Optimal Routing in a Small-World Network

Jian-Yang Zeng[1] (曾坚阳) and Wen-Jing Hsu[1,2] (许文经)

[1] School of Computer Engineering, Nanyang Technological University, Singapore 639798

[2] Computer Science Program, Singapore-MIT Alliance (SMA) Program, Singapore 117576

E-mail: zengjy@gmail.com; hsu@ntu.edu.sg

Revised May 15, 2006.

Abstract Substantial research has been devoted to the modelling of the small-world phenomenon that arises in nature as well as human society. Earlier work has focused on the static properties of various small-world models. To examine the routing aspects, Kleinberg proposes a model based on a d-dimensional toroidal lattice with long-range links chosen at random according to the d-harmonic distribution. Kleinberg shows that, by using only local information, the greedy routing algorithm performs in $O(\lg^2 n)$ expected number of hops. We extend Kleinberg's small-world model by allowing each node x to have two more random links to nodes chosen uniformly and randomly within $(\lg n)^{\frac{2}{d}}$ Manhattan distance from x. 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. Our routing algorithm keeps only $O((\lg n)^{\beta+1})$ bits of information on each node, where $1 < \beta < 2$, thus being scalable w.r.t. 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.

Keywords small-world model, augmented local awareness, decentralized routing, analysis of algorithms, distributed systems

1 Introduction

In [1] Milgram shows that there is a *small-world phenomenon* in the human society, namely, any two persons in the world can be connected by a chain of six acquaintances on average, and hence people can relay messages efficiently to any target via common acquaintances. The small-world phenomenon has also been shown to be pervasive in networks from nature and engineered systems, such as the World Wide Web[2,3], peer-to-peer systems[4–7], etc.

A number of network models has been proposed to study the small-world properties, see, e.g., [3, 8–10]. Newman and Watts[9] 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 edges to nodes uniformly distributed on a ring, where the nodes are connected with their near neighbors. However, Newman and Watts's model does not address the routing issues of small-world networks[10]. The poly-logarithmic diameter of some graphs does not imply the existence of efficient distributed routing algorithms[10]. For example, the random graph in [11] yields a logarithmic diameter, yet any routing algorithm equipped with only local information requires at least $\sqrt{n}$ expected number of hops (where n is the size of the network)[10].

In order to examine the routing issue, Kleinberg[10] develops a new model based on a d-dimensional torus with long-range links chosen randomly from the d-harmonic distribution, that is, a long-range link exists between nodes u and v with probability $\Theta(Dist(u,v)^{-d})$, where $Dist(u,v)$ denotes the Manhattan distance between u and v. Based on this model, Kleinberg shows that a simple greedy routing algorithm using only local information can route messages between any two nodes in $O(\lg^2 n)$. The symbol lg denotes base-2 logarithm. Also, we remove the ceiling or floor for simplicity throughout the paper. This bound is tightened to $\Theta(\lg^2 n)$ later by Barrière *et al.*[12] and Martel *et al.*[13], respectively. Further research[6,13–15] shows that in fact the $O(\lg^2 n)$ bound of the original greedy routing algorithm can be improved by injecting more information to the routing message. Manku *et al.*[6] show that if each message holder at a routing step takes its neighbors' neighbors into account when making routing decisions, the routing complexity can be improved to $O\left(\frac{\lg^2 n}{q \lg q}\right)$, where q denotes the number of long-range contacts for each node. Lebhar and Schabanel[14] propose a routing algorithm for 1-dimensional Kleinberg's model; they show that a routing path with expected length of $O\left(\frac{\lg n(\lg\lg n)^2}{\lg^2(1+q)}\right)$ can be found. Two research groups, namely, Fraigniaud *et al.*[15] 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 additional state information during routing, while [15] uses an oblivious greedy routing algorithm. In [13], Martel and Nguyen show that the expected diameter of a d-dimensional Kleinberg network is $\Theta(\lg n)$. However, it is so far unresolved whether the routing complexity can be matched, which motivates our work.

