Fault Tolerant Additive Spanners

Gilad Braunschvig*, Shiri Chechik*, and David Peleg*

Department of Computer Science and Applied Mathematics,
The Weizmann Institute, Rehovot, Israel
{gilad.braunschvig,shiri.chechik,david.peleg}@weizmann.ac.il

Abstract. Graph spanners are sparse subgraphs that preserve the distances of the original graph, up to some small multiplicative factor or additive term (known as the *stretch* of the spanner). A number of algorithms are known for constructing sparse spanners with small multiplicative or additive stretch. Recently, the problem of constructing *fault-tolerant multiplicative* spanners for general graphs was given some algorithms. This paper addresses the analogous problem of constructing *fault tolerant additive* spanners for general graphs.

We establish the following general result. Given an n-vertex graph G, if H_1 is an ordinary additive spanner for G with additive stretch α, and H_2 is a fault tolerant multiplicative spanner for G, resilient against up to f edge failures, with multiplicative stretch μ, then $H = H_1 \cup H_2$ is an additive fault tolerant spanner of G, resilient against up to f edge failures, with additive stretch $O(\tilde{f}(\alpha + \mu))$ where $\tilde{f}$ is the number of failures that have actually occurred ($\tilde{f} \leq f$).

This allows us to derive a poly-time algorithm $\texttt{Span}_{add}^{f-t}$ for constructing an additive fault tolerant spanner H of G, relying on the existence of algorithms for constructing fault tolerant multiplicative spanners and (ordinary) additive spanners. In particular, based on some known spanner construction algorithms, we show how to construct for any n-vertex graph G an additive fault tolerant spanner with additive stretch $O(\tilde{f})$ and size $O(fn^{4/3})$.

1 Introduction

1.1 Background and Motivation

The concept of spanners is a generalization of the notion of spanning trees. A spanner of a given graph is a subgraph that faithfully preserves the distances of the original graph. Two widely studied types of spanners are *multiplicative* spanners and *additive* spanners. A multiplicative spanner of the graph G is a subgraph H that preserves the distances between any two vertices in G up to a constant multiplicative factor (referred to as the stretch of the spanner), whereas an additive spanner of G preserves distances up to a constant additive term.

* Supported in part by the Israel Science Foundation (grant 894/09), the United States-Israel Binational Science Foundation (grant 2008348), the Israel Ministry of Science and Technology (infrastructures grant), and the Citi Foundation.

M.C. Golumbic et al. (Eds.): WG 2012, LNCS 7551, pp. 206–214, 2012.

More formally, a subgraph $H = (V, E_H)$ is a μ-multiplicative spanner of the graph $G = (V, E_G)$ if $E_H \subseteq E_G$ and $dist(u, v, H) \leq \mu \cdot dist(u, v, G)$ for every $u, v \in V$, where $dist(u, v, G')$ for a graph G' is the distance between u and v in G'. Similarly, a subgraph $H = (V, E_H)$ is an α-additive spanner of the graph $G = (V, E_G)$ if $E_H \subseteq E_G$ and $dist(u, v, H) \leq dist(u, v, G) + \alpha$ for every $u, v \in V$.

Additive spanners provide, in some sense, a much stronger guarantee than multiplicative ones, especially when dealing with long routes, because the *penalty* in taking the alternative route offered by the spanner is not proportional to the length of the original one, but bounded by a fixed term. Clearly any graph is a 1-multiplicative spanner and a 0-additive spanner of itself, so usually we are interested in computing spanners that are compact in the number of edges.

This paper considers settings in which the underlying graph G may occasionally suffer edge failures. In such settings, we are interested in *fault tolerant spanners*, both in the case of multiplicative and in the case of additive. These are spanners that keep the locality properties even after a number of faults occur. This robustness is important in systems that are prone to local malfunctions, like for example broken links in communication networks.

We say that a subgraph $H = (V, E_H)$ is a (μ, f)-multiplicative fault tolerant spanner of the graph $G = (V, E_G)$ if for every $F = \{e_1, \ldots, e_f\} \subseteq E_G$ and $u, v \in V$, $dist(u, v, H \setminus F) \leq \mu \cdot dist(u, v, G \setminus F)$.

