Graph Augmentation via Metric Embedding

Emmanuelle Lebhar[1] and Nicolas Schabanel[2]

[1] CNRS and CMM-Univesidad de Chile, Chile
elebhar@dim.uchile.cl
[2] CNRS and Université Paris Diderot, France
Nicolas.Schabanel@liafa.jussieu.fr

Abstract. Kleinberg [17] proposed in 2000 the first random graph model achieving to reproduce small world navigability, i.e. the ability to greedily discover polylogarithmic routes between any pair of nodes in a graph, with only a partial knowledge of distances. Following this seminal work, a major challenge was to extend this model to larger classes of graphs than regular meshes, introducing the concept of *augmented graphs navigability*. In this paper, we propose an original method of augmentation, based on metrics embeddings. Precisely, we prove that, for any $\varepsilon > 0$, any graph G such that its shortest paths metric admits an embedding of distorsion γ into $\mathbb{R}^d$ can be augmented by one link per node such that greedy routing computes paths of expected length $O(\frac{1}{\varepsilon}\gamma^d \log^{2+\varepsilon} n)$ between any pair of nodes with the only knowledge of G. Our method isolates all the structural constraints in the existence of a good quality embedding and therefore enables to enlarge the characterization of augmentable graphs.

Keywords: Small world, metrics embedding, greedy routing.

1 Introduction

The small world effect, or "six degrees of separation", is the well known property observed in social networks [9,21] that any pair of nodes in these networks is connected by a very short chain of acquaintances (typically polylogarithmic in the size of the network), that, moreover, can be discovered locally. In the literature, a *small world graph* can either refer to this property or to a graph with polylogarithmic diameter and high clustering (see e.g. [23]). In this paper, a small world graph refers to a graph of polylogarithmic diameter *and* whose short paths can be discovered locally, i.e. which is *navigable*. This surprising property has gained a lot of interest recently since Kleinberg [17] introduced the first analytical graph model for navigability, and because of its potential in the design of large decentralized networks with efficient routing schemes. The model proposed by Kleinberg in 2000 consists in a d-dimensional mesh augmented by one extra random link in each node, distributed according to the d-harmonic distribution. The local search is then modeled by greedy routing, which is the simple algorithm that, at each node, forwards the message to the neighbor that is the

T.P. Baker, A. Bui, and S. Tixeuil (Eds.): OPODIS 2008, LNCS 5401, pp. 217–225, 2008.

closest to the destination *in the mesh*. Kleinberg demonstrates that greedy routing computes paths of expected length $\Theta(\log^2 n)$ between any pair of nodes in his model, with the only knowledge of the distances in the mesh: the augmented mesh is *navigable*

Following this seminal work, a major challenge was to extend this model to larger classes of graphs than regular meshes, i.e. to determine which n-node graphs G admit an augmentation with one link in each node such that greedy routing with the only of G computes polylog(n)-length paths between any pair in the augmented graph. Kleinberg [18] and Duchon et al. [7] showed that this is possible for all graphs of bounded growth, i.e. where, for any node u and radius $r \geq 1$, the $2r$-neighborhood of u is of size at most a constant times its r-neighborhood. Fraigniaud [10] demonstrates that any bounded treewidth graph can also be augmented by one link per node to become navigable, and Abraham and Gavoille [4] showed that, more generally, this is possible for all graphs excluding a fixed minor. The definition of the problem can directly be extended to metric spaces by asking which n-points metric spaces[1] $M = (V, \delta)$ can be augmented by $O(\log n)$ links such that, in the resulting graph, greedy routing computes polylog(n) routes between any pair with the only knowledge of M. In this framework, Slivkins [22] showed that any doubling metric can be augmented to become navigable. A doubling metric is a metric where, for all $r \geq 1$, any ball of radius $2r$ can be covered by at most C balls of radius r, for some constant C.

