

Computing the Greedy Spanner in Near-quadratic Time^{*}

Prosenjit Bose¹, Paz Carmi¹, Mohammad Farshi¹, Anil Maheshwari¹, and Michiel Smid¹

School of Computer Science, Carleton University, Ottawa, ON, K1S 5B6, Canada.
jit@scs.carleton.ca, {paz,mfarshi}@cg.scs.carleton.ca,
{anil,michiel}@scs.carleton.ca

Abstract. It is well-known that the greedy algorithm produces high quality spanners and therefore is used in several applications. However, for points in d -dimensional Euclidean space, the greedy algorithm has cubic running time. In this paper we present an algorithm that computes the greedy spanner (spanner computed by the greedy algorithm) for a set of n points from a metric space with bounded doubling dimension in $\mathcal{O}(n^2 \log n)$ time using $\mathcal{O}(n^2)$ space. Since the lower bound for computing such spanners is $\Omega(n^2)$, the time complexity of our algorithm is optimal to within a logarithmic factor.

1 Introduction

A *network* on a point set V is a connected graph $G(V, E)$. When designing a network several criteria are taken into account. In particular, in many applications it is important to ensure a fast connection between every pair of points. For this it would be ideal to have a direct connection between every pair of points—the network would then be a complete graph—but in most applications this is unacceptable due to the very high costs associated with constructing such networks. This leads to the concept of spanners, as defined below.

Let (V, \mathbf{d}) be a metric space and $G(V, E)$ be a network on V such that the weight of each edge $e \in E$ is equal to the distance between its endpoints. We say that G is a *t -spanner* of V , for some constant $t > 1$, if for each pair of points $u, v \in V$, there exists a path in G between u and v of length at most $t \cdot \mathbf{d}(u, v)$. The *dilation* or *stretch factor* of G is the minimum t for which G is a t -spanner of V . Spanners were introduced by Peleg and Schäffer [12] in the context of distributed computing and by Chew [3] in the geometric context. Since then spanners have received a lot of attention, see [10, 11].

A classical algorithm for computing a geometric spanner for any set V of n points in \mathbb{R}^d , where \mathbf{d} is the Euclidean metric, and for any fixed $t > 1$, is the *greedy algorithm*, proposed independently by Bern in 1989 and Althöfer *et al.* [1]. The main steps of this algorithm are the following: first sort all the pairs of points in V with respect to their distances in increasing order and initialize the greedy

^{*} Research supported in part by NSERC and MRI.

graph $G(V, E)$ so that its edge set is empty. Next, the pairs are processed in sorted order. Processing a pair (u, v) entails a shortest path query in G between u and v . If there is no t -path between u and v (a path of length at most $t \cdot \mathbf{d}(u, v)$) in G then (u, v) is added to G , otherwise it is discarded. We will refer to the graph G generated by the greedy algorithm as the *greedy spanner*.

The greedy algorithm as stated above performs $\binom{n}{2}$ single-source shortest path queries. By employing Dijkstra’s single-source shortest path algorithm, the time complexity is $\mathcal{O}(mn^2 + n^3 \log n)$ using $\mathcal{O}(n^2)$ space, where n is the number of points and m is the number of edges in the spanner. It has been shown that for any set V of n points in \mathbb{R}^d and for any fixed $t > 1$ in the Euclidean metric, the greedy spanner has $\mathcal{O}(n)$ edges, bounded degree, and its total weight is $\mathcal{O}(wt(MST(V)))$, where $wt(MST(V))$ is the weight of the minimum spanning tree of V [4, 10]. Unfortunately, the naïve implementation of the greedy algorithm runs in cubic time.

Due to the high time complexity of computing the greedy spanner, researchers have proposed algorithms for computing other types of sparse t -spanners, see [10]. But it turns out that in practice the greedy algorithm produces t -spanners of high quality in comparison to other spanners [5, 6]. For example, they have been used for protein visualization as a low-weight data structure, which is used as a contact map, that allows approximate reconstruction of the full distance matrix [13].

