Improved Compact Routing Schemes for Power-Law Networks

Mingdong Tang, Jianxun Liu, Guoqing Zhang

HAL Id: hal-01593024
https://hal.inria.fr/hal-01593024

Submitted on 25 Sep 2017

Improved Compact Routing Schemes
for Power-Law Networks

Mingdong Tang[1,2], Jianxun Liu[1], Guoqing Zhang[2]

[1]Key Laboratory of Knowledge Processing and Networked Manufacturing,
Hunan University of Science and Technology, Xiangtan 411201, China
[2]Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100093, China
{tangmingdong, gqzhang}@ict.ac.cn

Abstract. Compact routing intends to achieve good tradeoff between the routing path length and the memory overhead, and is recently considered as a main alternative to overcome the fundamental scaling problems of the Internet routing system. Plenty of studies have been conducted on compact routing, and quite a few universal compact routing schemes have been designed for arbitrary network topologies. However, it is generally believed that specialized compact routing schemes for peculiar network topologies can have better performance than universal ones. Complex network research has uncovered that most real-world networks have degree distributions exhibiting power law tails, i.e., a few nodes have very high degrees while many other nodes have low degrees. High-degree nodes play a crucial role of hubs in communication and networking. Based on this fact, we put forward two highest-degree landmark based compact routing schemes, namely HDLR and HDLR$^+$. Theoretical analysis on random power law graphs shows our schemes can achieve a better space-stretch trade-off than prior compact routing schemes. Simulations conducted on random power law graphs and real-world AS-level Internet graph validate the effectivity of our schemes.

Keywords: compact routing, power law, stretch, routing table size, algorithm

1 Introduction

Routing is a basic function of a network on which many upper-layer applications rely. For the efficient use of network resources (link capacities, etc.) it is desirable to design schemes which route along paths as short as possible. This is measured by the metric called stretch, which can be represented by (α, β), meaning the path between any two nodes u, v is no longer than $\alpha d(u,v) + \beta$, where $d(u,v)$ is the length of the shortest path between u and v. A scheme with stretch $(\alpha, 0)$ is said to have multiplicative stretch, whereas one with stretch $(1, \beta)$ is said to have additive stretch β. In addition to stretch, with the exponential growth of network size, the routing table size also grows rapidly and becomes a bottleneck of the network, so, it becomes equally important to reduce the routing table size.

However, these two objectives are contradictive in that optimizing one will typically deteriorate the other. At one extreme, the shortest path routing ensures the minimum stretch (i.e. with a multiplicative stretch of 1) but requires a $O(n \log n)$ table size at each node (one entry for every node in the network). At the other extreme, the flooding scheme does not need to maintain any state information at each node, but has a worst-case multiplicative stretch $O(n)$ for general graphs. Both of these two solutions do not scale well. It is evident that there is a tradeoff between the space requirement and the stretch. It is desirable, therefore, to develop schemes with good tradeoff between them. Such schemes typically have the property of near-optimal stretch but with substantially smaller routing tables.

The routing schemes that balance between routing table size and stretch are usually called compact routing schemes. Compact routing schemes on arbitrary network topologies (also termed as universal schemes) [1-3] and specific networks such as trees [4-5], planar graphs[6], growth bounded graphs [7] and graphs with low doubling dimension [8-9] have attracted wide research attentions. Presently, the most optimized universal compact routing scheme could restrict the routing table size on each node to $\tilde{O}(n^{1/2})$ [1]bits, and meanwhile, achieve a multiplicative stretch of 3, i.e., the length of the path between any node pair taken by the routing is no more than 3 times the length of the shortest path between the node pair. Although it has been proved that no routing scheme can achieve multiplicative stretch less than 3 with sub-linear space for generic networks [3], it is generally believed that routing schemes specialized for peculiar network topologies can achieve better performance than universal ones. This is because, unlike the TZ scheme, specialized routing schemes can exploit the topological properties of special networks.

Recent researches on complex networks have uncovered that many real-world networks, such as the Internet [10], P2P [11], WWW [12], have scale-free topologies, i.e., the degree distribution follows a power law. Compact routing on power law networks has attracted increasingly more attention [13-17]. Many researchers even consider compact routing as a promising alternative to overcome the scaling limitations of the Internet routing system [18-20]. Although previous work has experimentally showed that specialized routing schemes can outperform the optimal universal routing scheme on power law networks, they seldom give theoretical analysis and guaranteed performance on stretch and routing table size.

In this paper, we focus on compact routing in power law graphs, and present two high-degree landmark based routing schemes, namely HDLR and HDLR[+]. Based on the RPLG (Random Power Law Graph) model [21-22], theoretical analysis shows: 1) HDLR can achieve a multiplicative stretch of 3 and average routing table size of $\tilde{O}(n^{1/3})$ with high probability. 2) HDLR[+] can achieve a stretch of $\min\{(3,0),(2,D_A)\}$ and average routing table size of $\tilde{O}(n^{1/3})$ with high probability, where n is the number of nodes in the graph, and D_A is the diameter of the subgraph induced by the landmark set, bounded by some constant with high probability. Simulations conducted on both random power law graphs and real-world AS graph validate our theoretical results.

---

[1] $\tilde{O}()$ notation indicates complexity similar to $O()$ notation up to poly-logarithmic factors.

2   Related Work

Since 1990s, several excellent compact routing schemes are proposed for generic networks [1-3][23], among which the Cowen scheme [2] and the TZ scheme [1] are two prominently referenced ones. Cowen scheme is the first universal compact routing scheme that reduces the stretch to (3,0), and produces sublinear routing table size whose upper bound is $\tilde{O}(n^{2/3})$. TZ scheme improves the Cowen scheme's routing table size upper bound to $\tilde{O}(n^{1/2})$. Both these two schemes are landmark-based and differ only in how landmarks are selected. Cowen scheme adopts a greedy algorithm to find the dominating set of the graph, and obtains the landmark set based on this dominating set, whereas the TZ scheme relies on random techniques for landmark selection. Since Gavoille and Gengler [3] showed that no universal routing scheme with a multiplicative stretch < 3 can produce $o(n)$ routing table size, and Thorup and Zwick [24] proved that any routing scheme with multiplicative stretch strictly below 5 cannot guarantee space smaller than $\tilde{O}(n^{1/2})$, it thus follows that the TZ scheme is nearly optimal as far as the stretch and routing table size are concerned.

All the above schemes are designed for generic networks. However, recent studies revealed that most real networks, for example, the Internet, WWW and unstructured P2P networks, can all be categorized as power-law networks [25-28], i.e., the degree distribution $p(k)$ satisfies $p(k) \sim k^{-\gamma}$, where γ is called the power-law exponent, typically between 2 and 3. It has been shown that on power-law networks, high degree nodes play the hub role for network connectivity and routing [16][21], most nodes are directly connected to these nodes or separated by only a few hops, and the path between these high degree nodes are typically very small[29]. So recently, compact routing schemes on power-law networks have attracted much attention. Krioukov [13] first evaluated the performance of the TZ scheme on Internet-like power-law graphs, and found that the performance is far better than the worst case theoretic results, e.g., the average multiplicative stretch is only 1.1. Brady and Cowen [14] introduced the concept of additive stretch, and proposed a compact routing scheme with additive stretch (1,d) and logarithmic routing table scaling, i.e. $O(e\log^2 n)$. Simulations showed that d and e can both take small values on power-law networks, and the average stretch is lower than the TZ scheme. Other specialized compact routing schemes on power-law networks can be found in recent works [15-17].

Although the aforementioned work is effective, none of them gave theoretical bounds on stretch and routing table size. Most recently, Chen et al. [28] proposes a compact routing scheme, which is the first routing scheme guarantees a better space-stretch tradeoff on power law graphs than universal ones. The theoretical analysis is based on the RPLG (Random Power Law Graph) model [21-22], and shows the routing table size can be reduced to $\tilde{O}(n^{1/3})$ or less using address length of $O(\log n \log \log n)$ bits and stretch of (3,0). However, Ref. [28] makes no progress in the worst-case stretch compared to the universal optimal scheme. Our proposed scheme HDLR[+] is the first compact routing scheme that achieves a better stretch than (3,0) on power law networks. Furthermore, HDLR and HDLR[+] are superior to the scheme by Ref. [28] by using addresses and packet headers both with $O(\log n)$ bits.

3 Definitions and Notations

We model a communication network as an undirected, unweighted graph $G=(V,E)$, with vertices representing routers and edges representing communication links. For convenience, we use the terms vertex and node, edge and link interchangeably throughout this paper. The distance between any two nodes $u, v \in V$, is defined as the number of edges on the shortest path between them, denoted by $d(u,v)$.

A routing scheme R is a distributed algorithm for packet delivery between any two nodes in the network. R can only be called compact if it has routing tables with sizes sublinear in n and message header sizes polylogarithmic in n. There are two classes of compact routing schemes: Name-dependent schemes are allowed to add labels to node addresses to encode useful information for routing purposes, where each label has length at most polylogarithmic in n. Name-independent schemes do not allow renaming of node addresses, instead they must function with all possible addresses. We consider name dependent routing in this paper. Denote the length of the path traversed by a packet from u to v according to a routing scheme R by $d_R(u,v)$.

The quality of a routing scheme is primarily measured by two metrics: the storage space of the distributed data structures established by the preprocessing procedure, typically referring to the size of routing table; and the routing efficiency, measured by the stretch. Other metrics include packet header length, routing table convergence time and communication cost. In this paper, we only consider routing on static topologies, hence the convergence time and communication cost are out of our consideration.

4 Our Routing Schemes

This section introduces the design of our routing schemes—HDLR and HDLR[+], which are both landmark-based schemes. The basic idea of landmark based routing is to select a small proportion of nodes from the network as landmarks. The routing table of each node only stores routing entries for the landmarks and its vicinity. As a result, suppose a node is sending a packet to some destination, if the destination node could be found in the source node's routing entries, then the packet can be delivered along the shortest path; otherwise, the packet is first sent towards the nearest landmark of the destination, and then will be delivered to the target by the landmark.

HDLR is adapted from the TZ scheme by using a highest degree landmark selection method instead of a random one, and HDLR[+] extends HDLR by optimizing its stretch.

4.1 HDLR

Let $A \subseteq V$ represent the landmark set. For any $u \in V$, define u's ball $B(u)$ to be $B(u) = \{v \in V | d(u, v) < d(u, A)\}$, where $d(u, A)$ is the minimum shortest path length among all the shortest paths from u to all the landmarks in A, and define u's cluster $C(u)$ to be

$C(u) = \{v \in V | u \in B(v)\}$. Take the graph in Fig.1 as an example, Table 1 illustrates these definitions.
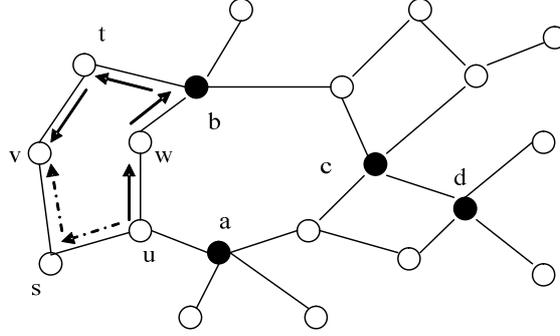

Fig.1. An example of the HDLR scheme. Filled nodes are landmarks. It illustrates the forwarding path form *u* to *v* taken by the HDLR scheme. Solid lines are the actual path used by HDLR scheme, and the dashed lines compose the shortest

Tabel 1. Balls and clusters of the nodes in Fig.1

	s	*t*	*u*	*v*	*w*
B()	{u,v}	φ	φ	{s,t}	φ
C()	{v}	{v}	{s}	{s}	φ

In HDLR, The preprocessing procedure consists of the following steps:
1) landmark selection: sort the nodes in decreasing order of their degrees, and choose the first n^x nodes as the landmarks, where $0 < x < 1$ is a tunable parameter;
2) address assignment: for each non landmark $u \in V - A$, assign $(u, A(u), e_{A(u)}(u))$ as its address, denoted as *Address*(u), where $A(u)$ is the name of the nearest landmark to *u*, and $e_{A(u)}(u)$ is the identifier of the port at $A(u)$ that lies on the shortest path from $A(u)$ to *u*; for each landmark, its address is identical to its name.
3) routing table setup: for each non landmark $u \in V - A$, its routing table is *Table* (u) $= \{(v, e_u(v)) | v \in A \cup C(u)\}$. In other words, *u* stores the routing entries to nodes in A and $C(u)$, where $e_u(v)$ is the identifier of the port at *u* that lies on the shortest path from *u* to *v*; for each landmark $l \in A$, its routing table only contains the routing information to other landmarks, i.e., *Table*(l) = $\{(v, e_l(v)) | v \in A - \{l\}\}$.

We can see the routing table at each node is much smaller than generic shortest path routing schemes. For example, in Fig. 1, the routing tables of *u* contain only 5 routing entries for {*a, b, c, d, s*}, as a contrast to 19 entries for all other nodes if a shortest path routing scheme is adopted.

The packet forwarding procedure of HDLR is very similar to the TZ scheme. Assuming a packet with a destination address $(v, A(v), e_{A(v)}(v))$ arrives at node u, u makes the following decision:

1) If $u = v$, then the packet reaches its destination;
2) otherwise, if $u = A(v)$, then forward the packet using the port $e_{A(v)}(v)$ that found in the packet header;
3) otherwise, if $Table(u)$ contains the routing entry for v, forward the packet using the port $e_u(v)$;
4) otherwise, forward the packet to $A(v)$ using the port $e_u(A(v))$ found in $Table(u)$.

The above forwarding algorithm is likely to use non-shortest paths, with a worst-case stretch of (3,0), as will been proved in Section 5. For example, in Fig. 1, when sending a packet from u to v, the actual forwarding path would be $u_- w_- b_- t_- v$, but the shortest path between u and v is $u_- s_- v$.

4.2   HDLR$^+$

The HDLR$^+$ scheme is an extension of HDLR, with the purpose to optimize the worst-case stretch for compact routing in power law graphs. Based on HDLR, we store limited more information at each node u, by adding entries for nodes in $B(u)$ to $Table(u)$. Therefore, for any non-landmark node u, $Table (u) = \{(v, e_u(v)) \mid v \in A \cup B(u) \cup C(u)\}$. While for each landmark l, its routing table is still only composed by entries for nodes in A, as both $B(l)$ and $C(l)$ are empty. By this means, as we will show through analysis in Section 5, the worst-case stretch of HDLR$^+$ is superior to HDLR.

Except the routing table setting, the other aspects of HDLR$^+$ such as the landmark selection and address assignment scheme, and the forwarding algorithm are identical to HDLR. Therefore, we omit unnecessary details of HDLR$^+$.

5   Analytical Performance Analysis

In this section, we use the random power law graph theory [21-22] to analyze the space-stretch tradeoff of the proposed two schemes, and present the bounds for address length, stretch and routing table size.

5.1   Theory of Random Power-Law Graphs

The random power-law graph model extends the generic random graph theory for power-law networks.

Definition 1(Random Power Law Graph [21,22]). Let $G(\mathbf{k})$ represent the ensemble of all random graphs satisfying the expected degree sequence $\mathbf{k} = (k_1, k_2, ..., k_n)$, where k_i is the expected degree of node v_i. The probability that there is an edge connecting two

nodes v_i and v_j is given by $k_i k_j / \sum_{z=1}^{n} k_z$. If $\mathbf{k} = (k_1, k_2, ..., k_n)$ follows power-law distribution, then $G \in G(\mathbf{k})$ is called a random power-law graph.

An important reason to work with this random graph model is that the edges are independent. This independence makes several graph properties easier to analyze. Actually, there are already many analytical results for the topological properties of random power-law graphs, in the following we introduce some useful ones, which offer basic proofs for the analysis in this section.

Let $S_{G(\mathbf{k})}^{k}$ denote the set of nodes in $G(\mathbf{k})$ whose expected degrees are no less than k, then the following property holds for the subgraph induced by $S_{G(\mathbf{k})}^{k}$.

Lemma 1. Suppose $G \in G(\mathbf{k})$ is a power-law graph instance with the power-law exponent satisfying $2 < \gamma < 3$. Given $t \geq n^{\delta}$ ($0 < \delta < \frac{3-\gamma}{\gamma-2}$), then the diameter of the subgraph induced by $S_{G(\mathbf{k})}^{t}$ has probability $1 - n^{-2}$ not exceeding $(1 + o(1)) \frac{\log n}{(3-\gamma)\log t}$.

Proof. Readers can refer to Ref. [21], page 62, Lemma 12, for the proof of this Lemma. □

From Lemma 1, we can see that the subgraphs induced by high-degree vertices of power law graphs usually have very small diameters, approaching some constant when the graph size increases.

Lemma 2. Let $\beta = \frac{\gamma-2}{2\gamma-3} + \varepsilon$, $\beta' = \frac{1-\beta}{\gamma-1}$, and $\alpha = \beta'(\gamma-2) + \frac{(2\gamma-3)\varepsilon}{\gamma-1}$, where ε is any positive real number that satisfies $n^{\frac{(2\gamma-3)\varepsilon}{\gamma-1}} \geq \frac{2(\gamma-1)}{\gamma-2} \ln n$ (apparently, when n is sufficiently large, ε could be very small). Let G is an instance of $G(\mathbf{k})$, then for all $u \in V(G)$, the following property holds with probability at least $1 - 3n^{-2}$: $|B(u)| = |\{u \in V(G): d(u, v) < d(u, S_{G(\mathbf{k})}^{\tau})\}| = O(n^{\alpha})$, where $\tau = n^{\beta'}$.

Proof. Please refer to Ref. [28], Lemma 7, for the proof of this Lemma. □

Lemma 2 indicates that for power law graphs, if high-degree nodes are selected as landmarks, the ball size of any node u is high likely to has an upper bound much smaller than the size of the graph. This might be partly caused by the reason that ordinary nodes in power law graphs usually have very small hop distances to high-degree nodes, which limits the range of neighborhood expansions and thus ball sizes.

5.2 Analysis

In the following, we analyse the routing performance of HDLR and HDLR[+] on random power law graphs, using the metrics such as address length, stretch and routing table size.

The address length of HDLR or HDLR[+] is $O(\log n)$ bits, which arises from the

following observations. For any node v, its address is $(v, L(v), e_{A(v)}(v))$. Since the node name and interface identifier can both be represented using $O(\log n)$ bits, so the address length is $O(\log n)$. Because each packet of HDLR or HDLR$^+$ only includes the destination address, the packet header length of HDLR or HDLR$^+$ is also in $O(\log n)$ bits. For the stretch of HDLR or HDLR$^+$, we have the following analytical results.
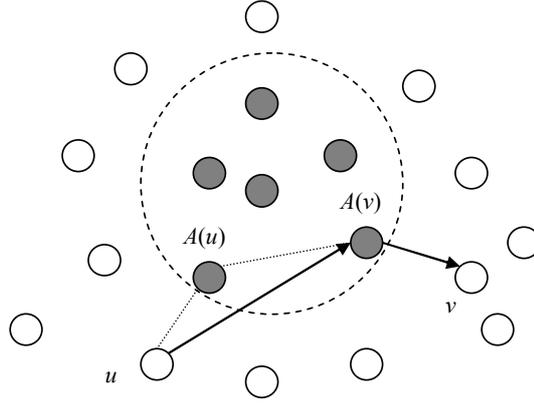

Fig.2. Proving the stretch of HDLR and HDLR$^+$. Filled nodes denote the landmark set. Assuming node u is to send a packet to node v, which is not stored in $Tabel(u)$.

Theorem 1. The worst-case stretch of HDLR and HDLR$^+$ are (3, 0) and $\min\{(3,0),(2,DA)\}$ respectively, where $D_A = \max_{u,v \in A} d(u,v)$ is the maximum distance between any two landmarks.

Proof. Let's first consider the HDLR scheme. Suppose a node u is to send a packet to node v. Let $d_R(u,v)$ denote the length of the route taken by HDLR between u and v. If v is stored in u's routing table, u can route along the shortest path to v, in this case the stretch is (1,0). Otherwise, u will forward the packet to $A(v)$, as is shown in Fig. 2. When the packet reaches $A(v)$, $A(v)$ will forward the packet along the shortest path to v. Therefore we have

$$d_R(u,v) \leq d(u, A(v)) + d(A(v), v) \tag{1}$$

According to the triangle inequality, we know

$$d(u, A(v)) \leq d(u,v) + d(A(v), v)) \tag{2}$$

Combine Equation (1) and (2), we can get

$$d_R(u,v) \leq d(u,v) + 2d(A(v), v) \tag{3}$$

Since v is not contained in u's routing table, i.e. $v \notin A \cup C(u)$, we certainly have $d(u,v) \geq d(A(v), v)$. Thus we get

$$d_R(u,v) \leq 3d(u,v) \tag{4}$$

This produces a stretch (3,0) for HDLR.

For the HDLR$^+$ scheme, It's not difficult to see, Equation (1), (2), (3) and (4) also hold according to its forwarding algorithm and routing table setting. This means the

HDLR$^+$ scheme also has a worst-case stretch (3,0). Besides, we'll prove HDLR$^+$ can achieve another stretch (2,D_A).

As the dash lines in Fig.2 show, according to the triangle inequality, we have

$$d(u, A(v)) \le d(u, A(u)) + d(A(u), A(v)) \tag{5}$$

Combine Equation (1) and Equation (5), we get

$$d_R(u,v) \le d(u, A(u)) + d(A(u), A(v)) + d(A(v), v) \tag{6}$$

Since v is not contained in u's routing table, i.e. $v \notin A \cup B(u) \cup C(u)$, thus both $d(u,v) \ge d(u, A(u))$ and $d(u,v) \ge d(A(v), v)$ hold. Therefore

$$d_R(u,v) \le 2d(u,v) + d(A(u), A(v)) \le 2d(u,v) + D_A \tag{7}$$

This produces a stretch (2, D_A). Combining Equation (4) and Equation (7), we can see HDLR$^+$ can achieve the stretch $\min\{(3,0),(2,D_A)\}$. □

In Theorem 1, one critical problem is what the D_A will be on a typical power-law network. In the following, we show that with high probability D_A is bounded by a constant in random power-law graphs.

Corollary 1. Let G to be an instance of $G(\mathbf{k})$, and the landmark set A to be composed of the $n^x (0 < x < 1)$ nodes with the highest expected degrees, then D_A has probability $1 - n^{-2}$ to be bounded by a constant.

Proof. Assuming the degree distribution of $G(\mathbf{k})$ follows $P(k) \propto k^{-r}$, then the cumulative degree distribution can be represented by $P(z \ge k) = ck^{1-r}$, where c is a constant. Denote the smallest expected node degree in A as τ. According to the cumulative degree distribution, we have $c\tau^{1-r} \approx n^x / n$, so

$$c\tau^{1-\gamma} \approx n^x / n$$
$$\Rightarrow \log c - (\gamma - 1) \log \tau \approx (x-1) \log n$$
$$\Rightarrow \log \tau \approx (1-x) \log n / (\gamma - 1) + \log c / (\gamma - 1)$$
$$\Rightarrow \tau \approx c_1 n^{(1-x)/(\gamma - 1)}$$

where c_1 is a constant depending on γ and c.

A can be treated as $S_{G(\mathbf{k})}^t$ of $G(\mathbf{k})$. According to Lemma 1, the diameter of the subgraph induced by A has probability $1 - n^{-2}$ not exceeding $(1+o(1))\frac{\log n}{(3-\gamma)\log \tau} = (1+o(1))\frac{(\gamma-1)\log n}{(3-\gamma)((1-x)\log n + \log c)}$. Since γ, c and x can all take constant values, this expression will be bounded by a constant as $n \to \infty$. According to the definition of D_A, D_A is no greater than the diameter of the subgraph induced by A, hence D_A is bounded by a constant with probability $1 - n^{-2}$. □

The routing table size of HDLR and HDLR$^+$ can be analyzed separately for landmarks and non landmarks. For each landmark l, HDLR and HDLR$^+$ have the same routing table setting, i.e., $Table_{HDLR}(l) = Table_{HDLR^+}(l) = \{(v, el(v)) | v \in A -$

$\{l\}\}$. Since the landmark set A has a size of $n^x (0 < x < 1)$, the routing table size of each landmark must be bounded by $O(n^x)$. In the following, we analyze the routing table size of non-landmark nodes for HDLR and HDLR$^+$.

We denote the routing table of node u as $Table_{HDLR}(u)$ according to HDLR and $Table_{HDLR^+}(u)$ according to HDLR$^+$, respectively. As defined earlier, for each non landmark node u, $Table_{HDLR}(u) = \{(v, e_u(v)) | v \in A \cup C(u) \}$, and $Table_{HDLR^+}(u) = \{(v, e_u(v)) | v \in A \cup C(u) \cup B(u)\}$. Apparently, since the size of the landmark set A is given, the routing table sizes critically depend on the size of $B(u)$ and $C(u)$. Again, we analyze the size of $B(u)$ and $C(u)$ by means of the power-law random graph theory.

Corollary 2. Let G be an instance of $G(\mathbf{k})$, A be the landmark set consisting of the n^x nodes with the highest expected degrees, then for any $u \in V(G)$, the probability that $|B(u)| = O(n^\alpha)$ is at least $1 - 3n^{-2}$, where $\alpha = \frac{(1-x)(\gamma-2)}{\gamma-1} + \xi$, ξ is sufficiently small.

Proof. According to Corollary 1, the least expected degree $\tau = c_1 n^{(1-x)/(\gamma-1)}$, which can be rewritten as $\tau = c_1 n^{\frac{1-x}{\gamma-1}+\varepsilon'}$, where $\varepsilon' = o(1)$ being sufficiently small. According to Lemma 2, we know that $\alpha = (\frac{1-x}{\gamma-1}+\varepsilon')(\gamma-2) + \frac{(2\gamma-3)\varepsilon}{\gamma-1}$, i.e., $\alpha = \frac{(1-x)(\gamma-2)}{\gamma-1} + \xi$, where $\xi = (\gamma-2)\varepsilon' + \frac{(2\gamma-3)\varepsilon}{\gamma-1}$, Since ε and ε' can both take sufficiently small values, so can ξ. □

For all $u \in V(G)$, Corollary. 2 gives the upper bound of $|B(u)|$. However, it is impossible to analyze the bound of $|C(u)|$ for each u in this way. We however, could provide the upper bound of the average size of $C(u)$.

Lemma 3. For any arbitrary graph G, $\sum_{u \in V(G)} |B(u)| = \sum_{u \in V(G)} |C(u)|$, i.e., $|C(u)|$ and $|B(u)|$ have the same average size.

Proof. According to the definitions of $C(u)$ and $B(u)$, it follows that for any $v \in B(u)$, we have $u \in C(v)$, and vice versa. Since these u, v occur pairwise, so Lemma 3 holds. □

Based on the above analysis, we have the following Theorem.

Theorem 2. Let G be an instance of $G(\mathbf{k})$, and let A be the landmark set containing the n^x nodes with the highest expected degrees, then for any non landmark node $u \in V(G)$, with probability at least $1 - 3n^{-2}$, the average routing table sizes of

$Table_{HDLR}(u)$ and $Table_{HDLR^+}(u)$ are both in $O(n^x + n^\alpha)$, where $\alpha = \frac{(1-x)(\gamma-2)}{\gamma-1} + \xi$ and ξ is sufficiently small.

Proof. Since $|A| = O(n^x)$, $|B(u)| = O(n^\alpha)$, thus both $|C(u)|$ and $|B(u)|$ have an average size $O(n^\alpha)$, so it is easy to see that $Table_{HDLR}(u)$ and $Table_{HDLR^+}(u)$ have an average size of $O(n^x + n^\alpha)$. □

Corollary 3. In HDLR and HDLR[+], For any γ satisfying $2 < \gamma < 3$, the average routing table size of non landmark nodes can be bounded by $O(n^{\frac{1}{3}+\xi})$ with high probability .

Proof. Since $2 < \gamma < 3$, $\alpha = \frac{(1-x)(\gamma-2)}{\gamma-1} + \xi \leq \frac{(1-x)}{2} + \xi$, so $Table_{HDLR}(u)$ and $Table_{HDLR^+}(u)$ have average sizes of $O(n^x + n^{\frac{1-x}{2}+\xi})$ with high probability, which results in $O(n^{\frac{1}{3}+\xi})$ by setting $x = 1/3$. □

So far, we have completed the analytical performance analysis for HDLR and HDLR[+]. The results show HDLR and HDLR[+] can achieve excellent bounds for address length, packet header size, stretch and routing table size.

6 Simulations

In order to provide an insight into the real performance of our proposed scheme on power-law networks, we perform simulations on both the synthesized graphs generated by the random power-law graph (RPLG) model and the real AS-level Internet topology. The AS graph is acquired from CAIDA [29], having a total of 9204 nodes, 28959 edges, and a power law component about 2.1. The random power law graphs are generated with different size and different power law component (e.g., $\gamma = 2.1, 2.3, 2.5$, and $n = 1000, 2000, ..., 10000$).

6.1 Diameter of the subgraph induced by landmarks

Fig.3 reports D_A, the diameter of the subgraph induced by the landmark set A, for random power law graphs with different network size and different power law exponents. We can see that D_A remains relatively stable as the network expands. Also, we observe that when γ increases, D_A increases accordingly. This is because when γ increases, the degrees of landmarks in A will decrease, making the edges connecting landmarks more sparse, thus D_A increases. It is worth noting, for AS graph, D_A is only 2, much smaller than those in the corresponding random

power law graphs. This is due to the rich-club structure in AS graph, i.e., high degree nodes are prone to connect with each other directly, which does not happen in random power law graphs.
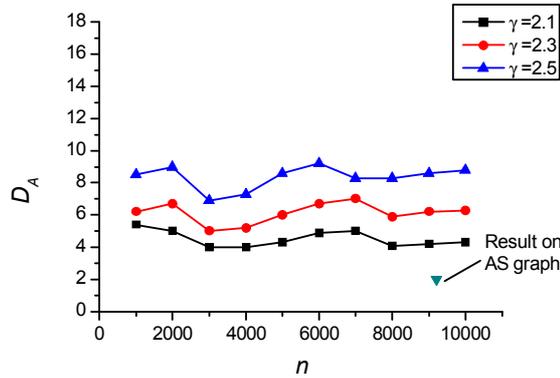

Fig.3. Diameter of landmarks on random power law graphs and AS graph.

6.2 Average Routing Table Size

Fig.4 shows the average routing table size obtained by our schemes and the TZ scheme on random power law graphs with $\gamma = 2.1$ and AS graph. For each random power law graph, the number of landmarks is set to be $|A| = n^{1/2}$, where n denotes the graph size. For AS graph with 9204 nodes, we set $|A| = n^{1/3} = 21$.
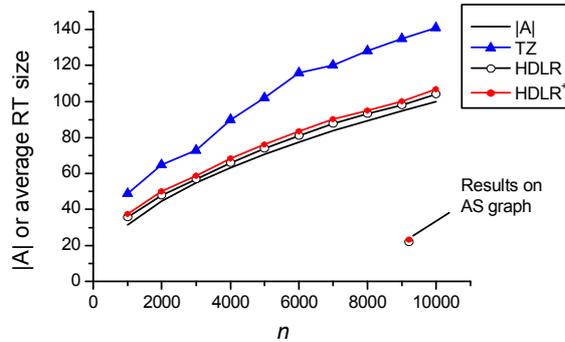

Fig.4. Average routing table size on random power law graphs and AS graph: comparison of our routing schemes and the TZ scheme.

It is evident that with our schemes, the routing table size can be significantly reduced compared to the TZ scheme. With our schemes, the average routing table size is very close to |A|, implying the size of the local cluster and the ball of a node are usually very small. For AS graph, the average routing table size achieved by HDLR

or HDLR[+] is about 22 or 23. It is not surprising that the routing table size produced on AS graph is remarkably smaller than those on random power law graphs. The reason may be that in AS graph, the highest degree nodes have much higher degrees than those in a random power law graph with comparable size. For example, the highest degree of AS graph with 9204 nodes is 2070, while the highest degree of a random power law graph with similar size is only about 60. This implies that in AS graph a majority of nodes are directly connected with a very small set of highest degree landmarks, thus markedly decreasing the size of the routing table.
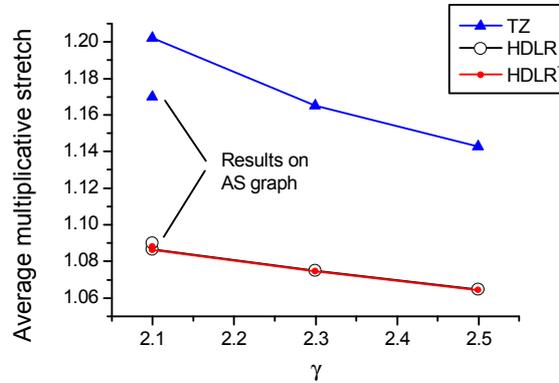
6.3 Average Stretch

Fig.5. Average stretch on random power law graphs and AS graph: comparison of our routing schemes and the TZ scheme.

Fig.5 shows the average multiplicative stretch achieved by our schemes and the TZ scheme on AS graph and random power law graphs with different exponents. Again, for random power law graphs we set $|A| = n^{1/2}$, and for AS graph we set $|A| = n^{1/3} = 21$. It is observed that our schemes outperform the TZ scheme with non marginal improvement For random power law graphs, HLDR or HDLR[+] achieves a 1.06~1.09 average stretch, whereas the TZ scheme can only achieve a much higher average stretch in the range of 1.14~1.20. For AS graph, our schemes have an average multiplicative stretch about 1.09 while the TZ scheme with an average multiplicative stretch about 1.17. The average stretch of HDLR[+] is very close to that of HDLR, with only negligible improvement. This is due to the fact that the routing table size of HDLR[+] is only slightly larger than HDLR. We also notice that, when γ increases, the average stretch of all routing schemes are decreasing. The reason may be, when γ increases, the average routing table size increases, thus a node knows more routing information, resulting more shortened stretched paths.

7   Conclusions

In this paper we present two compact routing schemes, HDLR and HDLR$^+$, aiming to improve the space-stretch tradeoff for routing in power law graphs. Based on the fact that in a power law network graph a few nodes usually have very high degrees, and play the important role of hub in communication and networking, we select high-degree nodes as landmarks.

Analytic study on random power law graphs shows: 1) HDLR can achieve a multiplicative stretch of 3 and an average routing table size of $\tilde{O}(n^{1/3})$ with high probability. 2) HDLR$^+$ can achieve a stretch of $\min\{(3,0),(2,D_A)\}$ and an average routing table size of $\tilde{O}(n^{1/3})$ with high probability, where D_A is the diameter of the sub-graph induced by the landmark set and is almost surely bounded by some constant. Both HDLR and HDLR$^+$ improve the space-stretch tradeoff obtained by the TZ scheme, whose bounds for stretch and routing table size are (3,0) and $\tilde{O}(n^{1/2})$. Simulations conducted on networks generated by synthetic random power law graph model with different parameter settings and real-world AS graph validate our results.

References

[1] Thorup M., Zwick U.: Compact routing schemes. In: 13th ACM Symposium on Parallel Algorithms and Architecture, pp.1-10. ACM Press, New York (2001)

[2] Cowen L.: Compact routing with minimum stretch. Journal of Algorithms, 38(1), pp.170-183 (2001)

[3] Gavoille C., Gengler M.: Space-Efficiency for routing schemes of stretch factor three. Journal of Parallel and Distributed Computing, 61(5), pp.679-687 (2001)

[4] Fraigniaud P., Gavoille C.: Routing in trees. In: 28th International Colloquium on Automata, Languages and Programming, pp. 757-772. Springer, Heidelberg (2001)

[5] Laing K. A.: Name-independent compact routing in trees. Information Processing Letters, 103(2), pp.57-60 (2007)

[6] Lu H.: Improved compact routing tables for planar networks via orderly spanning trees. In: 8th International Computing and Combinatorics Conference, pp. 57-66. Springer, Heidelberg (2002)

[7] Abraham I. and MalkhiD.: Name independent routing for growth bounded networks. In: 17th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA05), pp.49-55. ACM Press, New York (2005)

[8] KonjevodG., Richa A., Xia D.: Optimal scale-free compact routing schemes in networks of low doubling dimension. In: 18th ACM-SIAM Symposium on Discrete Algorithms, pages 939–948. ACM Press, New York (2007)

[9] Abraham I., Gavoille C., Goldberg A.V., and Malkhi D.: Routing in networks with low doubling dimension. In: 26th International Conference on Distributed Computing Systems. IEEE Computer Society, Washington, DC (2006)

[10] Faloutsos M., Faloutsos P., and Faloutsos C.: On power-law relationships of the Internet topology. Computer Communications Review, 29, pp.251-262 (1999)

[11] Ripeanu M., Iarmnitchi A., and Foster I.: Mapping the Gnutella network. IEEE Internet Computing, pp.50-57 (2002)

[12] Albert R., Jeong H., and Barabasi A. L.: Dimeter of the world wide web. Nature, 401, pp. 130-131 (1999)

[13] Krioukov D., Fall K., Yang X.: Compact routing on Internet-like graphs. In: IEEE INFOCOM 2004. IEEE Computer Society, Washington, DC (2004)

[14] Brady A., Cowen L.: Compact routing on power law graphs with additive stretch. In: 8th Workshop on Algorithm Engineering and Experiments, 119–128. SIAM (2006).

[15] Carmi S., Cohen R., Dolev D.: Searching complex networks efficiently with minimal information. Europhysics Letters, 74, 1102–1108 (2006)

[16] Enachescu M., Wang M., Goel A.: Reducing maximum stretch in compact routing. In: IEEE INFOCOM 2008, pp.977-985. IEEE Computer Society, Washington, DC (2008)

[17] Norros I.: Powernet: compact routing on Internet-like random networks. In: 5th Euro-NGI conference on Next Generation Internet Networks. IEEE Press , New York (2009)

[18] Meyer D., Zhang L., and Fall K.: Report from the IAB workshop on routing and addressing. RFC4984 (2007)

[19] Huston G.: Analyzing the Internet's BGP routing table. The Internet Protocol Journal, vol. 4, no.1 (2001)

[20] Krioukov D. and Claffy K.: On compact routing for the Internet. ACM SIGCOMM Computer Communication Review, 37(3), pp.43-52 (2007)

[21] Chung F., Lu L.: The average distances in random graphs with given expected degrees. Internet Mathematics, vol. 1, pp.15879-15882 (2002)

[22] Lu L.: Probabilistic methods in massive graphs and Internet computing. Ph D thesis, University of California San Diego (2002)

[23] Eilam T., Gavoille C., and Peleg D.: Compact routing schemes with low stretch factor. In: 17th Annual ACM Symposium on Principles of Distributed Computing, pp. 11-20. ACM Press, New York (1998)

[24] Thorup M. and Zwick U.: Approximate distance oracles. In: ACM Symposium on Theory of Computing, pp. 183-192. ACM Press, New York (2001)

[25] Barabasi A., Albert R.: Emergence of scaling in random network. Science 286, 509 (1999)

[26] Zhang G.Q., Zhang G.Q., Yang Q.F., Cheng S.Q., and Zhou T.. Evolution of the Internet and its cores. New Journal of Physics, 10, 123027 (2008)

[27] Zhou S., and Mondragon R.J. : The rich-club phenomenon in the Internet topology. IEEE Communications Letters, 3, pp.180-182 (2004)

[28] Chen W., Sommer C., Teng S. H., and Wang Y.: Compacting routing in power-law graphs. In: 23th International Symposium on Distributed Computing. Springer, Heidelberg (2009)

[29] CAIDA, http://www.caida.org