However, it was recently proven by Fraigniaud et al. [13] that such an augmentation is not possible for all graphs: there exist an infinite familiy of n-node graphs on which any distribution of augmented links will leave the greedy paths of expected length $\Omega(n^{1/\sqrt{\log n}})$ for some pairs. The best upper bound valid for arbitrary graphs up to our knowledge is an expected length $\tilde{O}(n^{1/3})$ between any pair, due to Fraigniaud et al. [12], with some specific link augmentation. The remaining gap between these two bounds is today still open and leaves a question mark on the limiting characteristics of a metric for the navigability augmentation.

Orthogonally to the navigability question, studies on embeddings of metric spaces have known huge developments this last decade (cf. Chapter 15 of [20] for a review), due in particular to their applications in approximation algorithms [15] and more recently in handling efficiently large decentralized networks [6]. An embedding σ of a metric $M = (V, \delta)$ into a metric $M' = (V', \delta')$ is an injective function σ on V into V'. Its quality is characterized by the distorsion it induces on the distances. For the sake of simplicity, we consider only non-contracting embeddings, we then say that σ has distorsion γ if and only iff for any $u, v \in V$, $\delta'(\sigma(u), \sigma(v)) \leq \gamma \cdot \delta(u, v)$. Crucial networking problems like routing, resource location or nearest neighbor search are easy to handle on a low dimensional euclidean space. However, large real networks like the Internet do not present

[1] A metric space $M = (V, \delta)$ is a set of points V associated with a distance function δ. Therefore, any weighted graph naturally defines a metric M on its set of nodes V with the distance function δ being the length of a shortest path between two nodes.

such a simple structure. The increasing interest for metrics embeddings comes therefore partially from the fact that, if the embedding is of good quality, it can provide a way to develop efficient algorithms on complex, or even arbitrary, metric spaces, by solving them on a simple metric space that approximates them well (cf. e.g. [14,15]). In addition, many good quality embeddings are computed with randomized local algorithms that only require a distance oracle, making them particularly appropriate to the large decentralized networks setting (cf. e.g. [5] for a seminal example).

In this paper, we propose a new way to tackle the augmented graphs navigability problem through the metric embedding setting.

1.1 Our Contribution

We introduce a generalized augmentation process. The main feature of our augmentation process is to use an embedding of the input graph shortest paths metric into a metric that is easy to augment into a navigable graph. This distinction between the augmentation process in itself (handled on the "easy" metric) and the structural characteristics of the input (captured by the embedding quality) provides a new way to characterize the classes of navigable graphs. We consider embedding into $(\mathbb{R}^d, \ell_p)$ which is the d-dimensional euclidean space associated to the ℓ_p norm, for $d, p \geq 1$: for any $u = (u_1, \ldots, u_d)$ and $v = (v_1, \ldots, v_d)$, we have $||u - v||_p = (\sum_{i=1}^{d} |u_i - v_i|^p)^{1/p}$. We prove the following theorem:

Theorem 1. *Let $p, n, \gamma, d \geq 1$. For any $\varepsilon > 0$, any n-node graph G whose shortest path metric $M = (V, \delta)$ admits an embedding of distorsion γ into $(\mathbb{R}^d, \ell_p)$ can be augmented with one link per node such that greedy routing in the resulting graph computes paths of expected length $O(\frac{1}{\varepsilon}\gamma^d \log^{2+\varepsilon} n)$ between any pair with the only knowledge of M.*

For instance, using the recent embedding result of Abraham et al. [3], we get as an immediate corollary that, for any $0 < \varepsilon \leq 1$ and any $n \geq 1$, any n-node graph G of doubling dimension D (cf. [14]) can be augmented so that the expected lengths of all greedy paths is $O((\log^{(1+\varepsilon)} n)^{O(D/\varepsilon)} \log^2 n) = O((\log n)^{O(D)})$ with the only knowledge of G. This provides a more direct proof to the fact that bounded doubling dimension graphs are navigable (proved in [22]).

Intuitively, if the metric considered is not too far from a metric M which can be easily augmented, we use a low distorsion embedding of the metric into M, draw the random links in M, and then map back appropriately these links to the original metric so that they will still be useful shortcuts for greedy routing.