For points in the plane under the Euclidean metric, Farshi and Gudmundsson [5, 6] introduced a speed-up strategy that generates the greedy spanner much faster in practice. They conjectured that their algorithm runs in $\mathcal{O}(n^2 \log n)$ time. However, as we will show in this paper, this conjecture is incorrect. They also showed that the greedy algorithm produces graphs whose size, weight, maximum degree and number of crossings are superior to the graphs produced from the other approaches which have near linear time complexity. For $t = 2$, $t = 1.1$ and $t = 1.05$ the number of edges in the greedy t -spanner is approximately $2n$, $4n$ and $6n$ respectively, which is surprisingly small. For comparison it is interesting to note that the Delaunay triangulation has approximately $3n$ edges and dilation bounded by 2.42 [9]. Also the maximum degree of the greedy 1.1-spanner, generated on a uniformly distributed set, of 8000 points is 14 and its weight is 11 times the weight of the minimum spanning tree of the point set. To have a rough comparison, the Θ -graph algorithm, which runs in $\mathcal{O}(n \log n)$ time, generated a graph containing 370K edges, the maximum degree was 144 and the weight was 327 times the weight of the minimum spanning tree.

For general metric spaces, there are cases where the complete graph is the only t -spanner of a point set. For example, assume V is a set of points from a metric space where the distance between any two distinct points is equal to 1. For any $1 < t < 2$, the complete graph is the only t -spanner of V . Therefore, in general metric spaces, we can not guarantee that the generated graph is sparse. The doubling dimension is defined as follows. Let λ be the smallest integer such that for each real number r , any ball of radius r can be covered by at most λ balls of radius $r/2$. The *doubling dimension* of V is defined to be $\log \lambda$.

The doubling dimension is a generalization of the Euclidean dimension, as the doubling dimension of d -dimension Euclidean space is $\Theta(d)$.

1.1 Main Results and Organization of the Paper

The main result of this paper is that for any metric space V of bounded doubling dimension, the greedy spanner of V has linear size and can be computed in $\mathcal{O}(n^2 \log n)$ time, where $n = |V|$. The organization of the remainder of this paper is as follows. In Sect. 2, we review the improved greedy algorithm of [5, 6] and give a counterexample to the conjecture that this algorithm only performs $\mathcal{O}(n)$ shortest path queries. In Sect. 3, we present an algorithm that generates the greedy spanner in near-quadratic time for some special cases. These results are generalized to metric spaces of bounded doubling dimension in Sect. 4.

2 The FG-greedy Algorithm

As mentioned above the running time of a naïve implementation of the greedy algorithm is $\mathcal{O}(mn^2 + n^3 \log n)$. Farshi and Gudmundsson [5, 6] introduced an improved version of the algorithm and showed that it improves the running time for constructing the greedy spanner considerably in practice on point sets in the plane with the Euclidean metric. We call this algorithm as the FG-greedy algorithm. The FG-greedy algorithm is the same as the original greedy algorithm except that it uses a matrix to save the length of the shortest path between every two points and updates the matrix only when it is required. Thus the matrix is not always up to date. Instead of computing a shortest path for each pair, it first checks the matrix to see if there is a t -path; if the answer is no, then it performs a shortest path query and updates the matrix which enables us to answer the distance queries correctly. They conjectured that the FG-greedy algorithm performs only $\mathcal{O}(n)$ shortest path queries, which would imply a running time of $\mathcal{O}(n^2 \log n)$.

2.1 A Counterexample

We give an example which shows that the FG-greedy algorithm may perform $\Theta(n^2)$ shortest path queries. Consider a set $S = \{p_0, p_1, \dots, p_{n-1}\}$ of n points on the real line such that $p_i = 2^i$. The algorithm sorts all pairs of points based on their distance. We also assume that for each pair (p_i, p_j) the index of the first point in the pair is less than the index of the second point, i.e. $i < j$. It is easy to see that the algorithm performs a shortest path query for each pair of points.