Analogously, we define the notion of additive fault tolerant spanners as follows. A subgraph $H = (V, E_H)$ is an (α, f)-additive fault tolerant spanner of graph $G = (V, E_G)$ if for every $F = \{e_1, \ldots, e_f\} \subseteq E_G$ and $u, v \in V$, $dist(u, v, H \setminus F) \leq dist(u, v, G \setminus F) + \alpha$.

Fault tolerant spanners were first considered by Levcopoulos, Narasimhan and Smid [11] in the context of geometric graphs (where the nodes are assumed to be in the Euclidean space and the distance between every two nodes is the Euclidean distance between them). Levcopoulos et al. [11] presented efficient constructions for fault tolerant spanners with $(1 + \epsilon)$ multiplicative stretch. The size of the spanner was later improved by Lukovszki [12] and then by Czumaj and Zhao [6].

Constructions for *multiplicative* fault tolerant spanners for general graphs that are robust to edge or vertex failures were presented in [5], later the construction for vertex failures was improved in [7]. In this paper we show a construction for *additive* fault tolerant spanners. We deal only with edge failures. Our result relies on the existence of fault tolerant multiplicative spanners and (ordinary) additive spanners and uses algorithms for constructing such spanners as subroutines.

1.2 Our Results

In this paper we prove the following general construction scheme.

Theorem 1. *Let $G = (V, E)$ be a general graph, $H_1 = (V, E_1)$ be an α-additive spanner of G and H_2 be a μ-multiplicative fault tolerant spanner of G, resilient against up to f edge failures. Then $H = H_1 \cup H_2$ is an α'-additive fault tolerant spanner resilient against up to f failures, with additive stretch $\alpha' \leq O(\tilde{f}(\mu + \alpha))$ where $\tilde{f} \leq f$ is the number of actual faults.*

Note that the stretch guarantee depends on the number of failures that have actually occurred. Hence if no failures occur, we get a stretch bound of α, independent of f, and the stretch deteriorates as the actual number of faults increases.

As a corollary, relying on existing spanner construction algorithms, we prove that for any graph $G = (V, E)$ there exists a poly-time constructible α'-additive fault tolerant spanner $H = (V, E')$, resilient against up to f edge failures, with additive stretch $\alpha' \leq O(\tilde{f})$ and size $|E'| \leq O(fn^{4/3})$.

1.3 Related Work

Graph Spanners were first introduced by Peleg and Ullman [13] as a technique for generating synchronizers. Later, spanners were used in various contexts including routing in communication networks and distributed systems [14,17], broadcasting [10], distance oracles [3,18], etc.

It is well known how to construct $(2k-1)$-multiplicative spanners with $O(n^{1+1/k})$ edges [2]. This size-stretch tradeoff is also conjectured to be optimal.

The picture for additive spanners is far from being complete, basically there are two known constructions for additive spanners. Aingworth et al. [1] presented a construction for 2-additive spanner with $O(n^{3/2})$ edges (for further follow-up see [8,9,19,16]). Later, Baswana et al. [4] presented an efficient construction for 6-additive spanner with $O(n^{4/3})$ edges.

In lack of truly understanding the complete picture for additive spanners, many papers consider the problem of constructing spanners with either non-constant additive stretch or with both multiplicative and additive stretch (e.g., [9,19,15,4]).

In order to achieve the constants mentioned above, we make use of existing constructions of ordinary additive spanners and multiplicative fault tolerant spanners. In practice, we may use the construction for additive spanners presented in [1,4] and the construction for multiplicative fault tolerant spanners presented in [5].

2 Preliminaries

Denote by $dist(u, v, G)$ the distance between u and v in G (if there is no path from u to v in G then $dist(u, v, G) = \infty$). Denote by $SP(u, v, G)$ the shortest path between u and v in G (if there is no path from u to v in G then $SP(u, v, G) = \emptyset$, if there is more than one such path then choose one arbitrarily). For a simple path P, denote by $|P|$ the number of edges in P. For a path P in the graph and vertices x, y on this path, denote by $P[x, y]$ the subpath of P from x to y. For a graph $G = (V, E)$ and a set of edges F, denote by $G \setminus F$ the graph $G' = (V, E \setminus F)$. Throughout this paper, when talking about fault tolerant additive spanners we distinguish between f, the maximum number of faults that the spanner can tolerate while keeping its stretch promise, and $\tilde{f}$, the number of edges that actually fail. The *size* of a graph $G(V, E)$ is defined to be its number of edges, $|E|$.

