

A Doubling Dimension Threshold $\Theta(\log \log n)$ for Augmented Graph Navigability^{*}

Pierre Fraigniaud^{1**}, Emmanuelle Lebhar^{2***}, and Zvi Lotker³

¹ CNRS, Laboratoire de Recherche en Informatique (LRI),
Université Paris-Sud, 91405 Orsay, France.

`pierre@lri.fr`

² LIP, École Normale Supérieure de Lyon, 46 allée d'Italie,
69364 Lyon Cedex 07, France.

`Emmanuelle.Lebhar@ens-lyon.fr`

³ Centrum voor Wiskunde en Informatica (CWI), Kruislaan 413,
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands.

`Z.Lotker@cwi.nl`

Abstract. In his seminal work, Kleinberg showed how to augment meshes using random edges, so that they become navigable; that is, greedy routing computes paths of polylogarithmic expected length between any pairs of nodes. This yields the crucial question of determining whether such an augmentation is possible for all graphs. In this paper, we answer negatively to this question by exhibiting a threshold on the doubling dimension, above which an infinite family of graphs cannot be augmented to become navigable whatever the distribution of random edges is. Precisely, it was known that graphs of doubling dimension at most $O(\log \log n)$ are navigable. We show that for doubling dimension $\gg \log \log n$, an infinite family of graphs cannot be augmented to become navigable. Finally, we complete our result by studying the special case of square meshes, that we prove to always be augmentable to become navigable.

Keywords: doubling dimension, small world, greedy routing.

1 Introduction

The *doubling dimension* [4, 13, 15] appeared recently as a key parameter for measuring the ability of networks to support efficient algorithms [7, 14] or to realize

^{*} The results contained in this paper were partially obtained when the two first authors were visiting CWI (Centrum voor Wiskunde en Informatica) in the Netherlands. Additional support from CWI.

^{**} Supported by the projects PairAPair of the ACI Masses de Données, and FRAGILE of the ACI Sécurité Informatique. Additional support from the INRIA project "Grand Large"

^{***} Supported by the CNRS Locoglobo grant and the project PairAPair of the ACI Masses de Données.

specific tasks efficiently [1, 2, 6, 24]. Roughly speaking, the doubling dimension of a graph G is the smallest d such that, for any integer $r \geq 1$, and for any node $u \in V(G)$, the ball $B(u, 2r)$ centered at u and of radius $2r$ can be covered by at most 2^d balls $B(u_i, r)$ centered at nodes $u_i \in V(G)$. (This definition can be extended to any metric, and, for instance, \mathbf{Z}^d with the ℓ_1 norm is of doubling dimension d). In particular, the doubling dimension has an impact on the analysis of the small world phenomenon [22], precisely on the expected performances of greedy routing in *augmented* graphs [17].

An augmented graph is a pair (G, φ) where G is an n -node graph, and φ is a collection of probability distributions $\{\varphi_u, u \in V(G)\}$. Every node $u \in V(G)$ is given an extra link pointing to some node v , called the *long range contact* of u . The link from a node to its long range contact is called a *long range link*. The original links of the graph are called *local links*. The long range contact of u is chosen at random according to φ_u as follows: $\Pr\{u \rightarrow v\} = \varphi_u(v)$. Greedy routing in (G, φ) is the oblivious routing protocol where the routing decision taken at the current node u for a message of destination t consists in (1) selecting a neighbor v of u that is the closest to t according to the distance in G (this choice is performed among all neighbors of u in G and the long range contact of u), and (2) forwarding the message to v . This process assumes that every node has a knowledge of the distances in G , or at least a good approximation of them. On the other hand, every node is unaware of the long range links added to G , but its own long range link. Hence the nodes have no notion of the distances in the augmented graph. Note that the knowledge of the distances in the underlying graph G is a reasonable assumption when, for instance, G is a network in which distances can be computed from the coordinates of the nodes (e.g., in meshes, as in [17]).