Moreover, the design of the augmented links in our process can be done in a fully decentralized way and only requires to know the embedding. In the case where the chosen embedding is itself local (like e.g. the seminal Bourgain embedding [5] if a distance oracle is available), we thus provide an algorithm which locally adds one address to each routing table in a network and guarantees a small number of hops decentralized routing between any pair for a large class of input graphs.

2 A Universal Augmentation Process via Metric Embedding

In this section, we present our augmentation process that adds one directed link per node. This process is universal in the sense that it only requires as an input the base graph (arbitrary) and an embedding function of this graph into $\mathbb{R}^d_{\ell_p}$, for some $p, d \geq 1$. Such a function exists for any graph and therefore the algorithm is not restricted to a specific graph class. However, as we will see in the next section, the analysis of greedy routing might give a poor routing time result if the embedding is not of good quality. There exists lower bound results on arbitrary metric embedding quality. A typical example is that embedding some n-node constant degree expander graph into $\mathbb{R}^d_{\ell_p}$ requires distortion $\Omega(\log n)$ [20] and dimension $d = \Omega(\log n)$ [2]. Nevertheless, expander graphs are always navigable without any augmentation given their polylogarithmic diameter.

The augmentation algorithm is based on the well known augmentation of d-dimensional meshes of the Kleinberg model, where the shortcuts are distributed according to the d-harmonic distribution. The idea is to map back these links to the original set of nodes. Given that not all the extremities of the shortcuts added in ℓ^d_p are images of the original nodes, this requires some careful rewiring.

Augmentation Process $\mathcal{AP}$
Input: An n-node graph $G = (V, E)$, an embedding σ of its shortest path metric $M = (V, \delta)$ into ℓ^d_p, and a constant $\varepsilon > 0$;
Output: G augmented with one directed link in each node.
Begin
For each $u \in V$ **do**
Pick a point $\tau_u \in \mathbb{R}^d_{\ell_p}$ with probability density:

$$\frac{1}{Z}\,\frac{1}{\big(||\sigma(u)-\tau||_p\big)^d \ln^{1+\varepsilon}(||\sigma(u)-\tau||_p+e)},$$

over all $\tau \in \mathbb{R}^d_{\ell_p}$.
Add a directed link from u to $v \in V$ where v is the node such that $\sigma(v)$ is the closest point to τ_u in $\sigma(V)$.
End.

Note: e stands here for exp(1) and is only used to allow distance to be zero in the formula. Z is the normalizing factor of the probability density described: $Z = \int_{t>0}^{\infty} \frac{S(t)}{t^d \ln^{1+\varepsilon}(t+e)} dt$, where $S(t)$ is the surface of an hypersphere of raius t in $\mathbb{R}^d$. Figure 1 illustrates the process $\mathcal{AP}$.

3 Navigability of Graphs Augmented with $\mathcal{AP}$

In this section, we demonstrate our main result. The *intrinsic dimension* [3], or *doubling dimension* [1] of a graph G characterizes its geometric property, this is

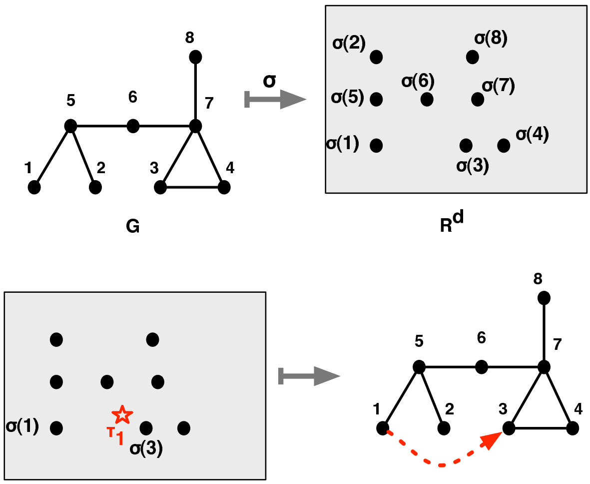

Fig. 1. Illustration of the augmented link from vertex 1 to 3 with process $\mathcal{AP}$

the minimum constant α such that any ball in G can be covered by at most 2^α balls of half the radius. We show that, if a graph has low intrinsic dimension, $\mathcal{AP}$ process provides augmented shortcuts that enables navigability. We have the following theorem:

Theorem 2. *Let $p, n, \gamma, d \geq 1, \varepsilon > 0$, G an n-node graph and σ an embedding of distorsion γ of the shortest path metric M of G into $(\mathbb{R}^d, \ell_p)$. Then, greedy routing in $\mathcal{AP}(G, \sigma, \varepsilon)$ computes paths of expected length at most $O(\frac{1}{\varepsilon}\gamma^d \log^{2+\varepsilon} n)$ between any pair, with the only information of the distances in M.*

Proof. In order to analyze greedy routing performances in $\mathcal{AP}(G, \sigma, \varepsilon)$, we begin by analyzing some technical properties of the probability distribution of the chosen points τ in $(\mathbb{R}^d, \ell_p)$. For any $u \in G$, we say that τ_u, as defined in algorithm $\mathcal{AP}$, is the *contact point* of u.

Let Z be the normalizing factor of the contact points distribution. We have:

$$Z = \int_{t>0}^{\infty} \frac{S(t)}{t^d \ln^{1+\varepsilon}(t+e)} dt,$$

where $S(t)$ stands for the surface of a sphere of radius t in $\mathbb{R}^d_{\ell_p}$. This surface is at most $c_p \cdot \big(2^d/(d-1)!\big) \cdot t^{d-1}$, where $c_p > 1$ is a constant depending on p. It follows:

$$Z \leq c_p \cdot \frac{2^d}{(d-1)!} \int_{t>1}^{\infty} \frac{dt}{t \ln^{1+\varepsilon}(t+e)} \leq c_p \cdot \frac{(1+e)}{\varepsilon} \cdot \frac{2^d}{(d-1)!}.$$

Let s and $t \in G$ be the source and the target of greedy routing in $\mathcal{AP}(G, \sigma, \varepsilon)$. Let $M = (V, \delta)$ be the shortest paths metric of G. Let v be the current node of greedy routing, and let $1 \leq i \leq \lceil \log \delta(s,t) \rceil$ such that $\delta(v,t) \in [2^{i-1}, 2^i)$.

Since σ has distorsion γ, we have:

$$\delta(v,t) \leq ||\sigma(v) - \sigma(t)||_p \leq \gamma \cdot \delta(v,t).$$

Let $X = ||\sigma(v) - \sigma(t)||_p$, and let $\mathcal{E}$ be the event: "$||\tau_v - \sigma(t)||_p \leq X/(4\gamma)$". Let $L(v)$ be the contact of v (i.e. the closest point to τ_v in $\sigma(V)$.

Claim. If $\mathcal{E}$ occurs, then $\delta(L(v),t) \leq \delta(v,t)/2$.

Indeed, assume that $\mathcal{E}$ occurs. From the triangle inequality, we have:

$$||\sigma(L(v)) - \sigma(t)||_p \leq ||\sigma(L(v)) - \tau_v||_p + ||\tau_v - \sigma(t)||_p.$$

And since $\sigma(L(v)) = \tau$ is closer to τ_v than $\sigma(t)$ by definition of $\mathcal{AP}$, we get:

$$||\sigma(L(v)) - \sigma(t)||_p \leq 2||\tau_v - \sigma(t)||_p \leq X/(2\gamma).$$

Finally:

$$\delta(L(v),t) \leq ||\sigma(L(v)) - \sigma(t)||_p \leq X/(2\gamma) \leq \delta(v,t)/2. \qquad \diamond$$

Claim. The probability that $\mathcal{E}$ occurs is greater than

$$C\,\frac{\varepsilon}{d5^d\gamma^d}\,\frac{1}{\ln^{1+\varepsilon}(2\gamma\delta(v,t)+e)},$$

for some constant $C > 0$.