3 Constructing (α, f)-Additive Fault Tolerant Spanners

3.1 The Construction

We start by describing the algorithm for constructing a fault tolerant additive spanner and continue with the analysis of the worst case additive stretch guaranteed by this construction. We rely on the existence of known algorithms $\texttt{Span}_{add}$ constructing an α-additive spanner for a given graph G (for certain values of α), and $\texttt{Span}_{mult}^{f-t}$ constructing a (μ, f)-multiplicative fault tolerant spanner for G (for certain values of μ) cf. [5,1,4].

Algorithm $\texttt{Span}_{add}^{f-t}$

1. Invoke Algorithm $\texttt{Span}_{add}$ to generate an α-additive spanner H_1 of G
2. Invoke Algorithm $\texttt{Span}_{mult}^{f-t}$ to generate a (μ, f)-multiplicative fault tolerant spanner H_2 of G
3. $H \leftarrow H_1 \cup H_2$
4. Return H

3.2 Analysis

We next analyze the additive stretch of the subgraph H constructed by Algorithm $\texttt{Span}_{add}^{f-t}$, and prove that it is bounded by a constant linear in μ, α and $\tilde{f}$, the number of actual failures.

Our analysis proceeds as follows. We inspect the shortest path P between two vertices s and t in the graph $G \setminus F$ and distinguish several key points on that path. Then we show that the additive spanner H_1 provides for each pair of these key points a fault-free detour that is not too long. In other parts along the path P we use the fault tolerant multiplicative spanner H_2 in order to progress while avoiding faults. Finally we show that the union of all of these detours provides a path in the constructed spanner H that is completely free of faults and is close in length to the shortest path P (up to an additive term).

Consider a source vertex s, a target vertex t and a set of $\tilde{f}$ edge faults $F = \{e_1, \ldots, e_{\tilde{f}}\}$ $(\tilde{f} \leq f)$. Let $P = SP(s, t, G \setminus F)$ be the shortest path from s to t after the failure event. Denote by $p(v)$ the *position* of v on P, where $p(v) = 0$ if v is the first vertex on P and $p(v) = |P|$ if v is the last vertex on P. Since H is a spanner of G, every pair of vertices $w_1, w_2 \in P$ s.t. $p(w_1) < p(w_2)$, has an alternative path in H. We refer to the shortest such path $SP(w_1, w_2, H)$, as the *bypass* of w_1 and w_2 in H.

We classify the bypasses as follows. If the bypass contains an edge in F, we say that the pair (w_1, w_2) belongs to *class* (u, v) if the first faulty edge that occurs on $SP(w_1, w_2, H)$ starting from w_1 is (u, v). Note that we take into consideration the direction of the edge, i.e., for every undirected edge e we have two different classes, one for each direction. For every pair of vertices $w_1, w_2 \in P$ s.t. $p(w_1) < p(w_2)$, if $SP(w_1, w_2, H)$ does not use any edge of F, we say that the pair (w_1, w_2) is of class Φ.

Note that if the pair (w_1, w_2) is of class Φ, then

$$\begin{aligned} dist(w_1, w_2, H \setminus F) &= dist(w_1, w_2, H) \\ &\le dist(w_1, w_2, G) + \alpha \\ &\le dist(w_1, w_2, G \setminus F) + \alpha \ , \end{aligned}$$

and therefore

$$dist(w_1, w_2, H \setminus F) \ \le \ |P[w_1, w_2]| + \alpha \ .$$

Next, order all pairs of vertices $(w_1, w_2) \in P$ s.t. $p(w_1) < p(w_2)$ in a lexicographic order according to the value $(p(w_1), p(w_2))$.