An infinite family of graphs $\mathcal{G} = \{G^{(i)}, i \in I\}$ is *navigable* if there exists a family $\Phi = \{\varphi^{(i)}, i \in I\}$ of collections of probability distributions, and a function $f(n) \in O(\text{polylog}(n))$ such that, for any $i \in I$, greedy routing in $(G^{(i)}, \varphi^{(i)})$ performs in at most $f(n^{(i)})$ expected number of steps where $n^{(i)}$ is the order of the graph $G^{(i)}$. More precisely, for any pair of nodes (s, t) of $G^{(i)}$, the expected number of steps $\mathbb{E}(\varphi^{(i)}, s, t)$ for traveling from s to t using greedy routing in $(G^{(i)}, \varphi^{(i)})$ is at most $f(n^{(i)})$. The *greedy diameter* of $(G^{(i)}, \varphi^{(i)})$ is defined as $\max_{s, t \in G^{(i)}} \mathbb{E}(\varphi^{(i)}, s, t)$.

In his seminal paper, Kleinberg [17] proved that, for any fixed integer $d \geq 1$, the family of d -dimensional meshes is navigable. Duchon et al [8] generalized this result by proving that any infinite family of graphs with bounded growth is navigable. Fraigniaud [11] proved that any infinite family of graphs with bounded treewidth is navigable. Finally, Slivkins [24] recently related navigability to doubling dimension by proving that any infinite family of graphs with doubling dimension at most $O(\log \log n)$ is navigable. All these results naturally lead to the question of whether all graphs are navigable.

Let $\delta : \mathbf{N} \mapsto \mathbf{N}$, let $\mathcal{G}_{n, \delta(n)}$ be the class of n -node graphs with doubling dimension at most $\delta(n)$, and let $\mathcal{G}_\delta = \cup_{n \geq 1} \mathcal{G}_{n, \delta(n)}$. By rephrasing Slivkins result [24], we get that \mathcal{G}_δ is navigable for any function δ bounded from above by $c \log \log n$

for some constant $c > 0$. This however lets open the case of graphs of larger doubling dimensions, namely the cases of all families \mathcal{G}_δ where δ is satisfying $\delta(n) \gg \log \log n$.

1.1 Our results

We prove a threshold of $\delta(n) = \Theta(\log \log n)$ for the navigability of \mathcal{G}_δ : below a certain function δ , \mathcal{G}_δ is navigable, while above it \mathcal{G}_δ is not navigable. More precisely, we prove that, for any function δ satisfying $\lim_{n \rightarrow \infty} (\log \log n) / \delta(n) = 0$, \mathcal{G}_δ is not navigable. Hence, the result in [24] is essentially the best that can be achieved by considering only the doubling dimension of graphs.

Our negative result requires to prove that for an infinite family of graphs in \mathcal{G}_δ , *any* distribution of the long range links leaves the expected number of steps of greedy routing above any polylogarithmic for some pairs of source and target. For this purpose, we exhibit graphs presenting a very high number of possible “directions” for a long range link to go. By a counting argument, we show that there exist pairs of source and target at distance greater than any polylogarithm, between which greedy routing does not use any long range link, whatever their distribution is. In other words, we exhibit an infinite family of graphs with non polylogarithmic greedy diameter for any augmentation. This negative results answers a question asked in [11, 18].

We also prove a somehow counter intuitive result by showing that a super-graph of a navigable graph is not necessarily navigable. In particular, we show that all square meshes are navigable, for all dimensions. Specifically, we prove that although the family of non navigable graphs that we use to disprove the navigability of all graphs contains the standard square meshes of dimension δ as subgraphs, this latter family of graphs is navigable.

1.2 Related works

Kleinberg showed that greedy routing performs in $O(\log^2 n)$ steps between any pair of nodes on d -dimensional meshes augmented by the d -harmonic distribution. I.e. the greedy diameter of these augmented meshes is $O(\log^2 n)$. Since then, several results have been developed to tighten the analysis of greedy routing on randomly augmented networks. Precisely, Barrière et al. [5] showed that the greedy diameter of the d -dimensional meshes augmented by the d -harmonic distribution is $\Theta(\log^2 n)$. In the special case of rings, Aspnes et al. [3] proved a lower bound on the greedy diameter of $\Omega(\log^2 n / \log \log n)$ for *any* augmentation. For paths, Flammini et al. [9] recently showed a lower bound on the greedy diameter of $\Omega(\log^2 n)$ in the case of symmetric and distance monotonic augmentations. Martel and Nguyen [21] showed however that the (standard) diameter of all these networks augmented by the harmonic distribution is $\Theta(\log n)$. In another perspective, several authors developed decentralized algorithms for the d -dimensional mesh augmented by the d -harmonic distribution. Lebar and Schabanel [19] presented a decentralized algorithm which performs in $O(\log n (\log \log n)^2)$ expected number of steps in this graph. The

