Optimal Routing in a Small-World Network

Jianyang Zeng* and Wen-Jing Hsu*

Abstract

Recently a bulk of research [14, 5, 15, 9] has been done on the modelling of the small-world phenomenon, which has been shown to be pervasive in social and nature networks, and engineering systems [16, 1, 2, 11]. In order to examine the navigating aspects of small-world graphs, Kleinberg [9] proposes a network model based on a d-dimensional torus lattice with long-range links chosen at random according to the d-harmonic distribution. Kleinberg shows that the greedy routing algorithm, by using only local information, performs in $O(\lg^2 n)$ expected number of hops. We extend Kleinberg's small-world model in that each node x has two more random links to nodes chosen uniformly and randomly within $(\lg n)^{\frac{2}{d}}$ Manhattan distance from x, where d denotes the dimension of the model. Based on this extended model, we then propose an oblivious algorithm that can route messages between any two nodes in $O(\lg n)$ expected number of hops, which is an optimal expected bound for routing. Our routing algorithm keeps only $O((\lg n)^{\beta+1})$ bits of information on each node, where $1 < \beta < 2$, thus being scalable with the network size. To our knowledge, our result is the first to achieve the optimal routing complexity while still keeping a poly-logarithmic number of bits of information stored on each node in the small-world networks.

Our results may be applied to the design of the logical overlay structure of large-scale distributed systems, such as peer-to-peer networks, in the same spirit as Symphony [11]. Since the links are randomly constructed according to the probabilistic distribution, our extended small-world network is less vulnerable to adversarial attacks, and thus provides good fault tolerance.

Key words: small-world networks, augmented local awareness, decentralized routing, design of algorithms, distributed systems.

1 Introduction

Milgram [16] shows the *small-world phenomenon* in the human society, that is, any two people in the world can be connected by a chain of six (on the average) acquaintances, and people can deliver messages efficiently to an unknown target via their acquaintances. The small-world phenomenon has also been shown to be pervasive in networks from nature and engineering systems, such as the World Wide Web [5, 1], peer-to-peer systems [2, 12, 11, 18], etc.

A number of network models have been proposed to study the small-world properties [14, 5, 15, 9]. Newman and Watts [15] propose a random rewiring model whose diameter is a poly-logarithmic function of the size of the network. The model is constructed by adding a small number of random

*Center for Advanced Information Systems, Nanyang Technological University, Nanyang Avenue, Singapore 639798. Email: {pg03858494, hsu}@ntu.edu.sg.

edges to nodes uniformly distributed on a ring, where the nodes are densely connected with their near neighbors. However, Newman and Watts's model does not address the routing issues of small-world networks [9]. The poly-logarithmic diameter of some graphs does not imply the existence of efficient routing algorithms [9]. For example, the random graph in [3] yields a logarithmic diameter, yet any routing using only local information requires at least $\sqrt{n}$ expected number of hops (where n is the size of the network) [9].

In order to examine the routing on random graphs, Kleinberg [9] develops a new model based on a d-dimensional torus lattice with long-range links chosen randomly from the d-harmonic distribution, that is, a long-range link between nodes u and v exists with probability proportional to $Dist(u, v)^{-d}$, where $Dist(u, v)$ denotes the Manhattan distance between nodes u and v. Based on this model, Kleinberg then shows that the simple greedy routing algorithm by using only local information can route messages between any two nodes in $O(\lg^2 n)$ [1] expected number of hops. This bound is tightened to $\Theta(\lg^2 n)$ later by Barrière et al. [4] and Martel et al. [13]. Further research [12, 10, 13, 6] shows that in fact the $O(\lg^2 n)$ bound of the original greedy routing algorithm can be improved by putting some extra information in each message holder. Manku et al. [12] show that if each message holder at a routing step takes its own neighbors' neighbors into account for making routing decisions, the bound of routing complexity can be improved to $O(\frac{\lg^2 n}{q \lg q})$, where q denotes the number of long-range contacts for each node. Lebhar and Schabanel [10] propose a routing algorithm for 1-dimensional Kleinberg's model, which visits $O(\frac{\lg^2 n}{\lg^2(1+q)})$ nodes on expectation before routing the message, and they show that a routing path with expected length of $O(\frac{\lg n(\lg \lg n)^2}{\lg^2(1+q)})$ can be found. Two research groups, Fraigniaud et al. [6] and Martel and Nguyen [13], independently report that if each node is aware of its $O(\lg n)$ closest local neighbors, the routing complexity in d-dimensional Kleinberg's small-world networks can be improved to $O((\lg n)^{1+1/d})$ expected number of hops. The difference is that [13] requires keeping additional state information, while [6] uses an oblivious greedy routing algorithm. In [13], Martel and Nguyen show that the expected diameter of a d-dimensional Kleinberg network is $\Theta(\lg n)$. As such, there is still some room for improving the routing complexity, which motivates our work.

In [17], we have proposed a one-dimensional extended small-world model with augmented local links, and presented both non-oblivious and oblivious routing algorithms that can route messages between any two nodes in $O(\lg n \lg \lg n)$ expected number of hops. In this paper, we propose a d-dimensional extended version of Kleinberg's small-world model, where each node has two more random links to nodes within certain Manhattan distance. Based on this model, we present an oblivious decentralized algorithm that can finish routing in $O(\lg n)$ expected number of hops, which is an optimal routing complexity.

Applications of small-world phenomenon in computer networks include efficient lookup in peer-to-peer systems [12, 2, 11, 18], gossip protocol in a communication network [8], flooding routing in ad-hoc networks [7], etc.

[1] The logarithmic symbol lg is with the base 2, if not otherwise specified. Also, we remove the ceiling or floor for simplicity throughout the paper.