Lemma 1. *Let x_1, x_2 and y_1, y_2 be two pairs of vertices on path P of the same class (v, u) and $p(x_1) < p(x_2) \le p(y_1) < p(y_2)$. Then*

$$dist(x_1, y_1, H \setminus F) \le |P[x_1, y_1]| + 2\alpha.$$

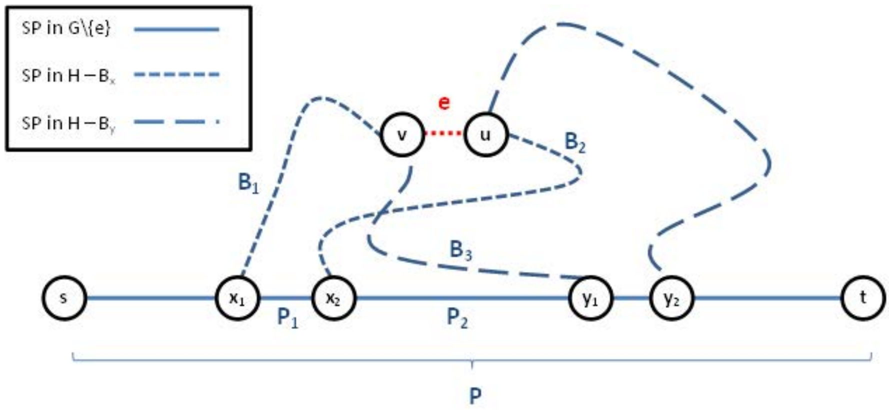

Fig. 1. Bypasses of class (v, u)

Proof. Consider the bypass $B_x = SP(x_1, x_2, H)$, $B_y = SP(y_1, y_2, H)$, and the subpaths $B_1 = B_x[x_1, v]$, $B_2 = B_x[v, x_2]$, $B_3 = B_y[y_1, v]$, $P_1 = P[x_1, x_2]$, $P_2 = P[x_2, y_1]$ (see Figure 1). By the definition of the class (v, u), the paths B_1 and B_3 do not contain any faults. Therefore,

$$dist(x_1, y_1, H \setminus F) \ \le \ |B_1| + |B_3| \ . \tag{1}$$

Since H contains H_1 and H_1 is an additive spanner of G, $dist(w_1, w_2, H) \le |Q| + \alpha$ for any two nodes w_1, w_2 and any path Q from w_1 to w_2 in G. In particular,

$$|B_1| + |B_2| \ \le \ |P_1| + \alpha \tag{2}$$

and also

$$|B_3| \;\leq\; |B_2| + |P_2| + \alpha\ . \tag{3}$$

Using Inequalities (1), (2) and (3), we get that

$$\begin{aligned} dist(x_1, y_1, H \setminus F) &\leq |B_1| + |B_3| \\ &\leq |B_1| + |B_2| + |P_2| + \alpha \\ &\leq |P_1| + |P_2| + 2\alpha \\ &= |P[x_1, y_1]| + 2\alpha \end{aligned}$$

Lemma 2. *Let (x_1, x_2) be the first pair in the lexicographic order of class different than Φ and let its class be (v, u). Let (y_1, y_2) be the last pair of class (v, u) in P. Then*

$$dist(x_1, y_1, H \setminus F) \leq |P[x_1, y_1]| + 2\alpha.$$

Proof. Note that $p(x_1) \leq p(y_1)$, since (x_1, x_2) is the first pair of class (v, u). If there is only one pair of class (v, u), then the analysis is the same as if $p(x_1) = p(y_1), p(x_2) = p(y_2)$. We consider two cases. The first case is where $p(y_1) < p(x_2)$. Then the pair (x_1, y_1) is of class Φ (because it appears before the pair (x_1, x_2) in the lexicographic order and (x_1, x_2) is the first pair of class different than Φ). It follows that $dist(x_1, y_1, H \setminus F) \leq |P[x_1, y_1]| + \alpha$. The second case is where $p(x_2) \leq p(y_1)$, and then it follows from Lemma 1 that $dist(x_1, y_1, H \setminus F) \leq |P[x_1, y_1]| + 2\alpha$.

Claim. Let (x_1, x_2) be the first pair in the lexicographic order of class different than Φ and let its class be (v, u). Let (y_1, y_2) be the last pair of the class (v, u), and let s_1 be the neighbor of y_1 on the path $P[y_1, y_2]$. Then either $dist(s, s_1, H \setminus F) \leq |P[s, s_1]| + 2\alpha + \mu - 1$ or $dist(s, t, H \setminus F) \leq |P| + \alpha$.

Proof. If the class of pair (s, t) is Φ, then the bypass from s to t in H contains no failures, so $dist(s, t, H \setminus F) = dist(s, t, H) \leq |P| + \alpha$ and we are done. So now suppose the pair (s, t) is not of class Φ. Then $x_1 = s$ since otherwise $p(x_1) > p(s)$ in contradiction to the assumption that (x_1, x_2) is the first pair of class different than Φ. According to Lemma 2,

$$dist(s, y_1, H \setminus F) \;\leq\; |P[s, y_1]| + 2\alpha\ . \tag{4}$$

Since H contains H_2, which is a (μ, f)-multiplicative fault tolerant spanner of G,

$$dist(y_1, s_1, H \setminus F) \;\leq\; |P[y_1, s_1]| \cdot \mu = 1 \cdot \mu \;=\; |P[y_1, s_1]| + \mu - 1\ . \tag{5}$$

Combining Inequalities (4) and (5), we get that

$$\begin{aligned} dist(s, s_1, H \setminus F) &\leq dist(s, y_1, H \setminus F) + dist(sy_1, s_1, H \setminus F) \\ &\leq |P[s, y_1]| + |P[y_1, s_1]| + 2\alpha + \mu - 1 \\ &= |P[s, s_1]| + 2\alpha + \mu - 1 \end{aligned}$$

Lemma 3. *Let N be the number of classes on $SP(s,t,G \setminus F)$. Then*

$$dist(s,t,H \setminus F) \leq dist(s,t,G \setminus F) + N(2\alpha + \mu - 1) + \alpha \ .$$

Proof. We prove the lemma by induction on N. For $N = 0$, the pair (s,t) is of class Φ and the lemma holds. Assume that the lemma holds for any $n < N$. By Claim 3.2, either $dist(s,t,H \setminus F) \leq |P| + \alpha$ in which case we are done, or

$$dist(s,s_1,H \setminus F) \ \leq \ |P[s,s_1]| + 2\alpha + \mu - 1 \ . \qquad (6)$$

Notice that the path $P[s_1,t]$ does not contain any pair of class (v,u). It follows that the number of classes on the path $P[s_1,t]$ is smaller than N, and clearly $P[s_1,t]$ is the shortest path from s_1 to t on $G \setminus F$. Therefore the induction assumption holds for the path $P[s_1,t]$, and it follows that

$$dist(s_1,t,H \setminus F) \ \leq \ dist(s_1,t,G \setminus F) + (N-1)(2\alpha + \mu - 1) + \alpha \ . \qquad (7)$$

Combining Inequalities (6) and (7), we get that

$$\begin{aligned} dist(s,t,H \setminus F) &\leq dist(s,s_1,H \setminus F) + dist(s_1,t,H \setminus F) \\ &\leq dist(s,s_1,G \setminus F) + dist(s_1,t,G \setminus F) \\ &\quad +(2\alpha + \mu - 1) + (N-1)(2\alpha + \mu - 1) + \alpha \\ &= dist(s,t,G \setminus F) + N(2\alpha + \mu - 1) + \alpha \ . \end{aligned}$$

Theorem 2. *H is an (α', f)-additive fault tolerant spanner of G with $\alpha' = O(\tilde{f}(\alpha + \mu))$, and its size is $|E(H)| = |E(H_1)| + |E(H_2)|$.*

Proof. The size bound is immediate from the construction. Since there are at most $2\tilde{f}$ different classes (excluding Φ), Lemma 3 implies that $dist(s,t,H \setminus F) \leq dist(s,t,G \setminus F) + 2\tilde{f}(2\alpha + \mu - 1) + \alpha$.

A poly-time algorithm $\texttt{Span}_{add}$ for constructing a 6-additive spanner of size $O(n^{4/3})$ for any n-vertex graph G is presented in [4]. In [5] a poly-time algorithm $\texttt{Span}^{f-t}_{mult}$ for constructing, for any n-vertex graph, a (μ, f)-multiplicative fault tolerant spanner of size $O(fn^{1+\frac{2}{\mu+1}})$ for every odd μ and every f. Using these two results and Theorem 2, choosing $\mu = 5$, yields the following,

Corollary 1. *For every f, every graph G contains a (poly-time constructible) (α', f)-additive fault tolerant spanner of size $O(fn^{4/3})$ with $\alpha' = 32\tilde{f} + 6$.*