algorithm Neighbor-Of-Neighbor presented by Manku et al. [20] performs in $O(\frac{1}{k \log k} (\log n)^2)$ expected number of steps, where k is the number of long range links per node in the mesh. Assuming some extra knowledge on the long range links, Fraigniaud et al. [12] described an oblivious routing which performs in $O((\log n)^{1+1/d})$ expected number of steps. Finally, Martel and Nguyen [21] presented a non oblivious routing protocol achieving the same performances under the same assumption as in [12].

1.3 Organization of the paper

The paper is organized as follows: Section 2 presents the main result of the paper by exhibiting the non navigability of graphs with doubling dimension $\gg \log \log n$. In Section 3, we study the special case of square meshes, and prove that they are all navigable.

2 Non navigable graphs

In this section, we prove that the result in [24] is essentially the best that can be achieved as far as doubling dimension is concerned.

Theorem 1. *Let $\delta : \mathbf{N} \mapsto \mathbf{N}$ be such that $\lim_{n \rightarrow \infty} \frac{\log \log n}{\delta(n)} = 0$. Then \mathcal{G}_δ is not navigable.*

Informally, the argument of the proof is that a doubling dimension $\gg \log \log n$ implies that the number of possible “directions” where a random link can go is greater than any polylogarithm of n . Therefore, for any trial of the long range links, there always exist a direction for which these long links do not help in the sense that there exist a source and a target between which greedy routing does not use any long range link.

Proof. We show that there exists an infinite family of graphs $\{G^{(n)}, n \geq 1\}$ indexed by their number of vertices, such that $G^{(n)} \in \mathcal{G}_{n, \delta(n)}$ and for any family $\Phi = \{\varphi^{(n)}, n \geq 1\}$ of collections of probability distributions, greedy routing in $(G^{(n)}, \varphi^{(n)})$ performs in an expected number of steps $t(n) \notin O(\text{polylog}(n))$ for some pairs of source and target.

Let $d : \mathbf{N} \mapsto \mathbf{N}$ be such that $d \leq \delta$, $\lim_{n \rightarrow \infty} \frac{\log \log n}{d(n)} = 0$, and $d(n) \leq \varepsilon \sqrt{\log n}$ for some $0 < \varepsilon < 1$. For the sake of simplicity, assume that $p = n^{1/d(n)}$ is integer. $G^{(n)}$ is the graph of n nodes consisting of $p^{d(n)}$ nodes labeled $(x_1, \dots, x_{d(n)})$, $x_i \in \mathbf{Z}_p$. Node $(x_1, \dots, x_{d(n)})$ is connected to all nodes $(x_1 + a_1, \dots, x_{d(n)} + a_{d(n)})$ where $a_i \in \{-1, 0, 1\}$, $i = 1, \dots, d(n)$, and all operations are taken modulo p (cf. Figure 1). Note that, by construction of $G^{(n)}$, the distance between two nodes $y = (y_1, \dots, y_{d(n)})$ and $z = (z_1, \dots, z_{d(n)})$ is $\max_{1 \leq i \leq d(n)} \min(|y_i - z_i|, p - |y_i - z_i|)$. Hence, the diameter of $G^{(n)}$ is $\lfloor p/2 \rfloor$.

Claim. $G^{(n)} \in \mathcal{G}_{n, \delta(n)}$.




Fig. 1. Example of graph $G^{(n)}$ defined in proof of Theorem 1 with $d(n) = 2$. Grey areas represent the various *directions* for the central node. The bold line represents a *diagonal* for the central node. Colored nodes belongs to a *line*.