Organization of the paper. The rest of the paper is organized as follows. Section 2 introduces the extended small-world model and the decentralized routing algorithm. Section 3 lists our main results and contributions. In Section 4, we theoretically analyze our decentralized routing algorithm. Section 6 gives a brief conclusion of this paper.

2 The Small-World Model and Decentralized Routing

Our small-world model is an extension of Kleinberg's d-dimensional model [9]. It is based on a d-dimensional torus $[n]_d = \{0, 1, \cdots, n\}^d$ with three extra links for each node, where $d \geq 2$. Firstly, as in Kleinberg's original model [9], each node has a long-range link to another node chosen randomly according to the d-harmonic distribution, that is, the probability that node u sends a long-range link to another node v is $\Pr[u \to v] = \frac{1}{Z_u \cdot Dist(u,v)^d}$, where $Dist(u, v)$ denotes the Manhattan distance between nodes u and v, and $Z_u = \sum_{z \neq u} \frac{1}{Dist(u,z)^d}$. To avoid confusing with the extra links to be introduced shortly, we refer to such long-range links as the **K-type links** or **K-links** for short (where K stands for Kleinberg), and refer to node v as a **K-neighbor** of node u if there exists a K-link from u to v. Here we will introduce two more extra links for each node u to nodes that are chosen uniformly at random from nodes within $(\lg n)^{2/d}$ Manhattan distance from u. We refer to these two links as **augmented local links** or **AL-links** for short, and refer to node v as an **AL-neighbor** of node u if there exists an AL-link from u to v. Finally, we refer to the local links on the torus as **torus-links** or **T-links** for short, and refer to the local neighbors of node u on the torus as u's **T-neighbors**. We refer to all nodes linked by u, including its K-neighbor, AL-neighbor and T-neighbor, as the *immediate neighbors* of node u.

We assume that all T-links on the torus are undirected, while all extra links including K-links and AL-links are directed. Obviously, there are $2d + 3$ immediate neighbors for each node in our extended small-world model. Thus, our extended model retains the same $O(1)$ order of node degree as that in Kleinberg's small-world model. Throughout this paper, we use the terms *model* and *network* interchangeablly.

In our decentralized routing algorithm, the message holder is also referred to as the *current node*. Given the current node x, let $\Gamma_x(0) = \{x\}$, and let $\Gamma_x(1)$ denote the AL neighborhood of all nodes in $\Gamma_x(0)$, and $\Gamma_x(2)$ denote the AL neighborhood of all nodes in $\Gamma_x(1)$, and so on. In other words, we refer to $\Gamma_x(i)$ as the ith level of AL neighborhood for node x, and refer to $A_x(i) = \bigcup_{j \leq i} \Gamma_x(j)$ as the first i levels of AL neighborhood for node x. At a certain level i of AL neighborhood, we may also refer to $A_x(i-1)$ as the set of *previously known nodes*. Let $L_x(i) = A_x(i) - A_x(i-1)$ denote the set of *new* nodes discovered during the ith level of AL neighborhood. We will call $A_x(k)$ the *AL awareness* of node x, if each node in our extended small-world model is aware of the first k levels of its AL neighborhood.

There are normally two approaches for decentralized routing: oblivious and non-oblivious schemes [6]. A routing algorithm is oblivious if the message holder makes routing decisions only based on its own routing table and the target node. On the other hand, a routing algorithm is said to be non-oblivious if the routing decisions of the message holder depend on the previous routing history stored in the message header, the target node and its routing table. The scheme in [6] is oblivious, while the schemes in [10] and [13] are non-oblivious. In this paper, we only consider the oblivious routing scheme.

The description of our oblivious routing algorithm is given as follows (Algorithm 1).

Algorithm 1

Input: the source s and the target t.

Initialization:

$x \leftarrow s$.

The first phase: $Dist(x,t) \geq (\lg n)^{\frac{2}{d}+1}$.

```
while Dist(x,t) ≥ (lg n)^(2/d+1) do
    node x detects in its AL awareness A_x(β lg lg n) whether there exists a node z that contains a K-neighbor within m/lg^τ n
    Manhattan distance from t, where 1 < β < 2 and τ denotes a certain constant which will be given later. Let Z_x denote
    the set of such nodes z.
    if Z_x is empty then
        Kleinberg's greedy algorithm is executed, that is, the message is routed to an immediate neighbor closest to t.
    else
        node x finds a node z in Z_x that is closest to x in terms of AL-links (ties are broken arbitrarily).
        node x computes a shortest path π : x_0 = x, x_1, ·, ·, x_t = z from x to z among A_x(β lg lg n).
    end if
    if the shortest path π consists of only x itself then
        node x routes the message to its K-neighbor.
    else
        node x routes the message to its next AL-neighbor x_1 along the shortest path π.
    end if
end while
```

The final phase: $Dist(x,t) < (\lg n)^{\frac{2}{d}+1}$

Kleinberg's greedy algorithm is executed, that is, the message is forwarded to an immediate neighbor closest to the target node t, until it reaches t.

3 Our Contributions

Our main result is given as follows.

***Theorem* 1** *If each node in the extended small-world model is aware of the first $\beta \lg\lg n$ levels of its AL neighborhood, where $1 < \beta < 2$, then there exists an oblivious algorithm that can route messages between any two nodes in $O(\lg n)$ expected number of hops.*

Since any graph with $O(1)$ node degree has $\Omega(\lg n)$ diameter, the $O(\lg n)$ expected bound of routing complexity in Theorem 1 is an optimal bound. In addition, since $A_x(\beta \lg\lg n) = 1 + 2 + \cdots + 2^{\beta \lg\lg n} \leq 2\lg^{\beta} n$, the number of bits required on each node is at most $O((\lg n)^{\beta+1})$, where $1 < \beta < 2$. To our knowledge, this is the first result that achieves the optimal routing complexity while still keeping a poly-logarithmic number of bits information stored on each node in the small-world networks. A comparison of our scheme with the other existing results is shown in Table 1.