4 Conclusions and Open Problems

Although the concept of spanners is well established and bounds have been proven for fault tolerance in the case of multiplicative spanners, up until now there were no known constructions or lower bounds on the space and stretch of fault tolerant *additive* spanners. Hopefully this paper will open the door for more research in the field, as it leaves open several interesting problems. Our construction is relatively simple and uses previously known constructions as a *black box*. This leaves the possibility that there might exist a more sophisticated construction for fault tolerant additive spanners, with stretch that is sublinear in the number of faults f. Moreover, our analysis deals only with edge failures, and future research may focus on overcoming vertex failures. Finally, it would be interesting to consider *fault tolerant* (α, β)*-spanners*. For example, by simply applying our construction and analysis and using any construction for (α, β)-spanners and (μ, f)-fault tolerant multiplicative spanner as building blocks, one can present an (α', β')-spanner that is robust to f faults, where $\alpha' = \alpha^2$ and $\beta' = O(\tilde{f}(\alpha\beta + \mu))$, but this is by no means known to be the best possible.

References

1. Aingworth, D., Chekuri, C., Indyk, P., Motwani, R.: Fast estimation of diameter and shortest paths (without matrix multiplication). SIAM J. Comput. 28(4), 1167–1181 (1999)
2. Althöfer, I., Das, G., Dobkin, D., Joseph, D., Soares, J.: On sparse spanners of weighted graphs. Discrete & Computational Geometry 9, 81–100 (1993)
3. Baswana, S., Kavitha, T.: Faster algorithms for approximate distance oracles and all-pairs small stretch paths. In: Proc. IEEE Symp. on Foundations of Computer Science (FOCS), pp. 591–602 (2006)
4. Baswana, S., Kavitha, T., Mehlhorn, K., Pettie, S.: Additive spanners and (α, β)-spanners. ACM Trans. on Algo. 7, A.5 (2010)
5. Chechik, S., Langberg, M., Peleg, D., Roditty, L.: Fault-tolerant spanners for general graphs. In: Proc. 41st ACM Symp. on Theory of Computing (STOC), pp. 435–444 (2009)
6. Czumaj, A., Zhao, H.: Fault-tolerant geometric spanners. Discrete & Computational Geometry 32, 2004 (2003)
7. Dinitz, M., Krauthgamer, R.: Fault-Tolerant Spanners: Better and Simpler. In: 30th ACM Symp. on Principles of Distributed Computing (PODC), pp. 169–178 (2011)
8. Dor, D., Halperin, S., Zwick, U.: All-pairs almost shortest paths. SIAM J. Computing 29(5), 1740–1759 (2000)
9. Elkin, M., Peleg, D.: $(1 + \epsilon, \beta)$-spanner constructions for general graphs. SIAM J. Computing 33(3), 608–631 (2004)
10. Farley, A.M., Proskurowski, A., Zappala, D., Windisch, K.: Spanners and message distribution in networks. Discrete Applied Mathematics 137(2), 159–171 (2004)
11. Levcopoulos, C., Narasimhan, G., Smid, M.: Efficient algorithms for constructing fault-tolerant geometric spanners. In: Proc. 30th ACM Symp. on Theory of Computing (STOC), pp. 186–195 (1998)
12. Lukovszki, T.: New Results on Fault Tolerant Geometric Spanners. In: Dehne, F., Gupta, A., Sack, J.-R., Tamassia, R. (eds.) WADS 1999. LNCS, vol. 1663, pp. 193–204. Springer, Heidelberg (1999)
13. Peleg, D., Ullman, J.D.: An optimal synchronizer for the hypercube. SIAM J. Computing 18(2), 740–747 (1989)
14. Peleg, D., Upfal, E.: A trade-off between space and efficiency for routing tables. J. ACM 36(3), 510–530 (1989)
15. Pettie, S.: Low distortion spanners. ACM Transactions on Algorithms 6(1) (2009)

16. Roditty, L., Thorup, M., Zwick, U.: Deterministic Constructions of Approximate Distance Oracles and Spanners. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 261–272. Springer, Heidelberg (2005)
17. Thorup, M., Zwick, U.: Compact routing schemes. In: Proc. 13th ACM Symp. on Parallel Algorithms and Architectures (SPAA), pp. 1–10 (2001)
18. Thorup, M., Zwick, U.: Approximate distance oracles. J. ACM 52(1), 1–24 (2005)
19. Thorup, M., Zwick, U.: Spanners and emulators with sublinear distance errors. In: 17th Symp. on Discrete Algorithms (SODA), pp. 802–809. ACM-SIAM (2006)