Clearly $G^{(n)}$ has n nodes. We prove that $G^{(n)}$ has doubling dimension $d(n)$, therefore at most $\delta(n)$. Let $\mathbf{0} = (0, \dots, 0)$. The ball $B(\mathbf{0}, 2r)$ can be covered by $2^{d(n)}$ balls of radius r , centered at the $2^{d(n)}$ nodes $(x_1, \dots, x_{d(n)})$, $x_i \in \{-r, +r\}$ for any $i = 1, \dots, d(n)$. Hence the doubling dimension of $G^{(n)}$ is at most $d(n)$. On the other hand, $|B(\mathbf{0}, 2r)| = (4r+1)^{d(n)}$ and $|B(\mathbf{0}, r)| = (2r+1)^{d(n)}$. Thus at least $(4r+1)^{d(n)}/(2r+1)^{d(n)}$ balls are required to cover $B(\mathbf{0}, 2r)$, since in $G^{(n)}$, for any node u and radius r , $|B(u, r)| = |B(\mathbf{0}, r)|$. This ratio can be rewritten as $2^{d(n)}(1 - \frac{1}{2(2r+1)})^{d(n)}$. For $2r = \frac{n^{1/d(n)}}{5}$, since $d(n) \leq \sqrt{\log n}$, we get that $(2r+1) > \frac{2^{\sqrt{\log n}}}{5} > d(n)$ for $n \geq n_0$, $n_0 \geq 1$. Then, for $n \geq n_0$,

$$\begin{aligned} \left(1 - \frac{1}{2(2r+1)}\right)^{d(n)} &> \left(1 - \frac{1}{2d(n)}\right)^{d(n)} \\ &= 2^{d(n) \log(1 - \frac{1}{2d(n)})} \\ &\geq 2^{d(n) \left(-\frac{1}{2d(n)} - \frac{4}{4(d(n))^2}\right)} = 2^{-\frac{1}{2} - \frac{1}{d(n)}} \end{aligned}$$

There exists $n_1 \geq n_0$, such that $2^{-\frac{1}{2} - \frac{1}{d(n)}} > \frac{1}{2}$ for $n \geq n_1$. Then, for $n \geq n_1$, $|B(\mathbf{0}, 2r)|/|B(\mathbf{0}, r)| > 2^{d(n)-1}$. Thus the doubling dimension of $G^{(n)}$ is at least $d(n)$, which proves the claim. \diamond

Definition 1. For any node $u = (u_1, \dots, u_{d(n)})$, and for any $D = (\nu_1, \dots, \nu_{d(n)}) \in \{-1, 0, +1\}^{d(n)}$, we call *direction* the set of nodes

$$\text{dir}_u(D) = \{v = (v_1, \dots, v_{d(n)}) : v_i = (u_i + \nu_i \cdot x_i) \bmod p, 1 \leq x_i \leq \lfloor p/2 \rfloor\}.$$

Note that, for any u , the directions $\text{dir}_u(D)$ for $D \in \{-1, 0, +1\}^{d(n)}$ partition the nodes of $G^{(n)}$ (see Figure 1). There are obviously $3^{d(n)}$ directions, and the $2^{d(n)}$ directions defined on $\{-1, +1\}^{d(n)}$ have all the same cardinality.

Definition 2. For any node $u = (u_1, \dots, u_{d(n)})$, and for any $D = (\nu_1, \dots, \nu_{d(n)}) \in \{-1, +1\}^{d(n)}$, we call diagonal the set of nodes

$$\text{diag}_u(D) = \{v = (v_1, \dots, v_{d(n)}) : v_i = (u_i + \nu_i \cdot x) \bmod p, 1 \leq x \leq \lfloor p/2 \rfloor\}.$$

The next claim shows that long range links are useless for greedy routing along a diagonal if they are not going in the direction of the diagonal.

Claim. Let u be any node and let v be the long range contact of u for some distribution $\varphi^{(n)}$ of the long range links. Assume $v \in \text{dir}_u(D)$ and let $t \in \text{diag}_u(D')$ for $D, D' \in \{-1, +1\}^{d(n)}$, $D \neq D'$. Greedy routing from u to t does not use the long range link (u, v) .