There are normally two approaches for decentralized routing: oblivious and non-oblivious schemes[15]. A routing algorithm is *oblivious* if the message holder

Short Paper

makes routing decisions only based on its own routing table and the target node. A routing algorithm is said to be *non-oblivious* if the routing decisions of the message holder also depend on the routing history stored in the message header. The scheme in [15] is oblivious, while the schemes in [13, 14] are non-oblivious. We will only consider the oblivious routing scheme.

In [16], 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 is augmented with 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 optimal.

Potential applications of the small-world model in computer networks include efficient lookup in peer-to-peer systems[4−7], gossip protocol in a communication network[17], flood routing in ad-hoc networks[18], etc.

The rest of the paper is organized as follows. Section 2 introduces the augmented small-world model and the decentralized routing algorithm. Section 3 lists our main results and contributions. Section 4 analyzes the complexity of decentralized routing algorithm. Section 5 briefly concludes this paper.

2 Small-World Model and Decentralized Routing

Our small-world model is an extension of Kleinberg's d-dimensional model[10]. It is based on a d-dimensional torus $[n]_d = \{0, 1, \ldots, n\}^d$ with three extra links for each node, where $d \geqslant 2$. Firstly, as with Kleinberg's original model[10], each node has a long-range link to another node chosen randomly according to the d-harmonic distribution, i.e., 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 the **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 the **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 the nodes linked by u, including its K-neighbor, AL-neighbor and T-neighbor, collectively 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* interchangeably.

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 $A_x(i) = \bigcup_{j \leqslant i} \Gamma_x(j)$ as the first i levels of AL neighborhood for node x. For a given level i of AL neighborhood, $A_x(i-1)$ is said to be the set of *previously known nodes*. The set $L_x(i) = A_x(i) - A_x(i-1)$ denotes the *new* nodes discovered during the ith level of AL neighborhood. We will call $A_x(k)$ the *AL awareness* of node x, where each node in our extended small-world model is aware of the first k levels of its AL neighborhood.

The description of our oblivious routing algorithm is given in Algorithm 1 in Fig.1.

Algorithm 1
Input: source s and target t.
Initialization: $x \leftarrow s$.
The first phase: $Dist(x, t) \geqslant (\lg n)^{\frac{2}{d}+1}$.

```
while Dist(x,t) ≥ (lg n)^(2/d+1) do
  node x checks 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 specified
    later. Let Z_x denote the set of such nodes z.
  if Z_x is empty then
    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 = x_0, 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 node x itself then
    The message is routed to node x's 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: $Dist(x, t) < (\lg n)^{\frac{2}{d}+1}$
The message is forwarded to an immediate neighbor closest to the target node t, until it reaches t.

Fig.1. Our oblivious routing algorithm (Algorithm 1).

3 Our Contributions

Our main result is 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 $\Omega(\lg n)$ expected path length, the $O(\lg n)$ expected bound of routing complexity in Theorem 1 is optimal. In addition, since the size of neighborhood $|A_x(\beta \lg\lg n)| = 1 + 2 + \ldots + 2^{\beta \lg\lg n} < 2\lg^{\beta} n$, and the address of each node is represented in a string of $O(\lg n)$ bits, 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 stored on each node in the small-world networks. A comparison of our scheme with the other existing results is shown in Fig.2.

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 $1 < \beta < 2$, then it is aware of at least $(\lg n)^{\delta+1}$ different nodes with certain probability, where $0 < \delta < \beta - 1$. Based on this result, in Lemma 3 we show that the AL awareness of each node is very likely to contain a K-neighbor that is close to the target node. Then in Lemmas 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 [4, 10] is useful for our subsequent analysis.

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

We first quantify the size of 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*

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

where $\xi > 0$.

Proof. The proof is divided into two parts. In Part 1, we show that with probability at least $1 - 8(\lg n)^{-\frac{4}{3}}$, $|L_x(\frac{1}{3}\lg\lg n)| \geqslant (\lg n)^{\frac{1}{3}}$, that is, every AL-link points to a new node during the first $\frac{1}{3}\lg\lg n$ levels of x's AL neighborhood. In Part 2, by using Chernoff's bound, we show that $|L_x(i)|$ still increases at an exponential rate for all $\frac{1}{3}\lg\lg n < i < \beta \lg\lg n$, based on which we have $|L_x(\beta \lg\lg n)| \geqslant (\lg n)^{1+\delta}$ 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 an AL-link points to a previously known node is so small that all AL-links tend to point to new nodes, and hence $|L_x(\frac{1}{3}\lg\lg n)| \geqslant (\lg n)^{1/3}$ w.h.p..

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

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

Thus, we have $|A_x(i)| \leqslant |A_x(\frac{1}{3}\lg\lg n)| < 2(\lg n)^{1/3}$ for all $0 \leqslant i \leqslant \frac{1}{3}\lg\lg n$. 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 $|A_x(i)| \leqslant 2(\lg n)^{1/3}$, 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

Scheme	#bits of awareness on each node	#hops expected	Oblivious or Non-oblivious?
Kleinberg's greedy[4,10,12]	$O(q \lg n)$	$O(\lg^2 n/q)$	Oblivious
Non-greedy[6]	$O(q^2 \lg n)$	$O(\lg^2 n/(q \lg q))$	Non-oblivious
Decentralized algorithm in [14]	$O(\lg^2 n/\lg(1+q))$	$O((\lg n)^2/\lg^2(1+q))$	Non-oblivious
Decentralized algorithm[13]	$O(\lg^2 n)$	$O((\lg n)^{1+1/d})$	Non-oblivious
Indirect-greedy algorithm[15]	$O(\lg^2 n)$	$O((\lg n)^{1+1/d})$	Oblivious
Near optimal algorithm for one-dimensional model with augmented awareness [16]	$O(\lg^2 n)$	$O(\lg n \ \lg\lg n)$	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

Fig.2. Comparisons of our decentralized routing algorithms with the other existing schemes. In the first three schemes (in [4, 6, 10, 12, 14]), we suppose that each node has q K-links, while in the next four schemes (in [13, 15, 16] and this paper), we suppose that each node has one K-link.

$2|A_x(\frac{1}{3}\lg\lg n)| \leqslant 4(\lg n)^{1/3}$ number of AL-links, so all AL-links point to new nodes with probability at least $(1-2(\lg n)^{-5/3})^{4(\lg n)^{1/3}} \geqslant 1-8(\lg n)^{-4/3}$ for sufficiently large n. Here we use the fact $(1+x)^a \geqslant 1+ax$ for $x > -1$ and $a \geqslant 1$. Thus, we have $\Pr[|L_x(\frac{1}{3}\lg\lg n)| \geqslant (\lg n)^{1/3}] \geqslant 1 - 8(\lg n)^{-4/3}$.

Part 2. Let $\mathcal{B}$ denote the event that $|L_x(\frac{1}{3}\lg\lg n)| \geqslant (\lg n)^{1/3}$. From Part 1, we have $\Pr[\mathcal{B}] \geqslant 1 - 8(\lg n)^{-4/3}$. Next, we will consider the sequence of $L_x(i)$ for $\frac{1}{3}\lg\lg n \leqslant i \leqslant \beta\lg\lg n$. We assume that $|A_x(i)| \leqslant (\lg n)^{1+\delta}$ for all $\frac{1}{3}\lg\lg n \leqslant i \leqslant \beta\lg\lg n$, otherwise the lemma holds true. Since each AL-link is connected to a node chosen uniformly and randomly 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 $\frac{1}{3}\lg\lg n \leqslant i \leqslant \beta\lg\lg n$. By Chernoff's bound, there exists a constant $0 < \epsilon < 1$ such that

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

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

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

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

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

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

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 $n > n_0$, $\frac{1}{3} + (\beta - \frac{1}{3})\lg\big(2((1-(\lg n)^{\delta-1})(1-\epsilon)\big) \geqslant 1+\delta$. Thus, there exists a constant $0 < \delta < \beta - 1$ such that $\Pr[|A_x(\beta\lg\lg n)| \geqslant (\lg n)^{1+\delta}] > \Pr[|L_x(\beta\lg\lg n)| \geqslant (\lg n)^{1+\delta}] > 1 - \frac{1}{\lg^{\xi} n}$, where $\xi > 0$. Therefore, the lemma is verified. □

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 $m \geqslant (\lg n)^{\frac{2}{d}+1}$. 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 $|A_x(\beta\lg\lg n)| \geqslant (\lg n)^{1+\delta}$. 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 $\Pr[u \rightarrow^K D_l(t)]$ 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 $Dist(x,t) = m \geqslant (\lg n)^{\frac{2}{d}+1}$, 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 $(m + \frac{m}{\lg^{\tau} n} + \beta\lg\lg n(\lg n)^{2/d}) \leqslant 2m$. 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

$$\begin{aligned}\Pr[u \rightarrow^K D_{\frac{m}{\lg^{\tau} n}}(t)] \geqslant & |D_{\frac{m}{\lg^{\tau} n}}(t)| \cdot \frac{c_1}{(2m)^d\lg n} \\ = & \Big(\frac{m}{\lg^{\tau} n}\Big)^d \cdot \frac{c_1}{(2m)^d\lg n} = \frac{c_3}{(\lg n)^{\tau d+1}},\end{aligned}$$

where $c_3 = c_1/2^d$ denotes a constant.

Let $\Pr[A_x(\beta\lg\lg n) \rightarrow^K D_{\frac{m}{\lg^{\tau} n}}(t)]$ 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

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

for a constant $\zeta > 0$. □

By using a similar technique as that in Lemma 3, we can obtain 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 $m \geqslant (\lg n)^{\frac{2}{d}+1}$. 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 $m \geqslant (\lg n)^{\frac{2}{d}+1}$. 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 $m \geqslant (\lg n)^{\frac{2}{d}+1}$. 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 $m \geqslant (\lg n)^{2/d+1}$, we have $(m-\beta(\lg n)^{2/d}\lg\lg n)^d \geqslant m^d - m^{d-1}(\beta(\lg n)^{2/d}\lg\lg n)$. Thus, we can obtain the following upper bound for $|I_1-I_2|$

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

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 $Dist(y_1,y_2) \geqslant m/2$ if $m \geqslant (\lg n)^{\frac{2}{d}+1}$. Since $|A_x(\beta\lg\lg n)| \leqslant 2\lg^{\beta} n$, we have

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

Thus, the lemma is verified. □

Lemma 7. *Suppose that the Manhattan distance between the current node x and the target node t in the extended small-world model is $m \geqslant (\lg n)^{\frac{2}{d}+1}$. 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. As 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 Lemmas 6 and 4, we have $\Pr[\mathcal{F}_1] \leqslant \frac{1}{\lg^{\varphi} n}$ and $\Pr[\mathcal{F}_2] \geqslant 1-\dfrac{1}{\lg^{\zeta} n}$ 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 $A_x(\lg\lg n)$ 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 $\pi: x_0=x, x_1,\ldots,x_t=z$ 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) it has a K-neighbor within $\frac{m}{\lg^{\tau} n}$ to the target node; 2) it is closest to node x in terms of AL-links.

We first consider the case where no tie arises, 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 $\pi: x_1,x_2,\ldots,x_t=z$ 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 the shortest paths. 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 $m \geqslant (\lg n)^{\frac{2}{d}+1}$. 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 $\big[\frac{m}{(\lg n)^{\tau i}}, \frac{m}{(\lg n)^{\tau(i-1)}}\big)$ Manhattan distance from t. Since the maximum Manhattan distance is n, there are at most $O\big(\frac{\lg n}{\lg\lg n}\big)$ phases. Because each phase takes at most $O(\lg\lg n)$ expected number of steps by Lemma 7, after $O\big(\frac{\lg n}{\lg\lg n}\big)\cdot O(\lg\lg n) = O(\lg n)$ 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. According to the time complexity of Kleinberg's routing algorithm, it takes $O(\lg^2 m)$ steps to route a message to a target node within m Manhattan distance. Thus, the final routing phase using Kleinberg's original algorithm takes at most $O(\lg^2 m) = O((\lg\lg^{\frac{2}{d}+1} n)^2) = O(\lg n)$ steps. Therefore, the lemma follows. □

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

5 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 $1 < \beta < 2$, there exists an oblivious decentralized algorithm that can finish routing in $O(\lg n)$ expected number of hops, which is optimal.

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[5]. Since the links in the model are randomized, our extended network is less vulnerable to adversarial attacks, and thus provides good fault tolerance.

Here we only focus on the oblivious routing scheme. A non-oblivious algorithm can be easily obtained based on the design and analysis of our oblivious scheme.

References

[1] Milgram S. The small world problem. *Psychology Today*, 1967, 61(2): 60–67.

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

[3] Watts D, Strogatz S. Collective dynamics of small-world networks. *Nature*, 1998, 393(6684): 440–442.

[4] Aspnes J *et al.* Fault-tolerant routing in peer-to-peer systems. In *Proc. The 21st Annual Symp. Principles of Distributed Computing*, Monterey, California, 2002, pp.223–232.

[5] Manku G S, Bawa M, Raghavan P. Symphony: Distributed hashing in a small world. In *Proc. The 4th USENIX Symposium on Internet Technologies and Systems*, Seattle, Washington, 2003, pp.127–140.

[6] Manku G S, Naor M, Wieder U. Know thy neighbor's neighbor: The power of lookahead in randomized p2p networks. In *Proc. The 36th Annual ACM Symposium on Theory of Computing*, Chicago, IL, USA, 2004, pp.54–63.

[7] Zhang H, Goel A, Govindan R. Using the small-world model to improve Freenet performance. In *Proc. The 21st Annual Joint Conference of the IEEE Computer and Communications Societies*, New York, USA, 2002, pp.1228–1237.

[8] Newman M E J. Models of the small world. *Journal of Statistical Physics*, 2000, 101(3-4): 819–841.

[9] Newman M E J, Watts D J. Scaling and percolation in the small-world network model. *Physical Review E*, 1999, 60(6): 7332–7342.

[10] Kleinberg J. The small-world phenomenon: An algorithmic perspective. In *Proc. The 32nd ACM Symposium on Theory of Computing*, Portland, Oregon, 2000, pp.163–170.

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

[12] Barriére L, Fraigniaud P, Kranakis E, Krizanc D. Efficient routing in networks with long range contacts. In *Proc. The 15th International Symposium on Distributed Computing*, Lisbon, Portugal, 2001, pp.270–284.

[13] Martel C, Nguyen V. Analyzing Kleinberg's (and other) small-world models. In *Proc. The 23rd Annual ACM Symposium on Principles of Distributed Computing*, St. John's, Newfoundland, Canada, 2004, pp.179–188.

[14] Lebhar E, Schabanel N. Almost optimal decentralized routing in long-range contact networks. In *Proc. The 31st International Colloquium on Automata, Languages and Programming*, Turku, Finland, 2004, pp.894–905.

[15] Fraigniaud P, Gavoille C, Paul C. Eclecticism shrinks even small worlds. In *Proc. The 23rd Annual Symposium on Principles of Distributed Computing*, St. John's, Newfoundland, Canada, 2004, pp.169–178.

[16] Zeng J, Hsu W J, Wang J. Near optimal routing for small-world networks with augmented local awareness. In *Proc. The 3rd Int. Symp. Parallel and Distributed Processing and Applications*, Nanjing, China, 2005, pp.503–513.

[17] Kemper D, Kleinberg J, Demers A. Spatial gossip and resource location protocols. In *Proc. The 33rd Annual ACM Symp. Theory of Computing*, Hersonissos, Greece, 2001, pp.163–172.

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