Proof of the claim. Let P be the probability that $\mathcal{E}$ occurs. P is the probability that τ_v belongs to the ball of radius $X/(4\gamma)$ centered at $\sigma(t)$ in $\mathbb{R}^d_{\ell_p}$. Let $\mathcal{B}$ be this ball. We have, by definition of $\mathcal{AP}$:

$$\begin{aligned} P &= \frac{1}{Z}\int_{\tau\in\mathcal{B}} \frac{1}{(||\sigma(v)-\tau||_p)^d \ln^{1+\varepsilon}(||\sigma(v)-\tau||_p + e)} \\ &\geq \frac{1}{Z}\int_{\tau\in\mathcal{B}} \frac{1}{\big((1+1/(4\gamma))X\big)^d \ln^{1+\varepsilon}((1+1/(4\gamma))X+e)}, \end{aligned}$$

since $(1+1/(4\gamma))X$ is the largest distance from $\sigma(v)$ to any point in $\mathcal{B}$.

On the other hand, the volume of $\mathcal{B}$ is at least $c'_p \cdot \frac{2^d}{d!}(X/(4\gamma))^d$, for some constant $c'_p > 0$. We get:

$$\begin{aligned} P &\geq \frac{1}{Z}\cdot\frac{c'_p 2^d (X/4\gamma)^d}{d!(1+\frac{1}{4\gamma})^d X^d}\cdot\frac{1}{\ln^{1+\varepsilon}((1+\frac{1}{4\gamma})X+e)} \\ &\geq \frac{c'_p}{c_p(1+e)}\cdot\frac{\varepsilon}{d5^d}\cdot\frac{1}{\gamma^d}\cdot\frac{1}{\ln^{1+\varepsilon}((1+\frac{1}{4\gamma})X+e)} \\ &\geq C\,\frac{\varepsilon}{d5^d\gamma^d}\,\frac{1}{\ln^{1+\varepsilon}(2\gamma\delta(v,t)+e)}. \qquad \diamond \end{aligned}$$

Claim. If the current node v of greedy routing satisfies $\delta(v,t) \in [2^{i-1}, 2^i)$ for some $1 \leq i \leq \lceil \log \delta(s,t) \rceil$, then after $O(\frac{1}{\varepsilon}\gamma^d(i-1)^{1+\varepsilon})$ steps on expectation, greedy routing is at distance less than 2^{i-1} from t.

Proof of the claim. Combining the claims, we get that, with probability $\Omega([\frac{1}{\varepsilon}\gamma^d \ln^{1+\varepsilon}(\gamma\delta(v,t))]^{-1})$ (where the Ω notation hides a linear factor in ε), the contact $L(v)$ of v is at distance at most 2^{i-1} to t. If this does not occur, greedy routing moves to a neighbor v' at distance strictly less than $\delta(v,t)$ to t and strictly greater than 2^{i-1} and we can repeat the same argument. Therefore, after $O(\frac{1}{\varepsilon}\gamma^d \ln^{1+\varepsilon}(\gamma\delta(v,t))) = O(\frac{1}{\varepsilon}\gamma^d (i-1)^{1+\varepsilon})$ steps, greedy routing is at distance less than 2^{i-1} to t with constant probability. $\diamond$

Finally, from this last claim, the expected number of steps of greedy routing from s to t is at most:

$$\sum_{i=1}^{\log(\delta(s,t))} O(\gamma^d (i-1)^{1+\varepsilon}) = O(\frac{1}{\varepsilon}\gamma^d \log^{2+\varepsilon} n).$$

From this theorem, results giving new insights on the navigability problem can be derived from the very recent advances in metric embeddings theory. In particular, graphs of bounded doubling dimension, that subsumes graphs of bounded growth, received an increasing interest recently. They are of particular interest for scalable and distributed network applications since it is possible to decompose them greedily into clusters of exponentially decreasing diameter.

Corollary 1. *For any $\varepsilon > 0$, any n-node graph G of bounded doubling dimension α can be augmented with one link per node so that greedy routing compute paths of expected length $O(\frac{1}{\varepsilon}\log^{(2+\varepsilon+2\alpha)} n)$ between any pair of vertices with the only knowledge of G.*

Indeed, from Theorem 1.1 of [3], it is known that, for every n-point metric space M of doubling dimension α and every $\theta \in (0,1]$, there exists an embedding of M into $\mathbb{R}^d_{\ell_p}$ with distorsion $O(\log^{1+\theta} n)$ and dimension $O(\alpha/\theta)$. Taking $\theta = 1$ gives the corollary. This result was previously proved in [22] by another method of augmentation, using "rings of neighbor". The originality of our method is that it is not specific to a given graph or metric class, this dependency lying only in the embedding function. Therefore, it enables to get more direct proofs that a graph is augmentable into a navigable small world than previous ones.

This new kind of augmentation process via embedding is also promising to derive lower bounds on metrics embedding quality. Indeed, since not all graphs can be augmented to become navigable, necessarily, if there exists a positive result on small world augmentation via some embedding, then this embedding cannot keep the same quality for all graphs. For the particular case of Theorem 2, we derive that any injective function σ that embeds any arbitrary metric into $\mathbb{R}^d_{\ell_p}$ with distorsion γ has to satisfy $\gamma^d = \tilde{\Omega}(n^{1/\sqrt{\log n}})$. This lower bound is however subsumed by the bound provided by the Johnson-Lindenstrauss flattening lemma [16]: $\gamma^d = O((1+\varepsilon)^{\log n/\varepsilon^2}) = O(n^{(1+\varepsilon)/\varepsilon^2})$ for any $0 < \varepsilon < 1$, which is essentially tight (cf. e.g. [20]).

It is worth to note that Fraigniaud and Gavoille [11] recently tackled the question of navigating in a graph that has been augmented using the distances in a spanner[2] of this graph. They remarked that greedy routing usually requires to know the spanner map of distances in order to achieve an efficient routing. On the contrary, our augmentation process does not requires greedy routing to be aware of distances in $\mathbb{R}^d$. This is due to the geography of the spaces considered: an embedding of a graph in $\mathbb{R}^d$ preserves geographical neighboring regions.

4 Discussion

The result presented in this paper gives new perspectives in the understanding of networks small world augmentations. Indeed, the augmentation process $\mathcal{AP}$ isolates all the dependencies on the graph structure in the embedding function.

On the other hand, such an augmentation process focuses on the geography of the graph and cannot capture the augmentation processes that are based on graph separator decomposition. It can be distinguished two main kinds of augmentation processes in the navigable networks literature. One kind of augmentation relies on the graph density and its similarity with a mesh (like augmentations in [7,17,18,22]), while the other kind relies on the existence of good separators in the graph (like augmentations in [4,10]). Augmentation via embedding cannot be directly extended to augmentations using separators because of the difficulty to handle the distortion in the analysis of greedy routing. Finally, the extension of $\mathcal{AP}$ to graphs that are close to a tree metric (using embeddings into tree metrics) could open the path to the exhaustive characterization of graph classes that can be augmented to become navigable, as well as provide new lower and upper bounds on embeddings as side results. More generally, the exhaustive characterization of the graphs that can be augmented to become navigable is still an important open problem, as well as the design of good quality embeddings into low dimensional spaces.