Let $u = (u_1, \dots, u_{d(n)})$, $v = (v_1, \dots, v_{d(n)})$ and $t = (t_1, \dots, t_{d(n)})$. Since $t \in \text{diag}_u(D')$, there exists $x \in \{1, \dots, \lfloor p/2 \rfloor\}$ such that $|t_i - u_i| = x$ for all $1 \leq i \leq d(n)$. Since $D \neq D'$, there exists $j \in \{1, \dots, d(n)\}$ such that:

$$t_j = u_j + \alpha \cdot x \tag{1}$$

$$v_j = u_j - \alpha \cdot y, \tag{2}$$

for some $\alpha \in \{-1, +1\}$ and $y \in \{1, \dots, \lfloor p/2 \rfloor\}$. Then,

$$\text{dist}(v, t) \geq |t_j - v_j| = x + y > x = \text{dist}(u, t).$$

Therefore greedy routing from u to t does not use the long range link (u, v) , which proves the claim. \diamond

Consider now a distribution $\varphi^{(n)}$ of long range links that belongs to some given collection of probability distributions $\Phi = \{\varphi^{(n)}, n \geq 1\}$. We prove that routing on the diagonal is hard. More precisely, let s be a source node and t be a target node, $t \in \text{diag}_s(D)$ for some D . If any node $u \in \text{diag}_s(D)$ between s and t has its long range contact $v \in \text{dir}_u(D_v)$ for some $D_v \neq D$, then, from the previous claim, greedy routing from s to t does not use any long range link and thus takes $\text{dist}(s, t)$ steps. We prove that this phenomenon occurs for at least one pair (s, t) such that $\text{dist}(s, t) \geq 2^{d(n)} - 3$.

Definition 3. An interval $I = [a, b]$ is a connected subset of a diagonal. Precisely, $[a, b]$ is an interval of $\text{diag}_a(D)$ if $b \in \text{diag}_a(D)$ and

$$I = \{c \in \text{diag}_a(D) \mid \text{dist}(a, c) + \text{dist}(c, b) = \text{dist}(a, b)\}.$$

We say that an interval I of $\text{diag}_u(D)$ is *good* if there exists $x \in I$ such that the long range contact y of x satisfies $y \in \text{dir}_x(D)$.

Definition 4. A line L of $G^{(n)}$ in direction $D \in \{-1, +1\}^{d(n)}$ is a maximal subset of $V(G^{(n)})$ such that for any two nodes $u, v \in L$, we have

$$\text{diag}_u(D) \cap \text{diag}_v(D) \neq \emptyset.$$

The set of all the lines in the same direction D partitions $G^{(n)}$ into n/p lines of size p .

Let us partition each line into p/X disjoint intervals of same length X . This results into n/X intervals per direction, thus in total into a set S of $\frac{n}{X} \cdot 2^{d(n)}$ intervals of length X , since there are $2^{d(n)}$ directions defined in $\{-1, +1\}^{d(n)}$. We show that if X is too small, then there is at least one of all the intervals in S which is not good.

There is a one-to-one mapping between intervals and nodes in the following sense. Each good interval $I = [a, b] \in S$ must contain a node u whose long range contact v satisfies $v \in \text{dir}_u(D)$. The node u is called the *certificate* for I . Node u cannot be the certificate of any other interval $J \in S$ with $J \neq I$, even for those such that $J \cap I \neq \emptyset$ when I and J are in two distinct directions.

We have $2^{d(n)} \cdot \frac{n}{X}$ intervals in S . Since a certificate certifies the goodness of exactly one interval, $2^{d(n)} \cdot \frac{n}{X}$ has to be at most n , that is: $X \geq 2^{d(n)}$. By the pigeonhole principle, if $X < 2^{d(n)}$, there is one interval $I = [s, t] \in S$ which is not good. From Claim 2, greedy routing from s to t takes $X - 1$ steps.

