

Compact Routing on Euclidian Metrics

Ittai Abraham
School of Computer Science and Engineering
Hebrew University of Jerusalem
Jerusalem, Israel
ittai@cs.huji.ac.il

Dahlia Malkhi
School of Computer Science and Engineering
Hebrew University of Jerusalem
Jerusalem, Israel
dalia@cs.huji.ac.il

ABSTRACT

We consider the problem of designing a compact communication network that supports efficient routing in an Euclidean plane. Our network design and routing scheme achieves $1 + \epsilon$ stretch, logarithmic diameter, and constant out degree. This improves upon the best known result so far that requires a logarithmic out-degree. Furthermore, our scheme is asymptotically optimal in Euclidean metrics whose diameter is polynomial.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Distributed networks*; G.2.2 [Discrete Mathematics]: Graph Theory—*Network problems, Graph labeling*.

General Terms

Algorithms, Theory.

Keywords

Compact Routing, Network Design.

1. INTRODUCTION

We study the problem of designing a communication network and a compact routing scheme for two dimensional Euclidean metrics¹. Given is a set V of n nodes situated on a two-dimensional plane. Each node $v \in V$ is defined by its coordinates $\langle v_x, v_y \rangle$. For any two nodes $u, v \in V$, let $\|uv\|$ denote the standard L_2 -norm distance, let D be the normalized diameter $\frac{\max_{u \neq v} \|uv\|}{\min_{u \neq v} \|uv\|}$. The problem of *compact routing on Euclidean metrics* is a combined problem of designing a network and a routing scheme on top of the network.

¹We present the two dimensional case for clarity, our construction can be extended to higher dimensional Euclidean metrics

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

PODC'04, July 25–28, 2004, St. Johns, Newfoundland, Canada.
Copyright 2004 ACM 1-58113-802-4/04/0007 ...\$5.00.

- *Network design*. Every node u must choose a set of out going neighbors. This induces a directed graph H on the set V .
- *Routing scheme*. A routing scheme, RS , on the graph H . The scheme allows any source node s that knows a target $t \in V$ and its coordinates $\langle t_x, t_y \rangle$, to route from s to t .

1.1 Complexity measures

The complexity of a solution is based on the following measures:

1. $Degree(H)$ - The maximal out going degree of H .
2. $Memory(RS)$ - The maximal number of bits used for storing routing information in a single node. A routing scheme is said to be *compact* if $Memory(RS) = o(n)$.
3. $Header(RS)$ - The maximal size of the header needed by RS to route from any source s to any destination t .
4. $Stretch(RS)$ - The maximal ratio, over all pairs, of the length of the routing path produced by RS by routing from any source s to any target t and the distance $\|st\|$.
5. $Diameter(RS)$ - The maximal number of hops, over all pairs, produced by RS by routing from any source s to any destination t .

1.2 Our contribution

In this paper, we present a network H and a routing scheme RS that for any $\epsilon > 0$ has the complexity measures below. The big O notations below indicate constant factors that depend on either ϵ or on the desired probabilistic guarantee. We emphasize that our construction does not attempt to optimize these constants, and prioritizes simplicity over tight constants.

1. $Degree(H) = O(1)$.
2. $Memory(RS) = O(\log D)$.
3. $Header(RS) = O(\log D)$.
4. $Stretch(RS) = 1 + \epsilon$.
5. $Diameter(RS) = O(\log D)$ with high probability.

1.3 Related Work

The problem of compact routing using geometric coordinates has been considered in a number of previous works. By simply linking each node to its immediate neighbor in every angle, say θ , one obtains a constant stretch θ -graph spanner as in [12], whose degree is constant but the number of hops may be $\Omega(n)$.

Combining constant stretch and low degree with low diameter requires a more involved construction. The paper of Hassin and Peleg [10] was the first to formally define the problem and give a solution based on building a routing hierarchy. Each node links in every angle θ to $\log D$ different nodes at geometrically increasing distances. Their algorithm achieves the following complexity measures. The stretch is $1 + \epsilon$, memory is $O(\log D)$, and diameter is $O(\log D)$. However the out degree of [10] is $O(\log D)$. Our work builds upon the ideas of Hassin and Peleg’s hierarchical construction and improves their solution by bringing down the node degree (any memory) to a constant. It is interesting to note that, when restricted to nodes on a bidirected ring, Hassin and Peleg’s network has the same topology as Chord [21]. In a similar manner, our network, when restricted to nodes on a bidirected ring, has similarities to Viceroy [17].

Hassin and Peleg present an alternative construction based on hierarchical tree covers [3] that has $O(\log n)$ diameter instead of $O(\log D)$ and achieves a stretch factor strictly larger than $1 + \sqrt{2}$. While the stretch cannot be arbitrarily close to 1, their construction has better asymptotic diameter when $D = \Omega(2^n)$, but in such a case simply storing the coordinates of a node requires $\Omega(n)$ bits which makes the whole scheme non compact.

Geometric routing was also studied in the context of mobile ad hoc networks (MANETs) that are enhanced with self-positioning devices such as GPSs. The model here is somewhat different than ours. It assumes that each node has a certain transmission range, and is linked directly to all nodes within this range (Unit Disk Graph). The first routing algorithm to guarantee delivery is *face routing*, due to Kranakis et al. [14]. However, face routing has no bound on the ratio between the cost of route and the minimal cost path. Both Bose et al. [4] (CGF) and Karp and Kung [11] (GPSR) propose an algorithm that combines greedy routing with face routing. In the MANET model, these algorithms guarantee delivery and for average case networks have expected cost $O(d)$ between a source s and a destination t , where d is the cost of minimal-cost path between s and t on the unit disk graph. The first algorithm that gives worst case guarantees is by Kuhn et al. [15]. They present a scheme in which, if the minimal cost path has cost d , then delivery with cost $O(d^2)$ is guaranteed, which is asymptotically optimal. In a follow up paper [16], they combine their bounded face routing with greedy routing to achieve a scheme that is both worst case asymptotically optimal and average case efficient. Due to the MANET model, all of the above algorithms have worst case diameter $\Omega(n)$.

A closely related problem is that of building sparse, low diameter geometric spanners, which are directed spanning graphs over a set of points on the plane that contain low stretch paths. Arya, Mount and Smid [1] construct a $1 + \epsilon$ stretch, $O(\log n)$ diameter, $O(n)$ edge spanner for any set of nodes in an Euclidean metric. Their construction is built with high probability and is asymptotically optimal. However, no compact decentralized routing algorithm is known

to route on the low stretch paths while requiring a non-trivial number of bits per node.

The general problem of designing labels and compact routing tables for low stretch routing on arbitrary weighted graphs is considered in [5, 24]. In these schemes the designer is allowed to give nodes labels with a polylogarithmic number of bits. The best labeled routing schemes for general weighted graphs achieve stretch $2k - 1$ for $k \geq 2$ with $O(n^{1/k} \log^2 n)$ bit routing tables and this is tight up to polylogarithmic factors [23, 24]. For comprehensive surveys on compact routing and compact network data structures, see [7, 8].

Recently, Talwar [22] gave a stretch $1 + \epsilon$ compact routing scheme with $O((\log D/k)^k \log^2 D)$ bit routing tables for graphs that induce a metric with *doubling dimension* k . A metric has a doubling dimension of k if every ball of radius $2r$ can be covered by at most 2^k balls of radius r , hence any finite dimension Euclidean Metric has a constant doubling dimension.

1.4 Applicability

Using geometric coordinates in general internets is made relevant not only by the ubiquity of GPS devices, but also by several recent techniques that embed internet nodes in a coordinate space. One of the pioneering mechanisms to predict network latency is based on the work of Ng and Zhang [19]. They embed the Internet latencies into a virtual geometric space (e.g., 3-D Euclidean) and characterize the position of any node with coordinates. The computed distances are used to predict the actual network distances. Following [19] other schemes were developed to improve the embedding of internet hosts into virtual geometric spaces, e.g., [9], [25], [6], and [20].

Our work embraces this promising direction, which opens new exciting opportunities. By utilizing geometric coordinates, not only can the distances between nodes be predicted accurately and efficiently, but in addition, the properties of the Euclidean space may be used. Most importantly, Euclidean spaces have a ‘sense of direction’ which allows to perform distance-preserving routing while maintaining a small number of links. We believe this approach may improve the design of overlay networks for various real-life systems.

Our work is also relevant to an on-going effort in designing geometric routing networks for peer-to-peer (p2p) applications, based on routing in “small worlds” [13].² The goal in this domain is for a dynamic set of nodes to jointly implement a shared data structure, such as a hash table. In order to realize a shared structure distributively, operations on data are routed among the nodes in order to dispatch where the data resides. The p2p works consider nodes dispersed uniformly on a Euclidean space (either real or virtual) of one or two dimensions and route in a distance-preserving manner. Chord [21] uses $O(\log n)$ links per node, and achieves an expected diameter $O(\log n)$. Symphony [18] uses a Kleinberg-like link distribution, achieving an expected diameter of $O(\log^2 n/k)$ with k links per node. The same complexity was achieved in [2] with a different Kleinberg-style randomized p2p network. Viceroy [17] achieves $O(\log n)$ diameter with 5 links.

²It is only fair to note that the importance of Kleinberg’s work lies completely outside the area of network design, and its focus is on modelling real-life small-world networks. Our work does not attempt to add to the understanding of social networks.

Using our network design with nodes that are dispersed uniformly on a uni-dimensional space matches the best complexity measures so far, i.e., of Viceroy. In this respect, our work extends all known p2p overlay network constructions into an arbitrary density space, while preserving locality, constant node degree, and logarithmic diameter.

As a special case, when considering the uniform “small world” network model analyzed in [13], our results imply that a constant number of carefully chosen long range links can significantly reduce the delivery time to $O(\log n)$ hops, while the scheme in [13] requires an expected $O(\log^2 n)$ delivery time.

2. PRELIMINARIES

Consider a set V of n nodes situated on a two-dimensional plane. Each node $v \in V$ is defined by its coordinates $\langle v_x, v_y \rangle$. For any two nodes $u, v \in V$, let $\|uv\|$ denote the standard L_2 -norm distance. For simplicity of presentation, assume that the distances are normalized so that the minimal distance between any pair of nodes is 1. Let D denote that maximal distance between pairs of nodes.

We consider two density models. The first one, a uniform grid, is presented for intuition. The second is a general plane, with arbitrary node densities.

We fix a parameter k , and set an angle $\alpha = 2\pi/k$. The construction will use three *sector* angles $\theta_1, \theta_2, \theta_3$, where $\alpha = \theta_1 \leq \theta_2 \leq \theta_3$, and three *overshoot* parameters o_1, o_2, o_3 , where $1 = o_1 \leq o_2 \leq o_3$. For the uniform grid, we set all sector angles to equal α and all overshoot parameters to equal 1, i.e. $\theta_1 = \theta_2 = \theta_3 = \alpha$ and $o_1 = o_2 = o_3 = 1$. We will later redefine these constants for the general case.

For the following definitions denote the following index sets: $J = \{1, 2, 3\}$ and $I = \{0, 1, \dots, k-1\}$.

Definition 2.1 (Ray of angle β). For any angle β , let $R(\beta)$ denote the ray from the origin with angle β from the x axis.

Definition 2.2 (Sector i of angle θ_j). For any index $j \in J$ and sector $i \in I$, let $S_j(i)$ denote the sector with angle θ_j whose boundary rays are $R(i\alpha - \frac{1}{2}\theta_j)$ and $R(i\alpha + \frac{1}{2}\theta_j)$.

Definition 2.3 (u 's shifted sector i of angle θ_j). For any node $u \in V$, index $j \in J$, and sector $i \in I$, define the shifted sector $S_j(u, i)$ as the sector $S_j(i)$ shifted by \vec{u} so that u is the new origin.

Note that $\alpha = \theta_1$ so $\{S_1(u, i) \mid i \in I\}$ forms a partition of the plane. For any two nodes $u, t \in V$ let $i(u, t)$ denote the index such that $t \in S_1(u, i(u, t))$ (with ties broken in any consistent manner).

Definition 2.4 (u 's ball of radius r). For a node $u \in V$ and distance $r \in \mathbb{R}^+$, let $B(u, r) = \{v \in V \text{ s.t. } \|uv\| \leq r\}$ denote the ball around u of radius r .

In our construction every node $u \in V$ has a *range* denoted $r(u) \in \mathbb{R}^+$ (we will later define exactly how this range is chosen). The range of a node induces several regions.

Definition 2.5 (u 's ℓ -th region in sector i). For a node $u \in V$ with range $r(u) = r$, sector $i \in I$, and $\ell \in \{0, 1, 2, 3, 4, 5\}$ we define the following regions:

$$r_1(u, i, \ell) = S_1(u, i) \cap B(u, r(\ell+1)/6) \setminus B(u, r\ell/6)$$

In addition for $\ell = 6$ we define:

$$r_1(u, i, 6) = S_1(u, i) \setminus B(u, r)$$

Definition 2.6 (u 's ℓ -th padded region in sector i). In the uniform density model, define $r_3(u, i, \ell) = r_1(u, i, \ell)$ as node u 's ℓ -th padded region in sector i .

For the general case, we will later on redefine the padded region $r_3(u, i, \ell)$ as a function of $S_3(u, i)$, $B(u, \cdot)$, $r(u)$ and o_3 . Denote $\bar{L} = \{0, 1, 2, 3, 4, 5\}$ and $L = \bar{L} \cup \{6\}$. Note that $\{r_1(u, i, \ell) \mid i \in I, \ell \in \bar{L}\}$ forms a partition of $B(u, r(u))$. For any two nodes u, t such that $t \in B(u, r(u))$ let $\ell(u, t)$ denote the index such that $t \in r_1(u, i(u, t), \ell(u, t))$ (with ties broken in any consistent manner). See Figure 1 for a pictorial example of the definitions above.




Figure 1: Node u and its division into regions.

We now define areas that depend on the current node u and the destination node t :

Definition 2.7 (j -th pad area of u and t). For any two nodes $u, t \in V$, and index $j \in J$ denote $s_j(u, t)$ as $s_j(u, t) = S_j(u, i(u, t)) \cap B(u, o_j \|ut\|)$ as the j -th pad area of u and t .

See Figure 2 for an example of $s_1(u, t)$ and $s_2(u, t)$. The pad areas have a central role in keeping the stretch low, our routing scheme always routes from u , towards destination t , to a node inside $s_3(u, t)$.

Definition 2.8 (halving area of u and t). For any two nodes $u, t \in V$ such that $t \in B(u, r(u))$ let $i = i(u, t)$ and $\ell = \ell(u, t)$ then denote $h(u, t)$ as

$$h(u, t) = s_3(u, t) \cap \bigcup_{\kappa \in \{-1, 0, 1\}} r_3(u, i, \ell + \kappa)$$

as the halving area of u and t .




Figure 2: The inner area is $s_1(u, t)$ and the outer area is $s_2(u, t)$.

The halving areas are used in our routing scheme to route to a node that decreases the distance to the target by half.

Finally, designate a parameter $\rho > 3$ that determines the high probability bounds that are desired.

3. THE NETWORK

Every node $u \in V$ chooses a *range*, denoted $r(u)$, to be equal to $D2^{-\ell}$, where ℓ is chosen uniformly and independently from $\{0, 1, \dots, \lceil \log_2 D \rceil\}$. We begin with an overview of the edge selection and routing scheme.

3.1 Network and Routing overview

In this section we present a high level overview of the network design and routing scheme. Routing towards a target t is done in three main stages: range adjustment, distance halving, and local routing.

1. Range adjustment. The goal of this stage is to reach a node u such that $t \in B(u, r(u))$. This is done using the *up* links. These links connect a node with range r to the closest node with range $2r$ in every sector $S_3(u, i)$.

Since node-ranges form a geometric series, this stage ends in at most $\log D$ hops. It is also possible that the network is too sparse and no appropriate *up* link is found in the vicinity of the target. In such a case routing proceeds directly to the third stage.

2. Distance halving. This stage begins with a node u such that $t \in B(u, r(u))$. Using the *down* links, a node u' with half the range is found and the invariant $t \in B(u', r(u'))$ is maintained. The stage ends when a node does not have an appropriate distance halving link in the desired vicinity of the target.

3. Local routing. The final stage begins in an area that contains the target and only $\rho \log^2 D$ other nodes. There are two sub stages: the *log* sub stage and the *closest* sub stage.

The *log* links connect to the $\log D$ closest node in every sector $S_2(u, i)$. Using the *log* links, after at most $\rho \log D$ hops an even sparser area with less than $\log D$

nodes is reached. Taking care that routing does not leave the sparse area incurs some technical details.

The second sub stage begins with an area that contains the target and at most $\log D$ other nodes. Using the *closest* links, the target is found after at most $\log D$ hops. Again care must be taken not to leave the sparse area.

In the following subsections, we will formally define for each stage, the network design (the links that build the graph H) and the routing scheme (the routing algorithm used to reach the target node).

3.2 Range adjustment

Network design. Every node u with range $r = r(u)$ chooses the following *up* links for range adjustment:

- For all $i \in I$, choose $up(i)$ as the closest node in $S_3(u, i)$ with range $2r$.

Note that some of the sectors may not contain a node with the desired range, in such a case the appropriate link is set to \perp .

Routing scheme. At a node u , range adjustment routing is done using the following rules:

- If $t \in B(u, r(u))$ then the range is good, and routing proceeds to the distance halving stage.
- Otherwise, if $up(i(u, t)) \in s_3(u, t)$ then route using $up(i(u, t))$.
- Otherwise, u and the target t are in a sparse area $s_3(u, t)$, and routing proceeds to the local routing stage.

3.3 Distance halving

Network design. Given a node u with range $u(r)$ there are $7k$ links of type *down* and $6k$ of type *local* chosen as follows,

- For all $i \in I$, and $\ell \in L$ choose $down(i, \ell)$ as the closest node in the padded region $r_3(u, i, \ell)$ with range $r(u)/2$. If there exists such a node then this link will be called *non-sparse*.
- For all $i \in I$, and $\ell \in \bar{L}$ choose $local(i, \ell)$ as the closest node in the region $r_1(u, i, \ell)$.

Note that some of the regions and padded regions may not contain a node with the desired range or any node at all, in such a case the appropriate link is set to \perp .

Routing scheme. At a node u , distance halving is done using the following rules:

- If exists $i \in I$ and $\ell \in L$ such that $down(i, \ell) \in h(u, t)$ then use the $down(i, \ell)$ link and continue with the distance halving stage.
- Otherwise, use $local(i, \ell)$ such that $i = i(u, t)$ and $\ell = \ell(u, t)$ and begin the local routing stage.

3.4 Local routing

Network design. Every node u chooses the following *log* links and *closest* links:

- For all $i \in I$, choose $log(i)$ as the $(\log D)$ 'th closest node in $S_2(u, i)$.

- For all $i \in I$, choose $closest(i)$ as the closest node in $S_1(u, i)$.

Routing scheme. Local routing is performed in two sub stages. The first sub stage uses the *log* links until a very sparse area is reached. Finally, the second sub stage uses the *closest* links in order to reach the target.

- Denote $i = i(u, t)$.
- If $log(i) \in s_2(u, t)$ then use $log(i)$.
- Otherwise use the $closest(i)$ link (which is clearly inside $s_1(u, t)$).

4. ANALYSIS FOR A UNIFORM GRID

In this section we present a simplified analysis for the case in which the n nodes are located on a $\sqrt{n} \times \sqrt{n}$ lattice with unit distances. Note that this simple network model is similar to the one studied in [13].

Recall that for the case of a uniform grid, we set $\alpha = \theta_1 = \theta_2 = \theta_3$, and $1 = o_1 = o_2 = o_3$ and hence $s_1(\cdot) = s_2(\cdot) = s_3(\cdot)$, and $r_1(\cdot) = r_3(\cdot)$. The following complexity measures are a direct result of our scheme.

Lemma 4.1. *The network and routing scheme have the following properties:*

1. $Degree(H) = O(1)$.
2. $Memory(RS) = O(\log D)$.
3. $Header(RS) = O(\log D)$.

PROOF. For the degree, there are k links of type *up*, $7k$ of type *down*, $6k$ of type *local*, k of type *log*, and k of type *closest*. As for local memory and header size, they designate grid coordinates, a pair of values between 1 and D . Hence, each memory entry of an out going link and every packet header is of size $O(\log D)$. \square

Note that, on a grid, $D = \sqrt{2n}$ and thus $O(\log D) = O(\log n)$.

For the uniform grid model note that every hop in every stage from u to v towards the destination t maintains the invariant $v \in s_1(u, t)$. Therefore, the stretch analysis can be derived by Lemma 3.6 of Hassin and Peleg [10].

Lemma 4.2. [10] *Let P be a path leading to t . If each hop $(u \rightarrow v) \in P$ has $\|uv\| \leq \|ut\|$ and $\angle tuv \leq \alpha$ then the stretch of the path is at most $\frac{1}{1-2\sin\alpha/2}$.*

Corollary 4.3. *The routing scheme has*

$$Stretch(RS) \leq \frac{1}{1-2\sin\alpha/2},$$

and for an appropriate choice of $\alpha = \alpha(\varepsilon)$, this is bounded by $1 + \varepsilon$.

In order to analyze the diameter we begin by showing that every hop in the distance halving stage really cuts the distance by half.

Lemma 4.4. *Given a target t , if u routes using a non-sparse link in the distance halving stage to node v then $t \in B(v, r(v))$.*

PROOF. Since u performs a distance halving step then $t \in B(u, r(u))$ hence $\|ut\| \leq r(u)$. Denote $i = i(u, t)$ and $\ell = \ell(u, t)$. In the uniform model due to the definition of $s_1(\cdot) = s_3(\cdot)$ then $v \in h(u, t)$ implies $v \in r_1(u, i, \ell) \cup r_1(u, i, \ell - 1)$.

The farthest distance between two nodes v, t in any $r_1(u, i, \ell) \cup r_1(u, i, \ell - 1)$ is maximized when $\ell = 5$, node v is at distance $(2/3)r(u)$ from u at one corner of $r_1(u, i, 4)$ and node t at distance $r(u)$ from u at the opposite corner of $r_1(u, i, 5)$. In such a case the remaining distance $\|v, t\|$, normalized to $r(u) = 1$, is

$$\sqrt{1 + (2/3)^2 - 2(2/3)\cos\theta_3}$$

which is smaller than $1/2$ when $k > 14$. \square

Lemma 4.5. *With high probability the routing scheme has*

$$Diameter(RS) = O(\log n).$$

PROOF. The range adjustment stage takes at most $O(\log n)$ hops until a node u is found such that $t \in B(u, r(u))$, this is due to the fact that each time an *up* link is used the range is doubled.

Then, from Lemma 4.4 each time a *down* link is used the distance halving stage reduces the level by one and maintains the invariant $t \in B(u, r(u))$. Thus distance halving takes at most $O(\log n)$ steps.

Consider the case that a *local* link is used and denote $i = i(u, t)$ and $\ell = \ell(u, t)$. The appropriate halving area $h(u, t)$ does not have a node of the requested level and hence the *local* link in $r_1(u, i, \ell) \cap s_1(u, t)$ is used. We claim that with high probability $r_1(u, i, \ell) \cap s_1(u, t)$ contains at most $\rho \log^2 n$ nodes. The probability that any region $r_1(u, i, \ell) \cap s_1(u, t)$ with more than $\rho \log^2 n$ nodes does not contain an appropriate link is at most

$$\left(1 - \frac{1}{\log n}\right)^{\rho \log^2 n} \leq n^{-\rho}.$$

The same analysis shows that if routing proceeds directly from range adjustment to local routing then the probability that the area $s_3(u, t)$ contains more than $\rho \log^2 n$ is at most $n^{-\rho}$.

For each of the n^2 pairs there are is one $s_3(u, t)$ sector to consider and at most one $r_1(u, i(u, t), \ell(u, t)) \cap s_1(u, t)$ padded region to consider. The lemma follows from the union bound by setting $\rho > 3$.

Once a *local* link is used the third stage begins. The first substage uses the *log* links towards the target. Since the distance is w.h.p. $\rho \log^2 n$ this stage will end after at most $O(\log n)$ hops when the distance is closer than $\log n$. Now the second substage starts using the *closest* links. This stage ends after at most $O(\log n)$ hops.

\square

5. ANALYSIS FOR GENERAL EUCLIDIAN METRICS

The degree, memory consumption and header size complexity in general metrics remain the same as in Lemma 4.1.

The main difficulty in the general case is bounding the diameter of the routing scheme. We overcome this difficulty by redefining the angles θ_2, θ_3 , overshoot parameters o_2, o_3 , region $r_3(u, i, \ell)$, and restricting k to be $k > 200$.

For simplicity, we present a design in which the *down* links cut the distance by a factor of $\frac{1}{2}$. Simple extensions to our scheme can lower k by a constant factor at the cost of slightly increasing the constants in the $O(\log D)$ diameter bound. However, we emphasize again that we do not attempt to tighten the constants in this exposition.

The parameters $\theta_1 = \alpha$ and $o_1 = 1$ remain unchanged. Consider the last sub-stage of the local routing stage which uses *closest* links. This sub-stage begins from a node u such that u 's appropriate *log* link is not inside $s_2(u, t)$, thus $s_2(u, t)$ has less than $\log D$ nodes. At this final sub-stage, a node searches for the next hop of type *closest* in the sector of angle θ_1 containing the target. Our aim is for this sub-stage to take $O(\log D)$ hops. In a non uniform Euclidean space, this requires a delicate construction. Special care needs to be taken in order to ensure that routing does not leave the sparse area $s_2(u, t)$. Otherwise, routing with *closest* links could lead to a node outside $s_2(u, t)$ and in general networks there is no bound on the densities in its vicinity and hence no bound on the number of hops. Our solution is to redefine $\theta_2 = 7\alpha/2$ and $o_2 = 1 + 2\sin(\alpha/2)$ which in turn increases the area $s_2(u, t)$. For this choice of parameters we prove that routing inside this area towards the target using *closest* links will always remain inside the initial sparse area $s_2(u, t)$ that contains only $O(\log D)$ nodes. Since the distance to the target always decreases this implies that the last sub-stage takes at most $O(\log D)$ hops.

Lemma 5.1. *Fix an initial node u , target t , and area $s_2(u, v)$. For any intermediate node, if $v \in s_2(u, t)$ then $s_1(v, t) \subset s_2(u, t)$.*

PROOF. The parameters $\theta_2 = 7\alpha/2$ and $o_2 = 1 + 2\sin(\alpha/2)$ were chosen specifically so that for $k > 200$, and any t we have $B(t, (o_2 - 1)\|ut\|) \subset s_2(u, t)$. See Figure 3. Indeed, we know that $t \in s_1(u, t)$, and the shortest distance from t to any point on the boundary of $s_2(u, t)$ is at least

$$\|ut\| \sin\left(\frac{\theta_2 - \theta_1}{2}\right) = \|ut\| \sin(5\alpha/4).$$

For $\alpha \leq 2\pi/200$, we have $\sin(5\alpha/4) \geq 2\sin(\alpha/2)$.




Figure 3: Nodes u, t, v, p and regions for Lemma 5.1.

Now consider any node $v \in s_2(u, t)$ and point p on the curved boundary of $s_1(v, t)$. Since $\|vp\| = \|vt\|$ and $\angle tvp \leq \alpha$, then from the cosine law and the identity $\sin(\frac{x}{2}) = \sqrt{\frac{1 - \cos(x)}{2}}$ the distance between point p and node t is at most

$$\sqrt{d^2 + d^2 - d^2 \cos(\alpha)} = d\sqrt{2(1 - \cos(\alpha))} \leq d \cdot 2\sin(\alpha/2)$$

where $d = \|vt\|$. Hence $\|pt\|$ is at most $(o_2 - 1)\|vt\|$. Given $\|vt\| \leq \|ut\|$ it follows that $p \in B(t, (o_2 - 1)\|ut\|)$ and thus $p \in s_2(u, t)$. This proves the lemma since $v \in s_2(u, t)$ and both $s_2(u, t)$ and $s_1(v, t)$ are convex areas. \square

We now proceed in redefining θ_3, o_3 , and $r_3(\cdot)$. Using the same arguments as in Lemma 4.5 the *log* routing sub-stage begins w.h.p. either (from the distance halving stage) in a region $h(u, t)$ that contains at most $\rho \log^2 D$ nodes, or (from the range adjustment stage) in a sector $s_3(u, t)$ that contains at most $\rho \log^2 D$ nodes. We must show two things about the *log* routing sub-stage:

1. Routing with *log* links towards the target always remains inside the original sparse area, either $h(u, t)$ or $s_3(u, t)$.
2. It takes at most $O(\log D)$ hops until an area $s_2(u', t)$ is reached with less than $\log D$ nodes (and from there the *closest* links can be used).

In order to prove (1.) we need to ensure that for every intermediate node v in this sub-stage, routing to the next hop in an angle θ_2 from t , at distance at most $\|vt\|o_2$ from v will reach a node that is still inside the original sparse area. The sparse area can either be $s_3(u, t)$ or $h(u, t)$. We guarantee staying inside the sparse area by redefining $\theta_3 = 12\alpha$ and $o_3 = 1 + 2\sin(2\alpha)$. See Table 1 for a summary of the values.

Table 1: Overshoot and angle values

overshoot	value
o_1	1
o_2	$1 + 2\sin(\alpha/2)$
o_3	$1 + 2\sin(2\alpha)$
angle	value
θ_1	α
θ_2	$7\alpha/2$
θ_3	12α

Finally, we need to redefine $r_3(\cdot)$ as follows.

Definition 5.2 (*u 's ℓ -th padded region in sector i*).

For a node $u \in V$ with range $r(u) = r$, sector $i \in I$, and index $\ell \in \bar{L}$ redefine $r_3(u, i, \ell)$ as

$$r_3(u, i, \ell) = \text{CH}(S_3(u, i) \cap B(u, o_3 r(\ell + 1)/6) \setminus B(u, r\ell/6))$$

Where $\text{CH}(A)$ is the Convex Hull of the set A .

See Figure 4 for an example of a padded region.

With these parameters and definitions we prove the following.

Lemma 5.3. *Fix an initial node u and target t . For any intermediate node, if $v \in h(u, t)$ and $\|vt\| \leq \|ut\|$ then $s_2(v, t) \subset h(u, t)$.*

Similarly, if $v \in s_3(u, t)$ and $\|vt\| \leq \|ut\|$ then $s_2(v, t) \subset s_3(u, t)$.

PROOF SKETCH. The proof is similar to Lemma 5.1. Denote $i = i(u, t)$ and $\ell = \ell(u, t)$. The parameters $\theta_3 = 12\alpha$ and $o_3 = 1 + 2\sin(2\alpha)$ were chosen specifically so that for $k > 200$, and for any $u, t \in V$:

$$B(t, (o_3 - 1)\|ut\|) \subset h(u, t).$$




Figure 4: Example of the region $r_1(u, i, \ell)$ and the padded region $r_3(u, i, \ell)$ for $\ell = 1$.

It is easy to show that the upper and lower boundary lines of $h(u, t)$ are far enough. We now show that the side rays are also far.

The shortest distance from t to any point on the two boundary rays of $h(u, t)$ is at least

$$\|ut\| \sin\left(\frac{\theta_3 - \theta_2}{2}\right) = \|ut\| \sin(17\alpha/4).$$

For $\alpha \leq 2\pi/200$, we have

$$\sin(17\alpha/4) \geq 2\sin(2\alpha).$$

Hence $B(t, (o_3 - 1)\|ut\|) \subset h(u, t)$.

Now consider any node $v \in h(u, t)$ such that $\|vt\| \leq \|ut\|$ and point p on the curved boundary of $s_2(v, t)$. Thus $\angle pvt \leq \theta_2$, and $\|vp\| = o_2\|vt\|$, by the cosine law on the triangle pvt , the point p is at a distance at most $\|vt\|\sqrt{1 + o_2^2 - 2o_2\cos\theta_2}$ from t and $\|vt\| \leq \|ut\|$. For $\alpha \leq 2\pi/200$, we have

$$2\sin(2\alpha) \geq \sqrt{1 + o_2^2 - 2o_2\cos\theta_2}.$$

So $p \in B(t, (o_3 - 1)\|ut\|)$ and thus $p \in h(u, t)$. This proves the lemma since both $h(u, t)$ and $s_2(v, t)$ are convex areas.

The proof that if $v \in s_3(u, t)$ and $\|vt\| \leq \|ut\|$ then $s_2(v, t) \subset s_3(u, t)$ is done in a similar manner. \square

We still need to prove that it takes at most $O(\log D)$ hops of type *log* until an area $s_2(u, t)$ is reached with less than $\log D$ nodes. A *log* step from a node u ‘skips’ $\log D$ nodes within $s_2(u, t)$ that are closer to u than the target t . However, because of the angle of the sector, it is possible that some skipped nodes are closer to t than the *log* link, and so we cannot simply discount the skipped nodes. Showing that within $O(\log D)$ hops the target is reached requires a more subtle argument, which uses the stretch analysis below.

Theorem 5.4. *In the general case, with $o_3 = 1 + 2\sin(2\alpha)$, the routing scheme has*

$$\text{Stretch}(RS) = \frac{o_3}{1 - \sqrt{1 + (o_3)^2 - 2o_3 \cos(13\alpha/2)}},$$

and for an appropriate choice of $\alpha = \alpha(\varepsilon)$, this is bounded by $1 + \varepsilon$.

PROOF. Denote $g(\alpha) = \frac{o_3}{1 - \sqrt{1 + (o_3)^2 - 2o_3 \cos(13\alpha/2)}}$. Let P denote the path taken by the routing algorithm from source node s to destination node t . For any hop $(u \rightarrow v) \in P$ from node u to a node v on the path P , denote $\text{cost}(u, v) = \|uv\|$ the cost of the hop, and denote $\text{gain}(u, v) = \|ut\| - \|vt\|$

the gain towards the target of the hop. Since for all stages $v \in s_3(u, t)$ and $k > 200$ then the gain is always positive.

In order to show low stretch we will show that for any hop $(u \rightarrow v)$ in P

$$\frac{\text{cost}(u, v)}{\text{gain}(u, v)} \leq g(\alpha). \quad (1)$$

This implies

$$\text{cost}(P) = \sum_{(u \rightarrow v) \in P} \text{cost}(u, v) \leq g(\alpha) \|st\|$$

and choosing $\alpha = g^{-1}(1 + \varepsilon)$ gives the desired stretch.

Let u be any node along P and v the next hop towards destination t , then our routing scheme maintains the following properties due to the fact that $v \in s_3(u, t)$:

1. The angle $\angle twv$ is at most $\theta_3/2 + \theta_1/2 = 13\alpha/2$. Let $i = i(u, t)$, in our construction, $t \in S_1(u, i)$ and v is always in $S_3(u, i)$ so $\angle twv$ is maximized when t is on one ray of $S_1(u, i)$ and v is on the opposite ray of $S_3(u, i)$.
2. The distance $\|uv\|$ is at most $o_3\|ut\|$. Note that this distance is even smaller for some routing stages.

Denote $d = \|ut\|$, $c = \|uv\|$, and $e = \|vt\|$. If c, d remain fixed then increasing $\angle twv$ increases $\text{cost}(u, v)/\text{gain}(u, v)$, hence the maximum is obtained when $\angle twv = 13\alpha/2$. This follows directly from the law of the cosines on $\angle twv$.

We now look for a maximum of the function $f(x) = x/(1 - \sqrt{1 + x^2 - 2x \cos(13\alpha/2)})$ in the range $0 \leq x \leq o_3$. There is minimum in the range and no other extreme points, hence the maximum of $f(c)$ is obtained when $c = o_3$. Therefore,

$$\frac{\text{cost}(u, v)}{\text{gain}(u, v)} = \frac{c}{d - e} \leq \frac{o_3}{1 - \sqrt{1 + (o_3)^2 - 2o_3 \cos(13\alpha/2)}}.$$

\square

Bounding the diameter in the *log* sub stage requires also a ‘dual’ to [Theorem 5.4](#).

Lemma 5.5. *If a path from u towards t reaches node v such that:*

1. The angle $\beta = \angle twv$ is at least $\cos^{-1}\left(\frac{1}{1+\varepsilon}\right)$ and at most $\pi/2$.
2. The gain $\text{gain}(u, v) = \|ut\| - \|vt\|$ is positive.

Then $\text{cost}(u, v)/\text{gain}(u, v) > 1 + \varepsilon$.

PROOF. Denote as above $d = \|ut\|$, $c = \|uv\|$, and $e = \|vt\|$ then $\frac{\text{cost}(u, v)}{\text{gain}(u, v)} = \frac{c}{d - e}$ and we want to prove that $\frac{c}{d - e} > 1 + \varepsilon$ which is equivalent to proving $e > d - \frac{c}{1 + \varepsilon}$, using the law of the cosine and gain positivity this is equivalent to proving the following inequality

$$c^2 + d^2 - 2cd \cos \beta > \left(d - \frac{c}{1 + \varepsilon}\right)^2 = d^2 + \frac{c^2}{(1 + \varepsilon)^2} - \frac{2cd}{1 + \varepsilon}$$

which is true if

$$c \left(1 - \frac{1}{(1 + \varepsilon)^2}\right) > 2cd \left(\cos \beta - \frac{1}{1 + \varepsilon}\right)$$

and this is true for any c if $\cos^{-1}\left(\frac{1}{1 + \varepsilon}\right) < \beta \leq \pi/2$ since RHS is negative and LHS is positive. \square

Finally we can conclude the analysis of the *log* stage by showing that routing with the *log* links takes $O(\log D)$ hops.

Lemma 5.6. *The log sub stage takes at most $O(\log D)$ hops.*

PROOF. We will show that for every hop, either the distance is cut by a constant factor, or there are $\log D - 1$ nodes that will never be considered during the route. Suppose node u uses its *log* link and moves to node w .

Since $w \in s_2(u, t)$ then by construction $\angle twu \leq \theta_2$. Now there are two cases to consider. If $\|uw\| \geq \|wt\|$ then by




Figure 5: Nodes in Lemma 5.6

Theorem 5.4. $\|uw\| \leq (1 + \varepsilon)(\|ut\| - \|wt\|)$, hence the distance has been shortened by a factor of at least $\frac{2+\varepsilon}{1+\varepsilon}$.

Otherwise, $\|uw\| < \|wt\|$ and so $\angle utw \leq \theta_2$. Therefore, $\angle uwt \geq \pi - 2\theta_2$. Consider any node $x \in s_2(u, t)$ such that $\|ux\| \leq \|uw\|$ (see Figure 5). Then $\angle uwx \leq \pi/2$, hence $\angle xwt = \angle uwt - \angle uwx \geq \pi/2 - 2\theta_2$.

Thus after visiting w , the routing scheme will never visit any node $x \in s_2(u, t)$ such that $\|ux\| \leq \|uw\|$ because the angle $\angle xwt$ would contradict Lemma 5.5 for $k > 200$. Hence all the $\log D - 1$ nodes whose distance to u is smaller than $\|wt\|$ will never be revisited. Therefore, each such *log* link skips a disjoint set of $\log D$ nodes, and there could be at most $\rho \log^2 D / \log D$ such steps.

□

In a similar manner to Lemma 4.4 we show that for large enough k , the distance is halved at each hop of the distance halving stage.

Lemma 5.7. *Given a target t , if u routes using a non-sparse down link in the distance halving stage to node v then $t \in B(v, r(v))$.*

PROOF SKETCH. Suppose $t \in B(u, r(u))$ and a non-sparse down link $v \in h(u, t)$ is used. Then the maximal distance between v and t is obtained when v is at distance $(2/3)r(v)$ from u in one corner of $r_1(u, i, 4)$, and v is in the opposite corner of $r_3(u, i, 5)$. In such a case $\angle vut \leq \theta_3/2 + \theta_1/2 = 13\alpha/2$ and the distance $\|vt\|$, normalized to $r(u)$, is at most

$$\sqrt{(2/3)^2 + (1 + 2 \sin(2\alpha))^2 - 2(2/3)(1 + 2 \sin(2\alpha)) \cos(13\alpha/2)}$$

which is smaller than $\frac{1}{2}$ for any $k > 200$. □

We conclude with the diameter analysis.

Theorem 5.8. *For the general case, the routing scheme has*

$$\text{Diameter}(RS) = O(\log D) .$$

PROOF. The range adjustment stage takes at most $\log D$ hops until it reaches either:

1. A sparse area $h(u, t)$ containing less than $\rho \log^2 D$ nodes, or
2. A node u such that $t \in B(u, r(u))$.

In case 2, the distance halving stage commences. In this stage, from Lemma 5.7 each non-sparse link that is used cuts the range in half. Hence the invariant $t \in B(v, r(v))$ is maintained. Distance halving continues until a *local* link is used, such that w.h.p. $h(u, t)$ contains less than $\rho \log^2 D$ nodes.

The local routing stage begins with *log* routing either from case 1 of the range adjustment stage or from the end of the distance halving stage. In both cases w.h.p. the areas (either $s_3(u, t)$ or $h(u, t)$) will contain at most $\rho \log^2 D$ nodes. Due to Lemma 5.3 the route will remain inside the sparse area, and from Lemma 5.6 the route will take at most $O(\log D)$ hops until a node u whose *log* link is outside $s_2(u, t)$ is reached.

Finally, the *closest* sub stage begins, and by Lemma 5.1, routing will remain inside the sparse area that contains at most $\log D$ nodes. The target will be reached in at most $\log D$ hops since each hop decreases the distance and so nodes will never be revisited. □

6. CONCLUSIONS

In this extended abstract we addressed an open question of Hassin and Peleg by showing that routing on the plane requires constant out degree instead of logarithmic. In the full paper we will extend our scheme to higher dimensional Euclidean metrics. The complexity measures of our scheme increase exponentially with the dimension. We will also show how to slightly alter the network design so that routing can proceed in purely greedy manner. In addition using $O(\log D)$ links we show how to reduce the diameter to $O(\log D / \log \log D)$.

We plan to further investigate the applicability of geometric routing for P2P networks. Specifically we have initial results for efficiently maintaining our scheme in a dynamic network and a scheme for making the network more fault tolerant.

Dynamism is also an issue for maintaining an Euclidean spanner [1] in a centralized manner. In the full paper we will show how our construction can be used to efficiently maintain a geometric spanner data structure against adversarial join and leave events that are oblivious to our scheme's randomization. This partially addresses the open question of Arya et al. [1].

From a practical point of view, the obvious open question is how to reduce the constants to a value that would give real world applicability.

Acknowledgements

The authors would like to thank Alex Samorodnitsky, Yair Bartal, and David Karger for useful discussions and the anonymous reviewers for their comments.

7. REFERENCES

- [1] S. Arya, D. M. Mount, and M. Smid. Dynamic algorithms for geometric spanners of small diameter: randomized solutions. *Comput. Geom. Theory Appl.*, 13(2):91–107, 1999.

- [2] J. Aspnes, Z. Diamadi, and G. Shah. Fault-tolerant routing in peer-to-peer systems. In *Proceedings of the Twenty-First ACM Symposium on Principles of Distributed Computing (PODC 02)*, pages 223–232, July 2002.
- [3] B. Awerbuch and D. Peleg. Routing with polynomial communication-space trade-off. *SIAM J. Discret. Math.*, 5(2):151–162, 1992.
- [4] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(6):609–616, 2001.
- [5] L. J. Cowen. Compact routing with minimum stretch. In *Proceedings of the tenth annual ACM-SIAM symposium on Discrete algorithms*, pages 255–260. Society for Industrial and Applied Mathematics, 1999.
- [6] R. Cox, F. Dabek, F. Kaashoek, J. Li, and R. Morris. Practical, distributed network coordinates. *SIGCOMM Comput. Commun. Rev.*, 34(1):113–118, 2004.
- [7] C. Gavoille. Routing in distributed networks: overview and open problems. *ACM SIGACT News*, 32(1):36–52, 2001.
- [8] C. Gavoille and D. Peleg. Compact and localized distributed data structures. *Distributed Computing*, 16(2-3):111–120, 2003.
- [9] K. P. Gummadi, S. Saroiu, and S. D. Gribble. King: Estimating latency between arbitrary internet end hosts. In *Proceedings of the SIGCOMM Internet Measurement Workshop (IMW 2002)*, 2002.
- [10] Y. Hassin and D. Peleg. Sparse communication networks and efficient routing in the plane. *Distributed Computing*, 14(4):205–215, 2001.
- [11] B. Karp and H. T. Kung. GPSR: greedy perimeter stateless routing for wireless networks. In *Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 243–254. ACM Press, 2000.
- [12] J.M. Keil and C.A. Gutwin. Classes of graphs which approximate the complete euclidean graph. *Discrete Computational Geometry*, 7:13–28, 1992.
- [13] J. Kleinberg. The Small-World Phenomenon: An Algorithmic Perspective. In *Proceedings of the 32nd ACM Symposium on Theory of Computing*, 2000.
- [14] E. Kranakis, H. Singh, and J. Urrutia. Compass routing on geometric networks. In *Proc. 11th Canadian Conference on Computational Geometry*, pages 51–54, Vancouver, August 1999.
- [15] F. Kuhn, R. Wattenhofer, and A. Zollinger. Asymptotically optimal geometric mobile ad-hoc routing. In *Proc. of the 6th international workshop on Discrete algorithms and methods for mobile computing and communications (Dial-M)*, pages 24–33. ACM Press, 2002.
- [16] F. Kuhn, R. Wattenhofer, and A. Zollinger. Worst-Case Optimal and Average-Case Efficient Geometric Ad-Hoc Routing. In *Proc. 4th ACM Int. Symposium on Mobile Ad-Hoc Networking and Computing (MobiHoc)*, 2003.
- [17] D. Malkhi, M. Naor, and D. Ratajczak. Viceroy: A scalable and dynamic emulation of the butterfly. In *Proceedings of the 21st ACM Symposium on Principles of Distributed Computing (PODC '02)*, pages 183–192, 2002.
- [18] G.S Manku, M. Bawa, and P. Raghavan. Symphony: Distributed hashing in small world. In *Proceedings of the 4th USENIX Symposium on Internet Technologies and Systems*, 2003.
- [19] E. Ng and H. Zhang. Predicting internet network distance with coordiantes-based approaches. In *Proceedings of IEEE Conference on Computer Communications (INFOCOM '02)*, 2002.
- [20] Y. Shavitt and T. Tankel. Big-bang simulation for embedding network distances in euclidean space. In *Proceedings of IEEE Conference on Computer Communications (INFOCOM '03)*, San Francisco, CA, 2003.
- [21] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the SIGCOMM 2001*, 2001.
- [22] K. Talwar. Bypassing the embedding: Approximation schemes and compact representations for growth restricted metrics. In *Preliminary version, to appear in STOC 2004*.
- [23] M. Thorup and U. Zwick. Approximate distance oracles. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 183–192. ACM Press, 2001.
- [24] M. Thorup and U. Zwick. Compact routing schemes. In *Proceedings of the thirteenth annual ACM symposium on Parallel algorithms and architectures*, pages 1–10. ACM Press, 2001.
- [25] S. Wilbur, T. H. Saleem, B. M. Pias, and J. Crowcroft. Lighthouses for scalable distributed location. In *The 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03)*, 2003.