Note that if we change the algorithm such that when a new edge is added to the greedy spanner, it performs a shortest path query for both endpoints of the new edge and updates the weight matrix then the new algorithm performs only $\mathcal{O}(n)$ shortest path queries on the above counterexample. However, in the full version of this paper, we give an example where we still need to perform $\Omega(n \log n)$ shortest path queries.

3 A Preliminary Algorithm

Let V be a set of n points in a metric space with distance function \mathbf{d} . As mentioned before, for generating the greedy t -spanner, we start with a graph $G(V, E)$ with no edges and we go through all the pairs of points in V (in increasing order based on their distances). For each pair we check if there exists a t -path between them in G , if not we add the edge between them in G .

In the new algorithm we basically do the same thing as the FG-greedy algorithm in the sense that we use the weight matrix to answer the shortest path queries. The differences are the following:

- We process the pairs whose distance is less than L , for a real number L , exactly in the way as in the original (or FG-) greedy algorithm.
- We divide the remaining pairs into buckets such that the i th bucket contains all the pairs whose distance is between $2^{i-1}L$ and 2^iL .
- During the processing of each bucket, we keep the shortest path between the pairs in the bucket up to date. To do this we update the shortest path between all pairs in the bucket before processing the pairs in the bucket and during the process we update it when we add an edge to the graph. Note that we keep the shortest path length between pairs in a weight matrix.

Without loss of generality we assume that the diameter of the point set is one. Let $0 < L < 1$ be a real number to be specified later. We split the pairs of points into $k = \mathcal{O}(\log(\frac{1}{L}))$ buckets such that the i th bucket, i.e. E_i , contains all the pairs with distance in $[2^{(i-1)}L, 2^iL)$. Let E_0 contain all the pairs with distance less than L .

The algorithm starts with the pairs in E_0 . It process all the pairs in the set E_0 in the same manner as the original (or FG-) greedy algorithm does. Therefore, if E_0 contains $\mathcal{O}(n^\beta)$ pairs then processing it takes $\mathcal{O}(n^{\beta+1} \log n)$ time. Let G denote the current greedy spanner after processing E_0 .

Now we show how to process the remaining buckets. Assume that we have processed buckets E_1, E_2, \dots, E_{i-1} and we need to process bucket E_i . Before processing the edges in this bucket, we compute the single-source shortest path with source at each point p of V and update the weight matrix. Then we make “local” updates when we add an edge to the graph. By “local” update, we mean we update the weight matrix for all points nearby each of the endpoints of the new edge. Since the weight matrix is up to date for all pairs in the current bucket we can answer the t -path queries in constant time using the weight matrix. For details see Algorithm 3.1.

Theorem 1. *Algorithm 3.1 generates the greedy spanner of the input point set.*

Proof. To prove the correctness of the algorithm, we need to prove that the t -path queries (line 17 of Algorithm 3.1) are answered correctly.

Let (p, q) be an arbitrary pair in E_i with $\mathbf{d}(p, q) \in [L_i, 2L_i)$ which is about to be processed in the algorithm. If there is no t -path between p and q in G then the algorithm answers the t -path query correctly since the entry in the weight

matrix corresponding to the pair is at least equal to the shortest path length between p and q in G . Assume that there is a t -path between p and q in G . We have two cases:

Case 1: The shortest path between p and q in G does not pass through any edges that were added during processing of pairs in E_i . In this case we are done because before processing the pairs in E_i , we updated all-pair shortest paths and adding new edges to the graph does not change the shortest path length between p and q .

Case 2: The shortest path π in G between p and q passes through some edges of E_i . Among all edges of $E_i \cap \pi$, let (u, v) be the one that was added last by the algorithm. We may assume without loss of generality that, starting at p , the path π goes to u , then traverses (u, v) , and then continues to q . We define

$$S_{(u,v)} = \{x \in V : \mathbf{d}(x, u) < (t - 1/2)L_i \text{ or } \mathbf{d}(x, v) < (t - 1/2)L_i\}.$$

We claim (and show below) that p or q belongs to $S_{(u,v)}$. This will imply that, in the iteration in which (u, v) is added to the graph, the algorithm computes the exact shortest-path length between p and all vertices of V , or between q and all vertices of V . Therefore, at the moment when (p, q) is processed, the value of $weight(p, q)$ is equal to the shortest-path length in G between p and q .