Since $d(n) \leq \varepsilon \sqrt{\log n}$, we have:

$$p = n^{1/d(n)} \geq 2^{\frac{1}{\varepsilon} \sqrt{\log n}} \geq 2^{\varepsilon \sqrt{\log n}} - 2 \geq 2^{d(n)} - 2.$$

Therefore, the value $X = 2^{d(n)} - 2$ can be considered for our partitioning. In this case, we obtain that greedy routing from s to t takes $2^{d(n)} - 3$ steps.

We complete the proof of the theorem by proving the following claim.

Claim. $2^{d(n)} \notin O(\text{polylog } n)$.

Let $\alpha \geq 1$, we have:

$$\frac{\log^\alpha n}{2^{d(n)}} = 2^{\alpha \log \log n - d(n)} = 2^{\alpha d(n) \left(\frac{\log \log n}{d(n)} - \frac{1}{\alpha} \right)}.$$

Since $\lim_{n \rightarrow \infty} \frac{\log \log n}{d(n)} = 0$, there exists $n_1 \geq 1$ such that for any $n \geq n_1$, $\left(\frac{\log \log n}{d(n)} - \frac{1}{\alpha} \right) \leq -\frac{1}{2\alpha}$, and thus $\frac{\log^\alpha n}{2^{d(n)}} \leq 2^{-d(n)/2}$. Moreover, $d(n) \geq \log \log n$, then, for $n \geq n_1$,

$$\frac{\log^\alpha n}{2^{d(n)}} \leq 2^{-\frac{\log \log n}{2}} = o(1).$$

In other words, $2^{d(n)}$ is not a polylogarithm of n , which proves the claim. \diamond

Remark. Note that the proof of Theorem 1 does not assume independent trials for the long range links.

3 Navigability of meshes

The family of non navigable graphs defined in the proof of Theorem 1 contains the standard square meshes of dimension $d(n)$ as subgraphs, where $d(n) \gg \log \log n$. Nevertheless, and somehow counter intuitively, a supergraph of a navigable graph is not necessarily navigable. In this section, we illustrate this phenomenon by focusing on the special case of d -dimensional meshes, the first graphs that were considered for the analysis of navigable networks [17]. Precisely, we show that any d -dimensional torus $C_{n^{1/d}} \times \dots \times C_{n^{1/d}}$ is navigable: either it has a polylogarithmic diameter, or it admits a distribution of links such that greedy routing computes paths of polylogarithmic length. This result has partially been proven in [9] for the case of constant dimensions. We give here a complete proof that holds for any dimension.

Theorem 2. *For any positive function $d(n)$, the n -node $d(n)$ -dimensional torus is navigable.*

Proof. We construct a random link distribution φ as follows. Let $u = (u_1, \dots, u_{d(n)})$ and $v = (v_1, \dots, v_{d(n)})$ be two nodes. If they differ in more than one coordinate, then $\varphi_u(v) = 0$; otherwise, i.e. they differ in only one coordinate, say the i th, then:

$$\varphi_u(v) = \frac{1}{d(n)} \cdot \frac{1}{2H_k} \cdot \frac{1}{|u_i - v_i|},$$

where $k = \frac{n^{1/d}}{2}$ and $H_k = \sum_{j=1}^k \frac{1}{j}$ is the harmonic sum. Note that this distribution corresponds to:

- picking a dimension uniformly at random (probability $\frac{1}{d(n)}$ to pick dimension i)
- and to draw a long-range link on this axis according to the 1-harmonic distribution over distances ($\frac{1}{2H_k}$ is the normalizing coefficient for this distribution), which is the distribution given by Kleinberg to make the 1-dimensional torus navigable.

Let now $s = (s_1, \dots, s_{d(n)})$ and $t = (t_1, \dots, t_{d(n)})$ be a pair of source and target in the mesh. Assume that the current message holder during an execution of greedy routing is $x = (x_1, \dots, x_{d(n)})$, at distance X from t . The probability that x has a long range link to some node $w = (x_1, \dots, x_{i-1}, w_i, x_{i+1}, \dots, x_{d(n)})$, $1 \leq i \leq d(n)$ such that $|t_i - w_i| \leq |t_i - x_i|/2$, is at least $\sum_{1 \leq i \leq d(n)} \frac{1}{3d(n)H_k} = \frac{1}{3H_k}$, along the same analysis as the analysis of Kleinberg one dimensional model, summing over the dimensions. If such a link is found, it is always preferred to the local contact of x that only reduces one of the coordinate by 1. Thus, after at most $3H_k$ steps on expectation, one of the coordinates has been divided by two. Note that since long range links only get to nodes that differs in a single coordinate from their origin, further steps cannot increase $|x_i - t_i|$ for any $1 \leq i \leq d(n)$, x being the current message holder. Repeating the analysis for all coordinates, we thus get that after $3d(n)H_k$ steps on expectation, all

the coordinates have been divided by at least two, and so the current distance to the target is at most $X/2$. Finally, the algorithm reaches t after at most $3d(n)H_k \log(\text{dist}(s,t))$ steps on expectation, which is $O((\log^2 n)/d(n))$.

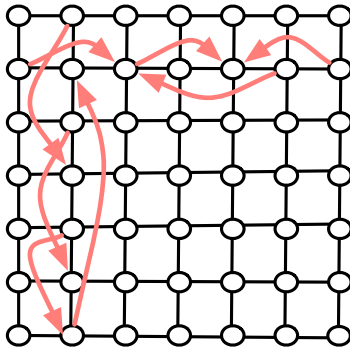


Fig. 2. Example of 2-dimensional mesh augmented as in proof of Theorem 2: 1-harmonic distribution of links on each axis. Bold links are long range links, they are not all represented.

Remark. Note that our example of non-navigability in Section 3 may appear somehow counter intuitive in contrast to our latter construction of long range links on meshes. Indeed, why not simply repeating such a construction on the graph $G^{(n)}$ defined in the proof Theorem 1? That is, why not selecting long range contacts on each "diagonal" using the 1-harmonic distribution, in which case greedy routing would perform efficiently between pairs (s,t) on the diagonals? This cannot be done however because, to cover all possible pairs (s,t) on the diagonals, $2^{d(n)}$ long range links per node would be required, which is larger than any polylogarithm of n when $d(n) \gg \log \log n$.

4 Conclusion

The increasing interest in graphs and metrics of bounded doubling dimension arises partially from the hypothesis that large real graphs do present a low doubling dimension (see, e.g., [10, 16] for the Internet). Under such an hypothesis, efficient compact routing schemes and efficient distance labeling schemes designed for bounded doubling dimension graphs would have promising applications. On the other hand, the navigability of a network is actually closely related to the existence of efficient compact routing and distance labeling schemes on the network. Indeed, long range links can be turned into small labels, e.g. via the technique of rings of neighbors [24]. Interestingly enough, our paper emphasizes

that the small doubling dimension hypothesis of real networks is crucial. Indeed, for doubling dimension above $\log \log n$, networks may become not navigable. It would therefore be important to study precisely to which extent real networks do present a low doubling dimension.

In a more general framework, our result of non navigability shows that the small world phenomenon, in its algorithmic definition of navigability, is not only due to the good spread of additional links over distances in a network, but is also highly dependent of the base metric itself, in particular in terms of dimensionality.

Peleg recently proposed the more general question of f -navigability. For a function f , we say that a n -node graph G is f -navigable if there exists a distribution φ of long range links such that the greedy diameter of the augmented graph (G, φ) is at most $f(n)$. From [23], all n -node graphs are \sqrt{n} -navigable by giving an uniform random distribution of the long range links. From Theorem 1, we get as a corollary that, for all graphs to be f -navigable, $f(n) = \Omega(2^{\sqrt{\log n}})$. It thus remains to close the gap between these upper and lower bounds for the f -navigability of arbitrary graphs.

Acknowledgments. We are thankful to Ph. Duchon for having pointed to us an error in an earlier version of this paper.

References

1. I. Abraham, C. Gavoille, A. V. Goldberg, and D. Malkhi. Routing in networks with low doubling dimension. In *26th International Conference on Distributed Computing Systems (ICDCS)*, to appear, 2006.
2. I. Abraham and D. Malkhi. Name independent routing for growth bounded networks. In *17th Annual ACM Symposium on Parallel Algorithms and Architecture (SPAA)*, pages 49–55, 2005.
3. J. Aspnes, Z. Diamadi, and G. Shah. Fault-tolerant routing in peer-to-peer systems. In *In 21st ACM Symp. on Principles of Distributed Computing (PODC)*, pages 223–232, 2002.
4. P. Assouad. Plongements lipshitzien dans \mathbf{R}^n . *Bull. Soc. Math.*, 111(4):429–448, 1983.
5. L. Barrière, P. Fraigniaud, E. Kranakis, and D. Krizanc. Efficient routing in networks with long range contacts. In *LNCS Proceedings of 15th International Symposium on Distributed Computing (DISC)*, volume 2180, pages 270–284, 2001.
6. H.T-H. Chan, A. Gupta, B. M. Maggs, and S. Zhou. On hierarchical routing in doubling metrics. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pages 762–771, 2005.
7. K. L. Clarkson. *Nearest-Neighbor Methods for Learning and Vision: Theory and Practice*, chapter Nearest-neighbor searching and metric space dimensions. (Survey). MIT Press, T. Darrell, P. Indyk, G. Shakhnarovich, and P. Viola, editors, 2006.
8. P. Duchon, N. Hanusse, E. Lebhar, and N. Schabanel. Could any graph be turned into a small world? *Theoretical Computer Science*, Special issue on Complex Networks, volume 355(1), pages 96–103, 2006.

9. M. Flammini, L. Moscardelli, A. Navarra, and S. Perennes. Asymptotically optimal solutions for small world graphs. In *LNCS Proceedings of the 19th International Symposium on Distributed Computing (DISC)*, volume 3724, pages 414–428, 2005.
10. M. Fomenkov, K. Claffy, B. Huffaker, and D. Moore. Macroscopic internet topology and performance measurements from the dns root name servers. In *USENIX LISA*, 2001.
11. P. Fraigniaud. Greedy routing in tree-decomposed graphs: a new perspective on the small-world phenomenon. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*, pages 791–802, 2005.
12. P. Fraigniaud, C. Gavoille, and C. Paul. Eclecticism shrinks even small worlds. In *Proceedings of the 23rd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 169–178, 2004.
13. A. Gupta, R. Krauthgamer, and J.R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 534–543, 2003.
14. S. Har-Peled and M. Mendel. Fast construction of nets in low dimensional metrics and their applications. In *Proceedings of the 21th ACM Symposium on Computational Geometry (SoCG)*, pages 150–158, 2005.
15. J. Heinonen. *Lectures on analysis on metric spaces*. Springer-Verlag, New York, 2001.
16. D. R. Karger and M. Ruhl. Finding nearest-neighbors in growth-restricted metrics. In *Proceedings of the 34th annual ACM symposium on Theory of computing (STOC)*, pages 741–750, 2002.
17. J. Kleinberg. The Small-World Phenomenon: An Algorithmic Perspective. In *Proceedings of the 32nd ACM Symposium on Theory of Computing (STOC)*, pages 163–170, 2000.
18. J. Kleinberg. Complex networks and decentralized search algorithm. In *Proceedings of the International Congress of Mathematicians (ICM)*, 2006. To appear.
19. E. Lebar and N. Schabanel. Close to optimal decentralized routing in long-range contact networks. *Theoretical Computer Science special issue on ICALP'04*, 348(2-3):294–310, 2005.
20. G. S. Manku, M. Naor, and U. Wieder. Know thy neighbor's neighbor: the power of lookahead in randomized p2p networks. In *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC)*, pages 54–63, 2004.
21. C. Martel and V. Nguyen. Analyzing Kleinberg's (and other) small-world models. In *In 23rd ACM Symp. on Principles of Distributed Computing (PODC)*, pages 178–187, 2004.
22. S. Milgram. The small world problem. *Psychology Today*, 61(1), 1967.
23. D. Peleg. Private communication, 2006.
24. A. Slivkins. Distance estimation and object location via rings of neighbors. In *Proceedings of the 24th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 41–50, 2005.