Our results may be applied to the design of the logical overlay structure of large-scale distributed systems, such as peer-to-peer networks, in the same spirit as Symphony [11]. Since the links in the model are randomly constructed according to the probabilistic distribution, our extended network is less vulnerable to adversarial attacks, and thus provides good fault tolerance.

Scheme	#bits of awareness	#hops expected	Oblivious or Non-oblivious?
Kleinberg's greedy [9, 2, 4]	$O(q \lg n)$	$O(\lg^2 n/q)$	Oblivious
NoN-greedy [12]	$O(q^2 \lg n)$	$O(\lg^2 n/(q \lg q))$	Non-oblivious
Decentralized algorithm in [10]	$O(\lg^2 n/\lg(1+q))$	$O((\lg n)^2/\lg^2(1+q))$	Non-oblivious
Decentralized algorithm [13]	$O(\lg^2 n)$	$O((\lg n)^{1+1/d})$	Non-oblivious
Indirect-greedy algorithm [6]	$O(\lg^2 n)$	$O((\lg n)^{1+1/d})$	Oblivious
Near optimal algorithm for one-dimensional model with augmented awareness [17]	$O(\lg^2 n)$	$O(\lg n \ \lg \lg n)$	Both are considered
Optimal algorithm for d-dimensional model with augmented awareness [this paper]	$O((\lg n)^{\beta+1})$ $(1 < \beta < 2)$	$O(\lg n)$	Oblivious

Table 1: Comparisons of our decentralized routing algorithms with the other existing schemes. In the first three schemes (in [9, 2, 11, 12, 10]), we suppose that each node has q K-links, while in the next four schemes (in [13, 6, 17] and this paper), we suppose that each node has one K-link.

4 Analysis of Decentralized Routing

In this section, we will give the proof of Theorem 1. A road map of the proof is given as follows. In Lemma 2, we first show that if each node is aware of the first $\beta \lg \lg n$ levels of its AL neighborhood, where $1 < \beta < 2$, the number of distinct nodes known to each node is at least $(\lg n)^{\delta+1}$ with certain probability, where $0 < \delta < \beta - 1$. Based on this result, in Lemma 3 we demonstrate that the AL awareness of each node is very likely to contain a K-neighbor that is close to the target node. Then in Lemma 7 and 9 we show that our oblivious routing algorithm can reduce the Manhattan distance effectively so that the $O(\lg n)$ expected bound of routing complexity can be achieved.

The following lemma from [2, 9] is useful for our subsequent analysis.

***Lemma* 1** *Let* $\Pr[u \xrightarrow{K} v]$ *denote the probability that node* u *sends a K-link to node* v *in a* d*-dimensional small-world model. Suppose that* $a \le Dist(u, v) \le b$*, then* $\frac{c_1}{b^d \lg n} \le \Pr[u \xrightarrow{K} v] \le \frac{c_2}{a^d \lg n}$*, where* c_1 *and* c_2 *are constants independent of* n*.*

We first have the following result on the number of distinct nodes in the AL awareness of each node.

***Lemma* 2** *Let* β *denote a constant such that* $1 < \beta < 2$*. Let* $A_x(\beta \lg \lg n)$ *denote the AL awareness of node* x *in the extended small-world model, where each node is aware of the first* $\beta \lg \lg n$ *levels of its AL neighborhood. Then there exists a constant* $0 < \delta < \beta - 1$ *such that*

$$\Pr[\ \ |A_x(\beta \lg \lg n)| \ge (\lg n)^{1+\delta}\ ] > 1 - \frac{1}{\lg^{\xi} n},$$

where $\xi > 0$*.*

Proof: The proof is divided into two parts. In Part 1, we show that with probability at least $1 - 8(\lg n)^{-\frac{4}{3}}$, $|L_x(\frac{1}{3} \lg \lg n)| \ge (\lg n)^{\frac{1}{3}}$, that is, every AL-link during the first $\frac{1}{3} \lg \lg n$ levels of x's AL neighborhood points to a new node. In Part 2, by using Chernoff's bound, we show that $|L_x(i)|$ increases at an exponential rate for all $\frac{1}{3} \lg \lg n < i < \beta \lg \lg n$, based on which we achieve $|L_x(\beta \lg \lg n)| \ge (\lg n)^{1+\delta}$ with probability at least $1 - \frac{1}{\lg^{\xi} n}$, where $\xi > 0$.

Part 1: We will show that at each level of AL neighborhood (among the first $\frac{1}{3}\lg\lg n$ levels of AL neighborhood), the probability that each AL-link points to a previously known node is so small that all AL-links tend to point to new nodes, and hence $|L_x(\frac{1}{3}\lg\lg n)| \geq (\lg n)^{1/3}$ very likely.

We first calculate the following upper bound for $|A_x(\frac{1}{3}\lg\lg n)|$.

$$|A_x(\frac{1}{3}\lg\lg n)| \leq 1 + 2 + 2^2 + \cdots + 2^{\frac{1}{3}\lg\lg n} = 2(\lg n)^{1/3} - 1 < 2(\lg n)^{1/3}.$$

Thus, we have $|A_x(i)| \leq |A_x(\frac{1}{3}\lg\lg n)| < 2(\lg n)^{1/3}$ for all $0 \leq i \leq \frac{1}{3}\lg\lg n$. Consider the construction of an AL-link of node x. Since each AL-link is connected to a node chosen randomly and uniformly from $\lg^2 n$ closest local nodes, each AL-link points to a node within the Manhattan distance $\lg^{2/d} n$ with an equal probability $(\lg n)^{-2}$. Since $|A_x(i)| \leq 2(\lg n)^{1/3}$, there are no more than $2(\lg n)^{1/3}$ previously known nodes at each level of AL neighborhood. Hence, the probability that any AL-link is connected to a previously known node is at most $2(\lg n)^{1/3} \cdot (\lg n)^{-2} = 2(\lg n)^{-5/3}$. Thus, the probability that an AL-link points to a new node is at least $1 - 2(\lg n)^{-5/3}$. There are in total at most $2|A_x(\frac{1}{3}\lg\lg n)| \leq 4(\lg n)^{1/3}$ number of AL-links, so all AL-links point to new nodes with probability at least $(1 - 2(\lg n)^{-5/3})^{4(\lg n)^{1/3}} \geq 1 - 8(\lg n)^{-4/3}$ for sufficiently large n. Here we use the fact $(1+x)^a \geq 1 + ax$ for $x > -1$ and $a \geq 1$. Thus, we have $\Pr[|L_x(\frac{1}{3}\lg\lg n)| \geq (\lg n)^{1/3}] \geq 1 - 8(\lg n)^{-4/3}$.

Part 2: Let $\mathcal{B}$ denote the event that $|L_x(\frac{1}{3}\lg\lg n)| \geq (\lg n)^{1/3}$. From Part 1, we have $\Pr[\mathcal{B}] \geq 1 - 8(\lg n)^{-4/3}$. Next, we will consider the sequence of $L_x(i)$ for $\frac{1}{3}\lg\lg n \leq i \leq \beta\lg\lg n$. We assume that $|A_x(i)| \leq (\lg n)^{1+\delta}$ for all $\frac{1}{3}\lg\lg n \leq i \leq \beta\lg\lg n$, otherwise the lemma holds true. Since each AL-link is connected to a node chosen randomly and uniformly from $\lg^2 n$ closest local nodes, and there are at most $(\lg n)^{1+\delta}$ previously known nodes at each level of AL neighborhood, each AL-link reveals a new node with probability at least $1 - (\lg n)^{-2} \cdot (\lg n)^{1+\delta} = 1 - (\lg n)^{\delta-1}$. Let X_i denote the sum of $2|L_x(i)|$ independent Bernoulli random variables each with expectation $1 - (\lg n)^{\delta-1}$. Then $|L_x(i+1)|$ stochastically dominates X_i for all $\frac{1}{3}\lg\lg n \leq i \leq \beta\lg\lg n$. By Chernoff's bound, there exists a constant $0 < \epsilon < 1$ such that

$$\Pr[\ |L_x(i+1)| \leq 2(1 - (\lg n)^{\delta-1})(1-\epsilon) \cdot |L_x(i)|\ ] \leq \exp\big(-\epsilon^2(1 - (\lg n)^{\delta-1})|L_x(i)|\big).$$

Let $\mathcal{E}_i$ denote the event that $|L_x(i+1)| \geq 2(1 - (\lg n)^{\delta-1})|L_x(i)|(1-\epsilon)$, where $\frac{1}{3}\lg\lg n \leq i \leq \beta\lg\lg n$, then we have $\Pr[\mathcal{E}_i] \geq 1 - \exp(-\epsilon^2(1 - (\lg n)^{\delta-1})|L_x(i)|)$. Let $\mathcal{E}$ denote the occurrences of the consecutive successful events $\mathcal{B}, \mathcal{E}_{\frac{1}{3}\lg\lg n}, \mathcal{E}_{\frac{1}{3}\lg\lg n+1}, \cdots, \mathcal{E}_{\beta\lg\lg n}$, then for large n, we have

$$\Pr[\mathcal{E}] \geq (1 - 8(\lg n)^{-\frac{4}{3}})(1 - \exp(-\epsilon^2(1 - (\lg n)^{\delta-1})(\lg n)^{\frac{1}{3}}))^{\beta\lg\lg n} > 1 - \frac{1}{\lg^{\xi} n},$$

where $\xi > 0$. At the last step, we applied the fact that $(1+x)^a \geq 1 + ax$ for $x > -1$ and $a \geq 1$.

When the event $\mathcal{E}$ occurs, we have

$$\begin{aligned}
|L_x(\beta \lg\lg n)| &\geq |L_x(\frac{1}{3}\lg\lg n)| \cdot \left(2(1-(\lg n)^{\delta-1})(1-\epsilon)\right)^{(\beta-\frac{1}{3})\lg\lg n} \\
&\geq (\lg n)^{\frac{1}{3}} \left(2(1-(\lg n)^{\delta-1})(1-\epsilon)\right)^{(\beta-\frac{1}{3})\lg\lg n} \\
&= (\lg n)^{\frac{1}{3}} (\lg n)^{(\beta-\frac{1}{3})\lg\left(2(1-(\lg n)^{\delta-1})(1-\epsilon)\right)} \\
&= (\lg n)^{\frac{1}{3}+(\beta-\frac{1}{3})\lg\left(2(1-(\lg n)^{\delta-1})(1-\epsilon)\right)}
\end{aligned}$$

Given a constant $1 < \beta < 2$, we can always find suitable constants $0 < \delta < \beta - 1$, $0 < \epsilon < 1$ and n_0 such that for all $n > n_0$, $\frac{1}{3} + (\beta - \frac{1}{3})\lg\left(2((1-(\lg n)^{\delta-1})(1-\epsilon)\right) \geq 1+\delta$. Thus, there exists a constant $0 < \delta < \beta - 1$ such that $\Pr[|A_x(\beta \lg\lg n)| \geq (\lg n)^{1+\delta}] > \Pr[|L_x(\beta \lg\lg n)| \geq (\lg n)^{1+\delta}] > 1 - \frac{1}{\lg^{\xi} n}$, where $\xi > 0$. Therefore, the proof of the lemma is completed. ■

Next, we will show that the AL awareness of the current node is very likely to contain a K-neighbor within $m/\lg^{\tau} n$ Manhattan distance from the target node t, where τ denotes a certain constant.

***Lemma* 3** *Suppose that the Manhattan distance between the current node x and the target node t in the extended small-world model is $m \geq (\lg n)^{\frac{2}{d}+1}$. Then there exists a constant τ such that with probability at least $1 - \frac{1}{\lg^{\zeta} n}$, where $\zeta > 0$, x's AL awareness $A_x(\beta \lg\lg n)$ contains a K-neighbor within $\frac{m}{\lg^{\tau} n}$ Manhattan distance from the target node t.*

Proof: Let $\mathcal{C}$ denote the event that $|A_x(\beta \lg\lg n)| \geq (\lg n)^{1+\delta}$. By Lemma 2, we have $Pr[\mathcal{C}] > 1 - \frac{1}{\lg^{\xi} n}$, where $\xi > 0$.

Let $D_l(t)$ denote the set of all nodes within l Manhattan distance from t. Given a node u in $A_x(\beta \lg\lg n)$, let $\Pr[u \xrightarrow{K} D_l(t)]$ denote the probability that u's K-neighbor is inside the ball $D_l(t)$.

Since each AL-link spans the Manhattan distance no more than $(\lg n)^{2/d}$, the nodes in x's AL awareness $A_x(\beta \lg\lg n)$ are all within $\beta \lg\lg n (\lg n)^{2/d}$ Manhattan distance from x. Since $Dist(x,t) = m \geq (\lg n)^{\frac{2}{d}+1}$, the maximum Manhattan distance between a node in $A_x(\beta \lg\lg n)$ and any node in $D_{\frac{m}{\lg^{\tau} n}}(t)$ is no more than $2m$. By Lemma 1, the probability for u's K-neighbor to be inside the ball $D_{\frac{m}{\lg^{\tau} n}}(t)$ is at least $\frac{c_1}{(2m)^d \lg n}$, so we have

$$\Pr[u \xrightarrow{K} D_{\frac{m}{\lg^{\tau} n}}(t)] \geq |D_{\frac{m}{\lg^{\tau} n}}(t)| \cdot \frac{c_1}{(2m)^d \lg n} = \left(\frac{m}{\lg^{\tau} n}\right)^d \cdot \frac{c_1}{(2m)^d \lg n} = \frac{c_3}{(\lg n)^{\tau d+1}},$$

where c_3 denotes a constant.

Let $\Pr[A_x(\beta \lg\lg n) \xrightarrow{K} D_{\frac{m}{\lg^{\tau} n}}(t)]$ denote the probability that at least one node in $A_x(\beta \lg\lg n)$ contains a K-neighbor within $D_{\frac{m}{\lg^{\tau} n}}(t)$. Then if $\tau d < \delta$, we have

$$
\begin{aligned}
\Pr[A_x(\beta \lg\lg n) \xrightarrow{K} D_{\frac{m}{\lg^\tau n}}(t)] &\geq \Pr[A_x(\beta \lg\lg n) \xrightarrow{K} D_{\frac{m}{\lg^\tau n}}(t) \mid \mathcal{C}] \cdot \Pr[\mathcal{C}] \\
&\geq \big(1 - (1 - \frac{c_3}{\lg^{\tau d+1} n})^{(\lg n)^{1+\delta}}\big) \cdot (1 - \frac{1}{\lg^\xi n}) \\
&\geq \big(1 - \exp(-c_3 (\lg n)^{\delta - \tau d})\big) \cdot (1 - \frac{1}{\lg^\xi n}) \\
&\text{(using the fact } (1+\tfrac{b}{x})^x \leq e^b,\ b \in \mathbb{R},\ x > 0) \\
&> 1 - \frac{1}{\lg^\zeta n} \quad \text{(because } \tau d < \delta),
\end{aligned}
$$

for a constant $\zeta > 0$. ■

By using a similar technique as that in Lemma 3, we can have the following lemma.

***Lemma* 4** *Suppose that the Manhattan distance between the current node x and the target node t in the extended small-world model is $m \geq (\lg n)^{\frac{2}{d}+1}$. Then there exists a constant τ such that with probability at least $1 - \frac{1}{\lg^\zeta n}$, where $\zeta > 0$, x's AL awareness $A_x(\beta \lg\lg n)$ contains a K-neighbor within $\frac{m - \beta(\lg n)^{2/d} \lg\lg n}{\lg^\tau n}$ Manhattan distance from the target node t.*

Given the current node x and the target node t, we refer to the set of nodes within $Dist(x,t)/\lg^\tau n$ Manhattan distance from t as the *influenced set* of node x, where τ is a certain constant as given above.

***Lemma* 5** *Suppose that the Manhattan distance between the current node x and the target node t in the extended small-world model is $m \geq (\lg n)^{\frac{2}{d}+1}$. Then a node within $(m - \beta(\lg n)^{2/d} \lg\lg n)/\lg^\tau n$ Manhattan distance from t is also inside the influenced set of any node in $A_x(\beta \lg\lg n)$.*

Proof: Since each AL-link spans no more than $(\lg n)^{2/d}$ Manhattan distance from the current node, and there are in total $\beta \lg\lg n$ levels of AL neighborhood for x, all nodes in $A_x(\beta \lg\lg n)$ span no more than $\beta(\lg n)^{2/d} \lg\lg n$ Manhattan distance from x. By this simple observation, the proof of the lemma can be easily obtained. ■

***Lemma* 6** *Suppose that the Manhattan distance between the current node x and the target node t in the extended small-world model is $m \geq (\lg n)^{\frac{2}{d}+1}$. Let I_1 denote the set of nodes within Manhattan distance $\frac{m}{\lg^\tau n}$ from t, and let I_2 denote the set of the nodes within Manhattan distance $\frac{m - \beta(\lg n)^{2/d} \lg\lg n}{\lg^\tau n}$ from t. Then the probability for $A_x(\beta \lg\lg n)$ to contain a K-neighbor within $I_1 - I_2$ is no more than $\frac{1}{\lg^\varphi n}$, where $\varphi > 0$.*

Proof: We first calculate an upper bound of $|I_1 - I_2|$. We have

$$
|I_1 - I_2| = \big(\frac{m}{\lg^\tau n}\big)^d - \big(\frac{m - \beta(\lg n)^{2/d} \lg\lg n}{\lg^\tau n}\big)^d = \frac{m^d - (m - \beta(\lg n)^{2/d} \lg\lg n)^d}{\lg^{\tau d} n}.
$$

When $m \geq (\lg n)^{2/d+1}$, we have $(m - \beta(\lg n)^{2/d} \lg\lg n)^d \geq m^d - m^{d-1}(\beta(\lg n)^{2/d} \lg\lg n)$. Thus, we can obtain the following upper bound for $|I_1 - I_2|$

$$|I_1 - I_2| \leq \frac{m^{d-1}(\beta(\lg n)^{2/d} \lg\lg n)}{\lg^{\tau d} n}.$$

By Lemma 1, the probability for a node y_1 in $A_x(\beta \lg\lg n)$ to send a K-link to a node y_2 in $I_1 - I_2$ is at most $\frac{c_2}{(m/2)^d \lg n}$, since $Dist(y_1, y_2) \geq m/2$ if $m \geq (\lg n)^{\frac{2}{d}+1}$. Since $|A_x(\beta \lg\lg n)| \leq 2\lg^\beta n$, we have

$$\begin{aligned}
&\Pr[A_x(\beta \lg\lg n) \text{ contains a K-neighbor within } I_1 - I_2\ ] \\
&\leq \frac{c_2}{(m/2)^d \lg n} \cdot |A_x(\beta \lg\lg n)| \cdot |I_1 - I_2| \\
&\leq \frac{2^d c_2}{m^d \lg n} \cdot 2\lg^\beta n \cdot \frac{m^{d-1}(\beta(\lg n)^{2/d} \lg\lg n)}{\lg^{\tau d} n} \\
&= \frac{2 \cdot 2^d c_2 \beta (\lg n)^{\beta + \frac{2}{d}} \lg\lg n}{m(\lg n)^{1+\tau d}} \\
&\leq \frac{2 \cdot 2^d c_2 \beta (\lg n)^{\beta} \lg\lg n}{(\lg n)^{2+\tau d}} \ \text{(since } m \geq (\lg n)^{\frac{2}{d}+1}\text{)} \\
&\leq \frac{1}{\lg^\varphi n} \ \text{(where } \varphi > 0\text{)}
\end{aligned}$$

Thus, the proof of the lemma is completed. ∎

Lemma 7 *Suppose that the Manhattan distance between the current node x and the target node t in the extended small-world model is $m \geq (\lg n)^{\frac{2}{d}+1}$. Then after at most $O(\lg\lg n)$ expected number of hops, the message will reach a node within $m/\lg^\tau n$ Manhattan distance from t, where τ denotes a certain constant.*

Proof: Like in Lemma 6, let I_1 denote the set of nodes within Manhattan distance $\frac{m}{\lg^\tau n}$ from t, and let I_2 denote the set of the nodes within Manhattan distance $\frac{m - \beta(\lg n)^{2/d} \lg\lg n}{\lg^\tau n}$ from t. Let $\mathcal{F}_1$ denote the event that $A_x(\beta \lg\lg n)$ contains a K-neighbor within $I_1 - I_2$, and let $\mathcal{F}_2$ denote the event that $A_x(\beta \lg\lg n)$ contains a K-neighbor within I_2. By Lemma 6 and Lemma 4, we have $\Pr[\mathcal{F}_1] \leq \frac{1}{\lg^\varphi n}$ and $\Pr[\mathcal{F}_2] \geq 1 - \frac{1}{\lg^\zeta n}$ respectively, where $\varphi, \zeta > 0$. Thus, with probability at least a positive constant, $A_x(\beta \lg\lg n)$ contains a K-neighbor within I_2, but no K-neighbor within $I_1 - I_2$.

We refer to the routing steps from a given node x to its intermediate node z in $A_x(\log\log n)$ as an indirect phase. The routings in different indirect phases are independent from each other. By above statement, after at most $O(1)$ expected number of indirect phases, i.e., at most $c' \cdot \beta \lg\lg n$ expected number of hops for a constant c', the message will be routed to a node x whose AL awareness contains a K-neighbor within I_2, but no K-neighbor within $I_1 - I_2$. Let z be the intermediate node in $A_x(\beta \lg\lg n)$ that contains a K-neighbor within the influenced set I_1 (or I_2) and is closest to node x in terms of AL-links. Next, we will show that after the event $\mathcal{F}_2 \bigcap \neg\mathcal{F}_1$ occurs, our oblivious algorithm will route the message to the intermediate node z along a shortest path $\pi : x_0 = x, x_1, \cdots, x_t = z$ among $A_x(\beta \lg\lg n)$.

We refer to a node z in $A_x(\lg\lg n)$ as a *good intermediate node* if it satisfies the following two conditions: (1) has a K-neighbor within $\frac{m}{\lg^\tau n}$ to the target node; (2) is closest to node x in terms of AL-links.

We first consider the case where no tie happens, that is, x's good intermediate node z is unique. By our oblivious algorithm, node x will route the message to its next AL-neighbor x_1 along a shortest path π. From Lemma 5, node z's K-neighbor is also inside the influenced set of x_1, and hence it satisfies the first condition of a good intermediate node for x_1. In addition, since z is an intermediate node closest to x, it is also an intermediate node closest to x_1. Thus, it also satisfies the second condition of a good intermediate node for x_1. Therefore, node x_1 will also regard node z as its good intermediate node, and find a shortest path $\pi : x_1, x_2, \cdots, x_t = z$ from x_1 to z, and then route the message to its next AL-neighbor x_2. Such a process is repeated for every node x_i on the shortest path π until the message reaches the intermediate node z. After that, the message will be routed to z's K-neighbor and hence reach a node within $m/\lg^\tau n$ Manhattan distance from the target node t.

When the ties happens, there may be more than one good intermediate nodes z for the current node x. However, it is easy to show that the message will be routed to a good intermediate node z along one of shortest pathes. This does not affect the result. Therefore, the proof of the lemma is completed. ∎

***Lemma* 8** *Suppose that the Manhattan distance between the current node x and the target node t in the extended small-world model is $m \geq (\lg n)^{\frac{2}{d}+1}$. Then after at most $O(\lg n)$ expected number of hops, our oblivious routing algorithm can reduce the Manhattan distance to within $(\lg n)^{\frac{2}{d}+1}$.*

Proof: By Lemma 7, after at most $O(\lg\lg n)$ expected number of hops, the message will reach a node within $\frac{m}{\lg^\tau n}$ Manhattan distance from t, where τ denotes a certain constant.

Divide the whole Manhattan distance $Dist(x,t)$ into phases such that the ith phase contains the nodes within $[\frac{m}{(\lg n)^{\tau i}}, \frac{m}{(\lg n)^{\tau(i-1)}})$ Manhattan distance from t. Since the maximum Manhattan distance is n, there are at most $O(\frac{\lg n}{\lg\lg n})$ phases. Because each phase takes at most $O(\lg\lg n)$ expected number of steps by Lemma 7, after $O(\frac{\lg n}{\lg\lg n}) \cdot O(\lg\lg n) = O(\lg n)$ hops, the Manhattan distance can be reduced to within $(\lg n)^{\frac{2}{d}+1}$. ∎

***Lemma* 9** *Suppose that the Manhattan distance between the current node x and the target node t in the extended small-world model is $m < (\lg n)^{\frac{2}{d}+1}$. Then using Kleinberg's greedy algorithm can route the message to the target node t in $O(\lg n)$ expected number of hops.*

Proof: When $Dist(x,t) < (\lg n)^{\frac{2}{d}+1}$, Kleinberg's greedy algorithm is executed, that is, the message is forwarded to an immediate neighbor closest to t.

Since each AL-neighbor is chosen uniformly at random from nodes within $(\lg n)^{2/d}$ Manhattan distance from x, the probability for an AL-link to jump over $(\lg n)^{2/d}/2$ Manhattan distance is $1 - 2^{-d}$. Thus, after at most $O(1)$ expected number of hops, the Manhattan distance will be reduced by $(\lg n)^{2/d}/2$. Therefore, after at most $O(\lg n)$ expected number of hops, the Manhattan distance will be reduced to within $(\lg n)^{2/d}$. After that, using the simple greedy algorithm via the local T-links can route the message to the target node t in $(\lg n)^{2/d}$ hops. Since $d \geq 2$, $(\lg n)^{2/d} = O(\lg n)$. Therefore, after in total $O(\lg n) + (\lg n)^{2/d} = O(\lg n)$ expected number of hops, the message will reach the target node t. ∎

Combining Lemma 7 and Lemma 9 together, we obtain the proof of Theorem 1.

5 Experimental Evaluation

In this section, we will conduct experiments to evaluate our schemes and other existing routing schemes for Kleinberg's small-world networks.

We focus on the following four schemes: (a) The original greedy routing algorithm [9] in Kleinberg's small-world network with only one long-range contact per node. Each node forwards the message to its immediate neighbor closest to the destination; (b) The greedy routing algorithm in Kleiberg's small-world network with two long-range contacts per node [2, 11]. In the experimental study, we would like to learn how much the additional number of long-range links can help routing. (c) The decentralized routing scheme with $O(\lg n)$ local awareness [6, 13]. With this scheme, we intend to evaluate the degree at which the local awareness improve the routing efficiency. (d) Our near-optimal routing scheme in [17]; (e) Our optimal routing scheme in this paper. We note that most schemes have both non-oblivious and oblivious versions. Here we only focus on the non-oblivious version for each scheme.

5.1 Experimental Setup

Network Construction: We construct the small-world network based on a ring $[0, 1, \cdots, n]$. Each node i is connected to its immediate neighbors $(i+1) \mod n$. Let $H_n = \sum_{i=1}^{n} 1/i$ denote the harmonic normalization factor. We then generate a sequence of intervals $I(i)$, which we call the probability intervals, where $1 \leq i \leq n-1$. Let $0 < I_1 \leq 1/H_n$, and $\frac{1}{(i-1)H_n} < I_i \leq \frac{1}{iH_n}$, where $2 \leq i \leq n-1$. Each node i uniformly generates a random number x in $(0, 1]$, and then finds the interval that contains x. Suppose that x is located in the interval I_k. Node i then forms a long-range link connected to a node with the distance k. When each node has multiple long-range contacts, it just generates more than one random numbers, and sets up the connections in the same way.

In the extension of Kleinberg's small-world networks, each node uniformly and randomly chooses two nodes within the Manhattan distance $\lg^2 n$ as its augmented local neighbors.

Messages Generation and Evaluation Metrics: We let each node generate a query message with a random destination, and then evaluate the following metrics.

(1) **Average length of routing path** is the average number of hops travelled by the messages from the source to the destination.

(2) **Storage requirement** for each node is the number of information bits required to be stored on each node.

5.2 Experimental Results

We vary the number of nodes in the network from 5,000 to 25,000, and evaluate different routing schemes, as shown in Figures 1 and 2. For large n, the greedy algorithm with increasing number of long-range contacts [2, 11], the decentralized routing algorithm with local awareness [6, 13], our near-optimal and optimal algorithms all improve Kleinberg's original greedy algorithm. Our optimal algorithm performs the best and it reduces the latency (in terms of number of hops) by around 40%, but at the cost of increased amount of information stored on each node. On the other hand, our near-optimal scheme can find a shorter routing path than the decentralized routing schemes with local awareness [6, 13], while keeping the same storage space on each node.

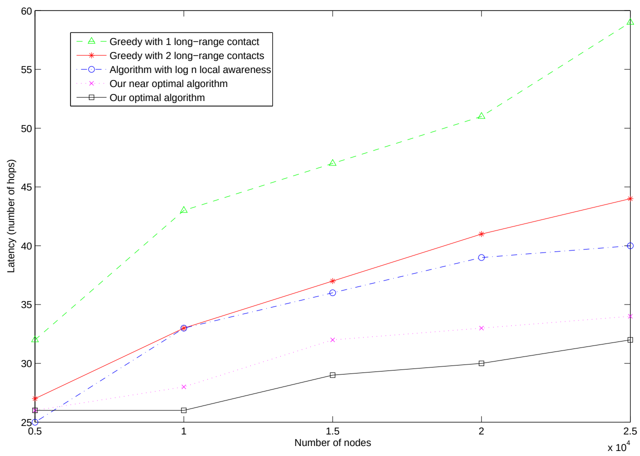

Figure 1: Average length of routing paths for different routing schemes.

6 Conclusion

We extend Kleinberg's small-world network with two more augmented local links, and show that if each node in the network is aware of $\beta \lg \lg n$ levels of augmented local neighborhood, where

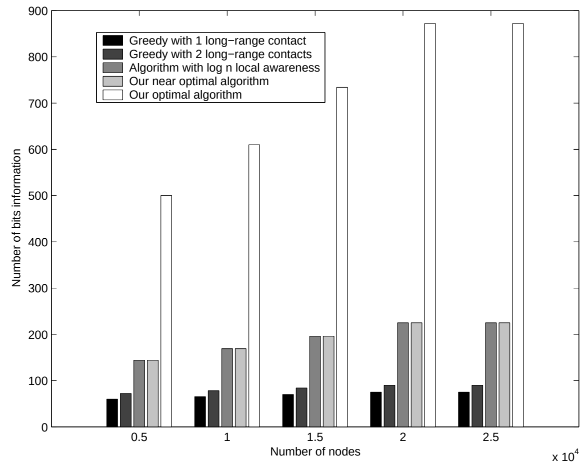

Figure 2: Storage requirement on each node for different routing schemes.

$1 < \beta < 2$, there exists an oblivious decentralized algorithm that can finish routing in $O(\lg n)$ expected number of hops, which is an optimal expected bound for routing.

In this paper, we only focus on the oblivious routing scheme. The non-oblivious algorithm can be easily obtained based on the design and analysis of our oblivious scheme.

References

[1] R. Albert, H. Jeong, and A.-L. Barabasi. The diameter of the World Wide Web. *Nature*, 401(9):130–131, 1999.

[2] J. Aspnes, Z. Diamadi, and G. Shah. Fault-tolerant routing in peer-to-peer systems. In *Proceedings of PODC'02*, 2002.

[3] B. Bollobás and F.R.K. Chung. The Diameter of a cycle plus a random matching. *SIAM Journal on Discrete Mathematics*, 1(3):328–333, 1988.

[4] L. Barriére, 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.

[5] D. Watts and S. Strogatz. Collective dynamics of small-world networks. *Nature*, 393:440–442, 1998.

[6] P. Fraigniaud, C. Gavoille, and C. Paul. Eclecticism shrinks even small worlds. In *PODC 2004*, 2004.

[7] C.M. Homan and G. Istrate. Small worlds, locality, and flooding on landscapes. Research Report TR-2003-796, Department of Computer Science, University of Rochester, USA, 2003.

[8] D. Kemper, J. Kleinberg, and A. Demers. Spatial gossip and resource location protocols. In *Proceedings of STOC*, 2001.

[9] J. Kleinberg. The Small-World Phenomenon: An Algorithmic Perspective. In *Proceedings of the 32nd ACM Symposium on Theory of Computing*, 2000.

[10] E. Lebhar and N. Schabanel. Almost optimal decentralized routing in long-range contact networks. In *ICALP 2004*, 2004.

[11] G. S. Manku, M. Bawa, and P. Raghavan. Symphony: Distributed hashing in a small world. In *Proceedings of the 4th USENIX Symposium on Internet Technologies and Systems*, 2003.

[12] G.S. Manku, M. Naor, and U. Wieder. Know thy neighbor's neighbor: The power of lookahead in randomized p2p networks. In *Proceedings of STOC 2004*, 2004.

[13] C. Martel and V. Nguyen. Analyzing Kleinberg's (and other) small-world models. In *PODC 2004*, 2004.

[14] M.E.J. Newman. Models of the small world. *J. Stat. Phys.*, 101, 2000.

[15] M.E.J. Newman and D.J. Watts. Scaling and percolation in the small-world network model. *Phys. Rev. E*, 60:7332–7342, 1999.

[16] S. Milgram. The small world problem. *Psychology Today*, 61, 1967.

[17] J. Zeng and W.-J. Hsu. Near Optimal Routing for Small-World Networks with Augmented Local Awareness. *arXiv cs.DM/0409017.*

[18] H. Zhang, A. Goel, and R. Govindan. Using the small-world model to improve Freenet performance. In *Proceedings of IEEE INFOCOM 2002*, 2002.