It remains to prove the claim. Assume that neither p nor q is contained in $S_{(u,v)}$. Then $\mathbf{d}(p, u) \geq (t - \frac{1}{2})L_i$ and $\mathbf{d}(q, v) \geq (t - \frac{1}{2})L_i$. Thus, if we denote the shortest-path length between p and q by $\mathbf{d}_G(p, q)$, then

$$\mathbf{d}_G(p, q) \geq \mathbf{d}(p, u) + \mathbf{d}(u, v) + \mathbf{d}(v, q) \geq 2(t - 1/2)L_i + L_i = 2tL_i > t \cdot \mathbf{d}(p, q),$$

which contradicts the fact that there is a t -path in G between p and q . \square

Algorithm 3.1: NEW-GREEDY(V, t, L)

```

Input:  $V, t > 1$  and  $L > 0$ .
Output:  $t$ -spanner  $G = (V, E')$ .
1 foreach  $(u, v) \in V^2$  do  $weight(u, v) := \infty$ ;
2  $E :=$  sorted list of all the pairs in  $V^2$  by increasing distance; /*ties are broken arbitrarily*/
3  $E_0 :=$  all the point pairs in  $E$  with distance in  $[0, L)$ ;
4  $j := 1$ ;
5 while  $E \setminus (\bigcup_{k=1}^{j-1} E_k) \neq \emptyset$  do
6    $E_j :=$  all the point pairs in  $E \setminus (\bigcup_{k=1}^{j-1} E_k)$  with distance in  $[2^{j-1}L, 2^jL)$ ;
7    $j := j + 1$ ;
8  $l := j - 1$ ;
9  $E' := \emptyset$ ;
10  $G := (V, E')$ ;
11 Process pairs in  $E_0$  in the same way as the original (or FG-) greedy algorithm;
12 for  $i := 1, \dots, l$  do
13    $L_i := 2^{i-1}L$ ;
14   foreach  $u \in V$  do
15     Compute single-source shortest paths with source at  $u$  and update  $weight$ ;
16   foreach  $(u, v) \in E_i$ ; /* in sorted order */
17   do if  $weight(u, v) > t \cdot \mathbf{d}(u, v)$  then  $E' := E' \cup \{(u, v)\}$ ;
18   foreach  $p \in V$  do
19     if  $\mathbf{d}(p, u) < (t - \frac{1}{2})L_i$  or  $\mathbf{d}(p, v) < (t - \frac{1}{2})L_i$  then
20       Compute single-source shortest paths with source at  $p$  and update  $weight$ ;
21 return  $G(V, E')$ ;

```

Now we show that the algorithm runs in near quadratic time in some special cases. First we need to recall the well-separated pair decomposition developed by Callahan and Kosaraju [2] in d -dimensional Euclidean space.

Definition 1. Let $s > 0$ be a real number, referred to as the separation constant. We say that two point sets A and B in \mathbb{R}^d are well-separated with respect to s , if there are two disjoint balls D_A and D_B of the same radius, r , such that

- (i) D_A contains A and D_B contains B ,
- (ii) the distance between D_A and D_B is at least $s \cdot r$.

Lemma 1 ([2]). Let A and B be two finite sets of points that are well-separated with respect to s , let x and p be points of A , and let y and q be points of B . Then (i) $|xy| \leq (1 + 4/s) \cdot |pq|$ and (ii) $|px| \leq (2/s) \cdot |pq|$.

Definition 2. Let V be a set of n points in \mathbb{R}^d and let $s > 0$ be a real number. A well-separated pair decomposition (WSPD) for V with respect to s is a collection $\mathcal{W} := \{(A_1, B_1), \dots, (A_m, B_m)\}$ of pairs of non-empty subsets of V such that

1. A_i and B_i are well-separated with respect to s , for all $i = 1, \dots, m$.
2. for any two distinct points p and q of V , there is exactly one pair (A_i, B_i) in the collection, such that (i) $p \in A_i$ and $q \in B_i$ or (ii) $q \in A_i$ and $p \in B_i$.

The number of pairs, m , is called the *size* of the WSPD. Callahan and Kosaraju [2] have shown that any set $V \subseteq \mathbb{R}^d$ admits a WSPD of size $m = \mathcal{O}(s^d n)$. Har-Peled and Mendel [8] generalized the results to metric spaces with doubling dimension λ . They showed that any set of n points from a metric space with doubling dimension λ admits a WSPD with respect to $s > 1$, of size $\mathcal{O}(s^{\mathcal{O}(\lambda)} n)$. In the rest of the paper, we assume that V is a set of n points from a metric space with doubling dimension λ .

Observation 1. If $\{(A_i, B_i)\}_i$ is a WSPD of a point set V with respect to $s = \frac{4(t+1)}{t-1}$, then for each i , there is at most one greedy edge between A_i and B_i in the t -spanner generated by the greedy algorithm.

Proof. Assume that we have a pair (A, B) in the WSPD such that there exist two edges (a_1, b_1) and (a_2, b_2) in the greedy t -spanner where $a_1, a_2 \in A$ and $b_1, b_2 \in B$. Without loss of generality assume the greedy algorithm process the pair (a_1, b_1) before the pair (a_2, b_2) .

Because A and B are s -well-separated pair, by Lemma 1, we have $\mathbf{d}(a_1, a_2) \leq \frac{2}{s} \mathbf{d}(a_2, b_2) < \mathbf{d}(a_2, b_2)$. By the same argument $\mathbf{d}(b_1, b_2) < \mathbf{d}(a_1, b_1)$. Therefore, there exists a t -path between a_1 and a_2 and a t -path between b_1 and b_2 when the greedy algorithm processes the pair (a_2, b_2) . Let G' be the graph just before processing (a_2, b_2) and let P be a path in G' between a_2 and b_2 generated by concatenating a t -path between a_2 and a_1 , the edge (a_1, b_1) and a t -path between b_1 and b_2 . If $|P|$ denotes the length of the path P , then

$$\begin{aligned}
|P| &= \mathbf{d}_{G'}(a_2, a_1) + \mathbf{d}(a_1, b_1) + \mathbf{d}_{G'}(b_1, b_2) \\
&\leq t \cdot \mathbf{d}(a_2, a_1) + \mathbf{d}(a_1, b_1) + t \cdot \mathbf{d}(b_1, b_2) \\
&\leq t \cdot \frac{2}{s} \mathbf{d}(a_2, b_2) + (1 + \frac{4}{s}) \mathbf{d}(a_2, b_2) + t \cdot \frac{2}{s} \mathbf{d}(a_2, b_2) \quad (\text{by Lemma 1}) \\
&= (\frac{4t}{s} + 1 + \frac{4}{s}) \mathbf{d}(a_2, b_2) \\
&= t \cdot \mathbf{d}(a_2, b_2).
\end{aligned}$$

This means that the greedy algorithm does not add (a_2, b_2) to the spanner since there already exist a t -path between them in G' . \square

As a corollary, since there exists a linear size WSPD for any point set in a metric space with bounded doubling dimension, the size of a greedy t -spanner on such a point set is linear.

Lemma 2. *While processing the pairs in E_i , for each point p , line 20 in Algorithm 3.1 is executed $\mathcal{O}(\frac{1}{(t-1)^{\mathcal{O}(\lambda)}})$ times.*

Proof. For simplicity, we first prove the lemma in the 2-dimensional Euclidean case. Assume the distance between the pairs in E_i is in $[L, 2L)$. Algorithm 3.1 performs a single-source shortest path computation with source at p , after adding an edge (u, v) to the graph, if $\mathbf{d}(p, u) < (t - \frac{1}{2})L$ or $\mathbf{d}(p, v) < (t - \frac{1}{2})L$. So if we draw a ball C with center at p and radius $(2t + 1)L$, then all the edges which affect p lies inside the ball C . So the number of times we need to execute line 20 for p is at most the number of edges in the greedy spanner with length between L and $2L$ which lie inside C . Now we show that the number of such edges is at most $\mathcal{O}(\frac{1}{(t-1)^2})$.

To show this, assume B is a square with side length $2(2t + 1)L$ which includes C . We cover the square B with cells of side length $\ell = \frac{L}{\sqrt{2}(s+4)}$ where $s = \frac{4(t+1)}{t-1}$. The number of such cells inside B is $(2\sqrt{2}(2t + 1)(s + 4))^2 = \mathcal{O}(\frac{1}{(t-1)^2})$. Let (u, v) be a greedy edge with distance in the interval $[L, 2L)$. First we show that the grid cells which contains u and v are s -well-separated. Assume C_1 and C_2 are balls with radius $\sqrt{2}\ell$ which contain the grid cell of u and the grid cell of v , respectively— see Fig. 1. Since $\mathbf{d}(u, v) \geq L$ and the radius of the circles are $\sqrt{2}\ell$, the distance between C_1 and C_2 is at least $L - 4\sqrt{2}\ell$. Therefore

$$d(C_1, C_2) \geq L - 4\sqrt{2}\ell = L - 4\sqrt{2} \frac{L}{\sqrt{2}(s+4)} = L \left(\frac{s}{s+4} \right) = s \left(\frac{L}{s+4} \right) = s \times \sqrt{2}\ell$$

which means the cells are s -well-separated.

Therefore, by Observation 1, we have at most one greedy edge between grid cells which are well-separated. This means that the number of the greedy edges with distance in $[L, 2L)$ inside the circle C is bounded by the number of cell pairs which is $\mathcal{O}(\frac{1}{(t-1)^4})$.




Fig. 1. Illustration for the proof of Lemma 2.

For the general case, the same argument works. In this case we use the property of doubling dimension which guarantees that each ball of radius $r > 0$ can be covered by 2^λ balls of radius $r/2$. This means that the number of balls with radius $\sqrt{2}\ell$ inside the ball C centered at p is $\mathcal{O}(\frac{1}{(t-1)^{\mathcal{O}(\lambda)}})$. \square

Now we are ready to compute the time complexity of Algorithm 3.1. Clearly lines 1–11 of the algorithm take $\mathcal{O}(n^2 \log n)$ time. For line 12, if the size of E_0 is β then it takes $\mathcal{O}(\beta(m + n \log n))$ since for each pair it performs a shortest path query.

For each interval, computing all-pairs shortest path, lines 15–17, takes $\mathcal{O}(mn + n^2 \log n)$ time, and by Lemma 2, the update procedure takes at most $\mathcal{O}(\frac{1}{(t-1)^{\mathcal{O}(\lambda)}}(mn + n^2 \log n))$. Since the number of intervals is $\mathcal{O}(\log(1/L))$, processing all intervals takes $\mathcal{O}\left(\frac{\log(1/L)}{(t-1)^{\mathcal{O}(\lambda)}}(mn + n^2 \log n)\right)$ time. For a metric space with doubling dimension λ , the size of the generated graph is $\mathcal{O}(\frac{n}{(t-1)^{\mathcal{O}(\lambda)}})$, the total running time of Algorithm 3.1 is $\mathcal{O}\left(\frac{\beta n + \log(1/L)n^2 \log n}{(t-1)^{\mathcal{O}(\lambda)}}\right)$. Therefore for a point set V with the property that there exists a real number L such that $\frac{1}{L} = \mathcal{O}(n^c)$ (c is a constant) and $\beta = \mathcal{O}(n \log^2 n)$, the greedy spanner can be computed in $\mathcal{O}(n^2 \log^2 n)$ time.

Points Sets with Polynomial Aspect Ratio. If the input point set has aspect ratio less than n^c , for some constant c , then by scaling the point set such that the longest distance is 1 and setting $L = \frac{1}{n^c}$, we have no pair of points in the scaled point set with distance less than L . