w} = w_n$ is a solution to (4). \square

Theorem 3.2. *Let G be an edge independent graph, and \mathcal{F} a forward search for a vertex $z \in V$. Then for all $x \in V$*

$$\mathbf{E}[T^z(x; \mathcal{F}_{\tilde{w}})] \leq \mathbf{E}[T^z(x; \mathcal{F})]$$

where $\mathcal{F}_{\tilde{w}}$ is cost-greedy search for z as in Definition 2.5.

Proof. Use the same definition of $r_z(x)$ as in the proof of Theorem 3.1. Like there, we will use induction on $r_z(x)$.

If $r_z(x) \leq 1$, then all forward searches from x are the same, and there is nothing to prove. Let \mathcal{F} be any forward search. Given $x \in V$, assume that for all $y \in V$ such that $r_z(y) < r_z(x)$, $\mathbf{E}[T^z(y; \mathcal{F}_{\tilde{w}})] \leq \mathbf{E}[T^z(y; \mathcal{F})]$.

Let $v = \mathcal{F}(x)$ and $v^* = \mathcal{F}_T(x)$, the places the respective algorithms choose as the next step. Below, we mean by ‘‘local knowledge’’ that which decentralized algorithm may use, as given in 2.3. We note, crucially, that because of our assumptions, $T^z(v; \mathcal{F})$ is independent of local knowledge at x , while $C(x, y)$ is measurable with respect to it.

$$\begin{aligned} \mathbf{E}[T^z(x; \mathcal{F})] &= \mathbf{E}[\mathbf{E}[T^z(x; \mathcal{F}) \mid \text{local knowledge at } x]] \\ &= \mathbf{E}[D(x, v) + \mathbf{E}[T^z(v; \mathcal{F})]] \\ &\geq \mathbf{E}[D(x, v) + \mathbf{E}[T^z(v; \mathcal{F}_{\tilde{w}})]] \\ &\geq \mathbf{E}[D(x, v^*) + \mathbf{E}[T^z(v^*; \mathcal{F}_{\tilde{w}})]] = \mathbf{E}[T^z(x; \mathcal{F}_{\tilde{w}})] \end{aligned}$$

where the first inequality follows by induction since $r_z(v) < r_z(x)$ and the last because the expression inside the first expectation is what v^* minimizes. \square

3.1 The Small-World Graph

A particular example graphs meeting the criteria of the last chapter are the small-world augmentations first introduced by Kleinberg [1]. This construction starts with a fixed finite grid H , letting $d = d_H$, and creating G by adding a random outgoing directed edge from each vertex x to destination y with probability

$$\mathbf{P}(x \rightsquigarrow y) \propto 1/d_G(x, y)^\alpha \quad (6)$$

d is naturally adapted for routing in G .

For simplicity, we let H be a ring of n vertices (Kleinberg originally used a two-dimensional square lattice, but the proofs are identical). Let $\mathcal{G}(n, \alpha)$ be the family of random graphs so constructed.

Using previous results about greedy routing on such graphs, we can calculate the cost order of cost-greedy search. In particular, we can see the order in n cannot be different from greedy routing.

Theorem 3.3. (Kleinberg) *If $G \in \mathcal{G}(n, \alpha)$ with $\alpha = 1$ there exists N_1 such that for $n \geq N_1$,*

$$\mathbf{E}[S^z(x; \mathcal{A}_G)] \leq k_1 \log n \log d(x, z)$$

where k_1 is a constant independent of x, z , and n .

Further results about such graphs, proved in [5] and [17] respectively are

Theorem 3.4. (Barriere et al.) *If $G \in \mathcal{G}(n, \alpha)$ with $\alpha = 1$, then there exists N_2 such that for $n \geq N_2$,*

$$\mathbf{E}[S^z(x; \mathcal{A}_G)] \geq k_2 \log n \log d(x, z)$$

where k_2 is a constant independent of x, z , and n .

Theorem 3.5. (Singh Manku) *If $G \in \mathcal{G}(n, \alpha)$, with $\alpha \geq 0$, then*

$$\mathbf{E}[S^z(x; \mathcal{A}_G)] \leq \mathbf{E}[S^z(x; \mathcal{A})]$$

for any $x, z \in V(G)$ and decentralized search \mathcal{A} .

Together, these allow us to prove the observation that

Proposition 3.6. *If $G \in \mathcal{G}(n, \alpha)$ and $0 < \mathbf{E}[C(x, y)] < \infty$ then*

$$\mathbf{E}[T^z(x; \mathcal{F}_{\bar{w}})] = \Theta(\log n \log d(x, z))$$

Proof. The upper bound comes directly from Theorems 3.2 and 3.3, since

$$\begin{aligned}\mathbf{E}[T^z(x; \mathcal{F}_{\bar{w}})] &\leq \mathbf{E}[T^z(x; \mathcal{A}_G)] \\ &= \mathbf{E}[C(x, y)]\mathbf{E}[S^z(x; \mathcal{A}_G)] \\ &\leq \mathbf{E}[C(x, y)]k_1 \log n \log d(x, y).\end{aligned}$$

where the middle equality follows from the fact \mathcal{A}_G routes independently of the costs, and the simple form of Wald's equation.

To prove the upper bound, consider all edges in the graph as directed, letting the edges of the H be denoted by double directed edges. Since a forward search can only ever traverse an edge in one direction, this does not affect its cost. Now let

$$R^z(x; \mathcal{F}) = \sum_{i=1}^{S^z(x; \mathcal{A})} C_{\min}(\mathcal{F}(x)^{i-1}) \quad (7)$$

where $C_{\min}(x) = \min\{C(x, y) : y \in N(x)\}$. This counts, at each step, the minimum cost of any outgoing edge, rather than the cost of the edge which was actually used.

Since the degree of each vertex is fixed, $C_{\min}(x)$ is i.i.d. for all x . Let \mathcal{S}_i be the σ -algebra generated by all the information seen in steps $1, 2, \dots, i$ of the search (as listed in Definition 2.3). Note:

- $C_{\min}(\mathcal{F}^{i+1}(x))$ is independent of \mathcal{S}_i .
- $S_z(x; \mathcal{A})$ is a Stopping Time with respect to the filtration $\{\mathcal{S}_i\}_{i=1}^{\infty}$.

Thus we may use Wald's Equation to conclude that

$$\mathbf{E}[R^z(x; \mathcal{F})] = \mathbf{E}[C_{\min}(x)]\mathbf{E}[S^z(x; \mathcal{F})]. \quad (8)$$

We now use the immediate fact that $R^z(x; \mathcal{F}) \leq T^z(x; \mathcal{F})$, followed by (8) and Theorems 3.5 and 3.4, to conclude

$$\begin{aligned}\mathbf{E}[T^z(x; \mathcal{F})] &\geq \mathbf{E}[R^z(x; \mathcal{F})] \\ &= \mathbf{E}[C_{\min}(x)]\mathbf{E}[S^z(x; \mathcal{F})] \\ &\geq \mathbf{E}[C_{\min}(x)]\mathbf{E}[S^z(x; \mathcal{A}_G)] \\ &\geq \mathbf{E}[C_{\min}(x)]k_2 \log n \log d(x, z)\end{aligned}$$

for sufficiently large n . Since this holds for any forward-search \mathcal{F} , it holds in particular for $\mathcal{F}_{\bar{w}}$. \square

Proposition 3.6 tells us that in this model, the order of cost-greedy routing will not be different from that of greedy routing. The proof of the lower bound assumes, however, that $\mathbf{E}[C(x, y)] < \infty$ and that the degree of each vertex is bounded as n grows. Neither of these things, and particularly not the latter, necessarily hold in applications.

3.2 Approximated Weights

We consider the situation when the solution \tilde{w} to (5) is not known exactly but approximated by another set of weights.

Proposition 3.7. *If \tilde{w} is the solution to (5) and w another set of positive weights such that*

$$\max_{x \in V} |w(x) - \tilde{w}(x)| \leq \epsilon$$

then for any edge-independent graph of size n

$$\mathbf{E}[T^z(x; \mathcal{F}_w)] - \mathbf{E}[T^z(x; \mathcal{F}_{\tilde{w}})] \leq 2n\epsilon \quad (9)$$

and more generally, for any $k \geq 0$

$$\mathbf{E}[T^z(x; \mathcal{F}_w)] - \mathbf{E}[T^z(x; \mathcal{F}_{\tilde{w}})] \leq 2\epsilon(k + n\mathbf{P}(S^z(x; \mathcal{F}_w) > k)). \quad (10)$$