References

1. Assouad, P.: Plongements lipschitzien dans $\mathbb{R}^n$. Bull. Soc. Math. France 111(4), 429–448 (1983)
2. Abraham, I., Bartal, Y., Neiman, O.: Advances in metric embedding theory. In: Proceeeding of the the 38th annual ACM symposium on Theory of Computing (STOC), pp. 271–286 (2006)
3. Abraham, I., Bartal, Y., Neiman, O.: Embedding Metric Spaces in their Intrinsic Dimension. In: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms (SODA), pp. 363–372 (2008)
4. Abraham, I., Gavoille, C.: Object location using path separators. In: Proceedings of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing (PODC), pp. 188–197 (2006)
5. Bourgain, J.: On Lipschitz embedding of finite metric spaces in Hilbert space. Israel Journal of Mathematics 52, 46–52 (1985)

[2] A k-spanner of a graph G is a subgraph G' such that for any $u, v \in G$, $\mathrm{dist}_{G'}(u,v) \leq k\mathrm{dist}_G(u,v)$.

6. Dabek, F., Cox, R., Kaashoek, F., Morris, R.: Vivaldi: A decentralized network coordinate system. In: ACM SIGCOMM (2004)
7. Duchon, P., Hanusse, N., Lebhar, E., Schabanel, N.: Could any graph be turned into a small-world? Theoretical Computer Science 355(1), 96–103 (2006)
8. Duchon, P., Hanusse, N., Lebhar, E., Schabanel, N.: Towards small world emergence. In: 18th Annual ACM Symp. on Parallel Algorithms and Architectures (SPAA), pp. 225–232 (2006)
9. Dodds, P.S., Muhamad, R., Watts, D.J.: An experimental study of search in global social networks. Science 301, 827–829 (2003)
10. Fraigniaud, P.: Greedy routing in tree-decomposed graphs: a new perspective on the small-world phenomenon. In: Brodal, G.S., Leonardi, S. (eds.) ESA 2005. LNCS, vol. 3669, pp. 791–802. Springer, Heidelberg (2005)
11. Fraigniaud, P., Gavoille, C.: Polylogarithmic network navigability using compact metrics with small stretch. In: 20th Annual ACM Symp. on Parallel Algorithms and Architectures (SPAA), pp. 62–69 (2008)
12. Fraigniaud, P., Gavoille, C., Kosowski, A., Lebhar, E., Lotker, Z.: Universal Augmentation Schemes for Network Navigability: Overcoming the $\sqrt{n}$-Barrier. In: Proceedings of the 19th Annual ACM Symposium on Parallel Algorithms and Architecture (SPAA), pp. 1–7 (2007)
13. Fraigniaud, P., Lebhar, E., Lotker, Z.: A Doubling Dimension Threshold $\Theta(\log \log n)$ for Augmented Graph Navigability. In: Azar, Y., Erlebach, T. (eds.) ESA 2006. LNCS, vol. 4168, pp. 376–386. Springer, Heidelberg (2006)
14. Gupta, A., Krauthgamer, R., Lee, J.R.: Bounded geometries, fractals, and low-distortion embeddings. In: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 534–543 (2003)
15. Indyk, P.: Algorithmic aspects of geometric embeddings. In: Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science, FOCS (2001)
16. Johnson, W.B., Lindenstrauss, J.: Extensions of Lipschitz maps into a Hilbert space. Contemporary mathematics 26, 189–206 (1984)
17. Kleinberg, J.: The Small-World Phenomenon: An Algorithmic Perspective. In: 32nd ACM Symp. on Theo. of Comp. (STOC), pp. 163–170 (2000)
18. Kleinberg, J.: Small-World Phenomena and the Dynamics of Information. Advances in Neural Information Processing Systems (NIPS) 14 (2001)
19. Kleinberg, J.: Complex networks and decentralized search algorithm. In: Intl. Congress of Math, ICM (2006)
20. Matousek, J.: Lectures on Discrete Geometry. Graduate Texts in Mathematics, vol. 212. Springer, Heidelberg (2002)
21. Milgram, S.: The Small-World Problem. Psychology Today, 60–67 (1967)
22. Slivkins, A.: Distance estimation and object location via rings of neighbors. In: 24th Annual ACM Symp. on Princ. of Distr. Comp. (PODC), pp. 41–50 (2005)
23. Watts, D.J., Strogatz, S.H.: Collective dynamics of 'small-world' networks. Nature 393, 440–442 (1998)