Proof. Let $\text{err}_w(x) = \mathbf{E}[T^z(x; \mathcal{F}_w)] - \mathbf{E}[T^z(x; \mathcal{F}_{\tilde{w}})]$. It follows that

$$\begin{aligned} \mathbf{E}[T^z(x; \mathcal{F}_w)] &= \mathbf{E}[C(x, \mathcal{F}_w(x)) + \mathbf{E}[T^z(\mathcal{F}_w(x); \mathcal{F}_w)]] \\ &= \mathbf{E}[C(x, \mathcal{F}_w(x)) + \mathbf{E}[T^z(\mathcal{F}_w(x); \mathcal{F}_{\tilde{w}})]] + \text{err}_w(\mathcal{F}_w(x)) \\ &= \mathbf{E}[C(x, \mathcal{F}_w(x)) + \tilde{w}(\mathcal{F}_w(x))] + \text{err}_w(\mathcal{F}_w(x)). \end{aligned}$$

Now, since by the definition of a weighted greedy search $C(x, \mathcal{F}_w(x)) + w(\mathcal{F}_w(x)) \leq C(x, y) + w(y)$ for all $y \in N(x)$

$$\begin{aligned} &\leq \mathbf{E}[C(x, \mathcal{F}_{\tilde{w}}(x)) + w(\mathcal{F}_{\tilde{w}}(x)) - w(\mathcal{F}_w(x)) + \tilde{w}(\mathcal{F}_w(x))] \\ &\quad + \text{err}_w(\mathcal{F}_w(x)) \\ &\leq \mathbf{E}[C(x, \mathcal{F}_{\tilde{w}}) + \tilde{w}(\mathcal{F}_{\tilde{w}}(x))] + \mathbf{E}|w(\mathcal{F}_{\tilde{w}}(x)) - \tilde{w}(\mathcal{F}_{\tilde{w}}(x))| \\ &\quad + \mathbf{E}|\tilde{w}(\mathcal{F}_w(x)) - w(\mathcal{F}_w(x))| + \text{err}_w(\mathcal{F}_w(x)) \\ &\leq \mathbf{E}[T^z(x; \mathcal{F}_{\tilde{w}})] + 2\epsilon + \text{err}_w(\mathcal{F}_w(x)) \end{aligned}$$

It follows that for any $k \geq 0$

$$\text{err}_w(x) \leq 2\epsilon k + \text{err}_w(\mathcal{F}_w^k(x)). \quad (11)$$

If $k > S^z(x; \mathcal{F}_w)$ then $\mathcal{F}_w^k(x) = z$ and $\text{err}_w(\mathcal{F}_w^k(x)) = 0$, so (9) follows since $n > S^z(x; \mathcal{A})$ for all searches.

To prove (10), note that by the same reasoning

$$\text{err}_w(\mathcal{F}_w^k(x)) = \text{err}_w(\mathcal{F}_w^k(x) \mid \mathcal{F}_w^k(x) \neq z) \mathbf{P}(S^z(x; \mathcal{F}_w) > k).$$

Since the graph is edge independent $\text{err}_w(\mathcal{F}_w^k(x) \mid \mathcal{F}_w^k(x) \neq z)$ is simply the error from some point which is not z , but where (9) still applies. \square

What the proposition, and in particular (11) says is that if an approximation w of \tilde{w} is off by ϵ , then each step in the routing adds at most 2ϵ to the optimal routing time. This is intuitively clear, since while \mathcal{F}_w may choose the wrong vertex in a given step, it can only do so when the total (actual) cost of routing via that vertex is within 2ϵ of the cost of routing via the real one. For the same reason, it is unlikely that a better bound can be achieved without further assumptions on the graph and the cost distribution.

4 Calculating the Weights

Theorem 3.1 provides us with a method of calculating the weights \tilde{w} for cost-greedy search. One can start by assigning any initial weighting w_0 , and then calculate w_1, w_2, \dots using (5).

A closed analytic form for $\mathbf{E}[T^z(x; \mathcal{F}_w)]$ as a function of the vector w is probably very difficult to find, even in the most simple situations. One can note however that it can be written recursively as

$$\mathbf{E}[T^z(x; \mathcal{F}_w)] = \sum (\mathbf{E}[T^z(x; \mathcal{F}_w)] + \mathbf{E}[C(x, y) | \mathcal{F}_w(x) = y]) \mathbf{P}(\mathcal{F}_w(x) = y)$$

where the sum is over all $y \in V$ such that $d(y, z) < d(x, z)$. In the very simplest cases (such as a directed loop with one augmented outgoing shortcut chord per vertex) it is possible to calculate $\mathbf{P}(\mathcal{F}_w(x) = y)$ and $\mathbf{E}[C(x, y) | \mathcal{F}_w(x) = y]$ analytically, in which case $\mathbf{E}[T^z(x; \mathcal{F}_w)]$ can be calculated numerically by recursion. Because this is complicated, and unlikely to be of much interest in practice, we do not linger on it.

A much more rewarding strategy is to calculate the weights empirically. That is, start by simulating a large number of searches from randomly chosen points using \mathcal{F}_{w_0} . While this is being done, sample the average routing cost to z from each vertex (due to the Markovian nature of forward search on an edge independent graph, a vertex may take a sample every time a query passes through it). After a sufficient number of queries, the average should be an estimate at $\mathbf{E}[T^z(x; \mathcal{F}_{w_0})]$ by the law of large numbers. One may then take the average costs from each point as w_1 , and continue in this manner. Proposition 3.7 indicates how close an approximation is needed, but unfortunately it is not strong enough to derive a rigorous bound using a polynomial number of samples.

Further, we note two things about the sampling implemented. Firstly, one needs to be careful about the way the repeated queries are done. Since we want the edge costs to be random, $C(x, y)$ must be picked anew, independently, for each query sampled. If the graph is random, the edges may be redrawn, but must not – it simply depends on whether they are to be seen as random or fixed edges in the G above.

Secondly, the proof of Theorem 3.1 guarantees convergence in n steps, meaning that an optimal routing is achieved once w_n has been calculated (if an empirical method is used, the resulting weighting may still suffer inaccuracies due to the sampling). This is an unfortunately large number of iterations, especially given that each may require simulating a large number of queries, but we find that in practice, much fewer iterations (typically two or three, even for very large networks) are needed, see Section 6 below.

5 Practicality and Decentralization

We proceed to discuss actual applications of Definition 2.5. On the face of it, the routing method described does not seem particularly practical. Even if we can calculate the weighting \tilde{w} , this gives a routing table of size of n , and such a table is needed for each z we wish to route for. The complete table of weights needed to route between any two vertices is thus of size n^2 .

Several assumptions can help here however.

Centralization: Translation Invariance

If we assume that the graph is translation invariant, then $\mathbf{E}[T^z(x; \mathcal{F})] = \mathbf{E}[T^0(x-z; \mathcal{F})]$ so x needs only know the routing cost from each starting vertex to a distinguished vertex 0. In fact, in many cases (such as the common case of augmenting single cycle with random outgoing edges) $\mathbf{E}[T^z(x; \mathcal{F}_T)]$ may be exactly, or at least approximately, a function of $d(x, z)$, in which case x need only know the expected cost of routing a given distance. This knowledge is the same for all x , so may be calculated as a single, global, vector.

Decentralization

If one wishes for a completely decentralized search system, as, for instance in peer-to-peer systems such as [10], then one cannot store a global vector of weights. Instead, each vertex must store the weights needed to route to every other vertex. In particular, each vertex x needs to be able to calculate $w_z(y)$ for each $y \in N(x)$ and $z \in V$. If one assumes translation invariance, x need only store one such weight vector, and can translate it to apply to its neighbors. Without such invariance, it needs to store $|N(x)|$ vectors.

Weight Vector Compression

In both cases above, however, we are still left with a routing table size of at least n , which is definitely not desirable. The heart of what makes our




Figure 2: Rather than storing every value of w vector (the dotted line) as a routing table, we store values at exponentially increasing positions, and use these to approximate the values between them.

method practically useful comes from the fact that the previous theory about decentralized search makes compression to a logarithmic size possible.

If we consider graphs of type $\mathcal{G}(n, \alpha)$ described in Section 3.1, we know from Proposition 3.6 that

$$\mathbf{E}[T^z(x, \mathcal{F}_{\tilde{w}})] \approx c\mathbf{E}[C] \log(n) \log(d(x, z)).$$

The utility of this is that if we know that $\mathbf{E}[T^z(x, \mathcal{F}_T)]$ grows logarithmically with $d(x, z)$ (as indicated by Proposition 3.6), we are motivated to assume that it, and thus the weights \tilde{w} in Definition 2.5, can be approximated by assuming $\tilde{w}(x)$ and $\tilde{w}(y)$ have similar values if x and y are such that $\log(d(x, z)) \approx \log(d(y, z))$. In particular if $0 \leq d(y, z) - d(x, z) \leq r$ we get

$$|\tilde{w}(y) - \tilde{w}(x)| \approx c_1 \mathbf{E}[C] \log(n) \frac{r}{d(x, z)}. \quad (12)$$

It is easy to prove, using the same methods as in the proof of Theorem 3.3 that for greedy routing in $\mathcal{G}(n, \alpha)$

$$\mathbf{P}(S^z(x; \mathcal{A}_G) \geq \log^3 n) \leq c_2 \frac{\log n}{n}.$$

Assuming that a similar bound holds for $\mathcal{F}_{\tilde{w}}$, equation (10) with $k = \log^3 n$ in Proposition 3.7, gives that

$$\mathbf{E}[T^z(x; \mathcal{F}_w)] - \mathbf{E}[T^z(x; \mathcal{F}_{\tilde{w}})] \leq 4 \max_{x \in V} |\tilde{w}(x) - w(x)| \log^3 n$$

so if $r < \epsilon d(x, z) / 4 \log^3 n$ in (12), then

$$\mathbf{E}[T^z(x; \mathcal{F}_w)] - \mathbf{E}[T^z(x; \mathcal{F}_{\tilde{w}})] < \epsilon.$$

Thus, the weight $w(x)$ can be substituted by the weight of a vertex $\epsilon d(x, z) / \log^3 n$ steps from x . To do this, we divide the routing distances into zones of size 2^i for $i = 0, 1, 2, \dots$, and record only the weights of $\log^3 n / \epsilon$ evenly spaced vertices within each zone (Figure 2). The routing table thus contains a




Figure 3: The steps taken by greedy and cost-greedy search when all edge costs are fixed to 1. In theory, the latter should converge to the former (which is optimal) but due to the inaccuracy of the estimates and the routing-table compression cost-greedy here performs about 5% worse than the optimum.

polylogarithmic number of entries ($O(\log^4 n)$), and yet by using the closest recorded weight as a substitute for $w(x)$, we incur only an ϵ -error on the total routing cost. Proving this rigorously, however, depends on tighter bounds than Proposition 3.6 or even Theorems 3.3 and 3.4 provide.

We will see experimentally in Section 6 that a routing table of size around $E[S^z(x; \mathcal{A}_G)]$ works well in practice, both when using a single vector and in a decentralized system.

While it may seem like a limitation that this will only work on graphs where routing in a logarithmic number of steps is possible, those are likely to account for most situations where decentralized search is of interest. Beyond Kleinberg’s small-world model, other cases where decentralized routing is expected to take a logarithmic number of steps are hypercubes (where the hamming distance is adapted for routing), and Chord networks [18] (where the circular distance is).

6 Experiments

6.1 Direct Applications

We start by simulating the algorithm under the most basic conditions. We let G consist of a single directed cycle of n vertices, augmented with $\log n$ outgoing shortcuts from each vertex, according to Kleinberg’s small-world model. That is, each shortcut from x is to an independently chosen vertex selected with according to (6) with $\alpha = 1$, which in this case translates to $1/h_n d(x, y)$, where d is distance in H , and $h_n \approx \log n$ is a normalizer.




Figure 4: The cost and steps taken by cost-greedy and greedy search on networks with $\text{Exp}(1)$ distributed costs along each edge, plotted against the size of the network. Networks consist of directed rings with $\log_2 n$ directed shortcuts per vertex.

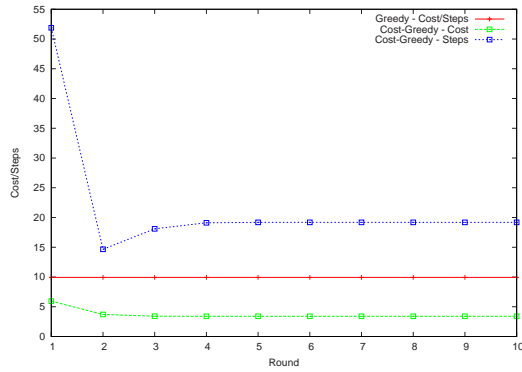


Figure 5: The cost and steps taken by cost-greedy and greedy routing on a network of size 262144 with $\text{Exp}(1)$ distributed costs along each edge. Cost-greedy performance is plotted against iteration of the system in Theorem 3.1, starting with all zero weights.

We start by assigning $w_0(x) = 0$ for all $x \in V$, and calculate the expected routing times by simulating $20n$ queries between randomly chosen points (this number of iterations is probably excessive). We re-sample the costs for each query, but the graph is kept the same. However, because only one sample vector of the expected routing times over each distance is kept, we still end up marginalizing over the shortcuts. We use the logarithmic compression of w described in Section 5 (in practice, we find this outperforms using a full w vector except when an extremely large number of queries is simulated), and use w_{10} as an estimate of the final value.

The difference between cost-greedy and standard greedy search in terms of query cost depends crucially on the distribution of $C(x, y)$. Quantitatively, it is possible to make the benefit of cost greedy as large (or small) as one wishes by a strategic choice of this distribution. For example, if

$$C(x, y) = \begin{cases} 2 & \text{with probability } \frac{1}{2} \\ 0 & \text{otherwise.} \end{cases}$$

then cost-greedy search will most often incur zero cost assuming the vertex degree is large enough (as will a simple lowest cost routing). It would thus be dishonest of us to claim that our methods are motivated based on the performance achieved with delays chosen by us. The experiments in this section are thus meant to verify that cost-greedy search behaves as expected, rather than to illustrate its benefit: the potential benefit of the algorithm must be evaluated for every particular situation where it may be applied.

Our first experiment, shown in Figure 3, is thus to see what happens if we fix the costs to 1 for all the edges. In this case cost and steps are the same, and since it is known (Theorem 3.5) that greedy search is optimal in the expected number of steps, the theory tells us that cost-greedy should, ideally, give the same value. In fact we find that it under-performs by about 5% in all the sizes tested – presumably due to the empirical estimate of the expected value, and the losses due to the logarithmic compression of the weight vector.

Figure 4 shows the performance of cost-greedy when the costs are exponentially distributed as a function of the graph size. We choose an exponential distribution simply because it is a common model for waiting times, and the mean of 1 means that the cost and steps of a route are of the same scale. We see, as expected, that cost-greedy search is able to produce routes that cost less by taking more steps than normal greedy search does. In Figure 5 we plot performance for a single network size against the iterations of (5) when starting with all zeros. We see that even in a network of hundreds of thousands of vertices, no measurable performance is gained after the fourth round - supporting our hypothesis that convergence is a lot faster than the bound given above.




Figure 6: Example of a network with two different vertex types - one tenth of the vertices are augmented by 55 outgoing edges, whereas the rest get just 5. All edge costs are exactly one, and the values are averaged over four simulations to decrease variance. “SJ Method” is the method of Şimşek and Jensen described in Section 6.2.

6.2 Out-degree Distribution

Another question that has been asked about navigability is how to route in a network if the vertices have variable degree, and if the degree as well as the position of the neighbors is known when the routing decision is made. This problem is motivated by the nature of social networks, which appear to be navigable, but where it is known that the vertex degree follows a heavy-tailed power-law. This problem is in many ways similar to that which we discuss above: like with the edge costs, degree distributions may incentivize away from a pure greedy strategy, and instead call for a trade-off between getting close to the destination, and other factors (in this case, wanting to route a vertex with high degree).

Şimşek and Jensen [13] have studied this problem by simulation. Their method is fundamentally similar to ours: they also seek to choose the neighbor which minimizes the expected number of steps to the destination (as we do if the costs are fixed to a unit value). However, rather than attempting to calculate the fix-point of the weights, they make a rough approximation of the value using the right-hand side of the inequality

$$\mathbf{E}[S^z(x)] \geq \mathbf{P}(S^z(x) > 1) = \mathbf{P}(z \notin N(x)) \quad (13)$$

to estimate the left.

To apply our methods above to the problem, we let the weights be a function not only of the distance to the destination, but also of the degree of the vertex. Since this question related only to the number of steps, we fix all the costs to 1. We expect the weight for any particular distance to be smaller

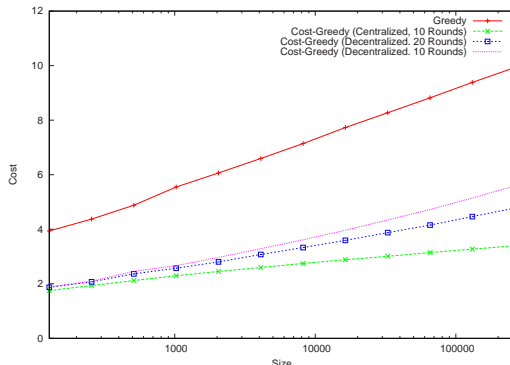


Figure 7: The same situations as in Figure 4, but now including the results when using a separate weight vector at every vertex in the manner described in Section 6.3. A round is $20 \times n$ simulated queries as above.

for vertices with higher degree (since the amount of ground gained in the first step should be better).

Figure 6 shows a simple example of this. In that case we have exactly two possible out-degrees: a few ($0.1n$) of the vertices have 55 shortcuts, while the rest have just 5. We compare the cost-greedy search as used above with regular greedy search and the method of Şimşek and Jensen. The results seem to vindicate the approximation used in the latter method, with this distribution it slightly out-performs cost-greedy search, meaning that the numerical losses in estimating the true weights are greater than the analytic loss of the approximation. This seems to be the case for most sensible such distributions, we find that cost greedy search only takes a slight lead when the popular vertices have more than a hundred times the degree of unpopular ones. If we presume that cost-greedy search can come within 5% of being optimal also here, we are forced to conclude that so does the SJ method.

One advantage that cost-greedy search has over the method of Şimşek and Jensen, is that their method requires detailed knowledge about the model in order to calculate the right-hand side of (13), which cost-greedy search does not. Algorithmically, any vertex may implement cost greedy search for its queries, and it needs only have the ability to measure the cost of the queries it sends to its neighbors, nothing more.

6.3 A Decentralized and Generalized Implementation

To look at the practical viability of the algorithms described above, we also simulate a completely distributed variant. In the decentralized variant, we equip each vertex with its own weight vector, measuring the mean cost of routing from it to destinations at varying distances. Like before we use a




Figure 8: Using a separate weight vector at every vertex in the manner described in Section 6.3. Costs are $\text{Exp}(1)$ distributed, and out-degrees are distributed according to a power-law with tail-exponent 2. “SJ Method” is the method of Şimşek and Jensen described in Section 6.2.

$\log_2 n$ compression of the weights – coalescing all distances between 2^k and 2^{k+1} into the same entry – but unlike above we do not calculate each weight by sampling over a fixed number of queries. Instead, we let the weight vector at each vertex x be calculated as the mean of the entries in a FIFO buffer, showing the cost of the last m queries x has routed destined for vertices of every distance category. As before we do not change the edges of the graph between the queries from which the weights are estimated. Because no marginalizing is occurring here, the graph model is actually one *fixed* realization for each size – the expectation is actually taken only over the costs.

To route a query, x uses the weight vectors of each of the vertices in $N(x)$ to minimize (2). In a real world implementation, these values could be periodically copied between neighbors. One problem we find with this method is that if x initially has m queries in a certain distance category that incur a very high cost, he will not attract more queries in that category from his neighbors (who see it as very costly to send such queries to him), meaning it takes a long time to clear the errant values from the FIFO buffer. Eventually the buffer will be replaced, if not otherwise then by the cost of the queries initiated at x itself, but in our simulations we find that this slows the convergence. To alleviate this, we keep a count of the number of queries x receives for each distance category during an interval. The theory says that these should be equal, so if one of the counts has fallen a lot behind (is less than a quarter of the queries x receives for itself) we set all the values in the buffer to 0.

Even with this method, the convergence is, as expected, slower than in the centralized version. Figure 7 shows the equivalent of Figure 4 but using

local weight vectors at each vertex. Here we let $m = 20$. We can note three things: there is an absolute performance cost of the decentralized version, the cost seems to get worse for larger sizes, but it still considerably outperforms greedy search. The first is probably due to each estimate of the expected routing time being based on far fewer values, while the second is due to us not simulating enough queries for full convergence at the large sizes, as seen by the increasing difference between the values as 10 and 20 rounds. We note that even 20 rounds is actually only 400 queries initiated at each vertex – a large number when we must simulate it for a quarter of a million vertices, but very little compared to the number of queries one would expect in most DHT’s or other deployments of distributed networks.

Finally, in Figure 8 we use the decentralized method to route in a situation when we both have exponential edge costs, and vertices of varying out-degree (in this case a power-law with $\mathbf{P}(|N(x)| > t) \approx t^{-2}$). Decentralized cost-greedy search can optimize both for varying costs and vertex degrees at the same time.

7 Conclusion

We have presented a method for decentralized search that takes varying costs of routing down different edges into account. We have showed that this method is optimal among all such algorithms that monotonically approach the destination of the query, and that the necessary weights can be calculated iteratively. On small-world graphs, we can calculate the order of costs, and say something about the approximation cost. Beyond these analytic facts, we have presented a number of techniques which make the algorithm practical, and experimented with actual implementations using simulation.

It would be very desirable to be able to better motivate our approximations rigorously. To do requires strengthening propositions 3.6 and 3.7, and perhaps a lot of work beyond that. In the short term, proving that any polylogarithmic routing table, and any polynomial number of samples, is sufficient would be a big improvement.

The long term goal of these studies is to try to find adaptive methods for decentralized search when the edge costs are not independent random values, but depend, for instance, on the number of queries that have passed down the edges recently. Such methods for routing with congestion are of interest to deployments of peer-to-peer and other distributed systems.

References

- [1] J. Kleinberg. The small-world phenomenon: an algorithmic perspective. In *Proceedings of the 32nd ACM Symposium on Theory of Computing (STOC)*, 2000.
- [2] J. Kleinberg. Complex networks and decentralized search algorithms. In *Proceedings of the International Congress of Mathematicians (ICM)*, 2006.
- [3] S. M. Ross. *Introduction to Stochastic Dynamic Programming*. Academic Press, 1983.
- [4] J. Kleinberg. Navigation in a small world. *Nature*, page 845, 2000.
- [5] L. Barriere, P. Fraigniaud, E. Kranakis, and D. Krizanc. Efficient routing in networks with long range contacts. In *Proceedings of the 15th International Symposium on Distributed Computing, DISC'01*, 2001.
- [6] J. Kleinberg. Small-world phenomena and the dynamics of information. In *Advances in Neural Information Processing Systems (NIPS) 14*, 2001.
- [7] D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, and A. Tomkins. Geograph routing in social networks. In *Proceedings of the National Academy of Science*, volume 102, pages 11623–11628, 2005.
- [8] P. Fraigniaud, E. Lebhar, and Z. Lotker. A doubling dimension threshold $\theta(\log \log n)$ for augmented graph navigability. In *Proceedings of the 14th European Symposium on Algorithms (ESA)*, 2006.
- [9] O. Sandberg. Neighbor selection and hitting probability in small-world graphs. To appear in *The Annals of Applied Probability*, 2007.
- [10] I. Clarke, T. Hong, S. Miller, O. Sandberg, and B. Wiley. Protecting free expression online with Freenet. *IEEE Internet Computing*, 6:40–49, 2002.
- [11] I. Clarke. Freenet's next generation routing protocol, 2003. <http://freenetproject.org/ngrouting.html>.
- [12] O. Sandberg. Distributed routing in small-world networks. In *Proceedings of the Eighth Workshop on Algorithm Engineering and Experiments (ALENEX06)*, 2006.
- [13] Ö. Şimşek and D. Jensen. Decentralized search in networks using homophily. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, 2005.

- [14] M. Franceschetti and R. Meester. Navigation in small world networks, a scale-free continuum model. *Journal of Applied Probability*, 43:1173–1180, 2006.
- [15] T. Stern. A class of decentralized routing algorithms using relaxation. *IEEE Transactions on Communications*, 25:1092–1102, 1977.
- [16] C.M. Newman and L.S. Schulman. One dimensional $1/|j - i|^s$ percolation models: The existence of a transition for $s \leq 2$. *Communications in Mathematical Physics*, 104:547–571, 1986.
- [17] G. Singh Manku. 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)*, 2004.
- [18] I. Stoica, R. Morris, D. Karger, F. Kaashoek, and H. Balakrishnan. Chord: A scalable Peer-To-Peer lookup service for internet applications. In *Proceedings of the 2001 ACM SIGCOMM Conference*, pages 149–160, 2001.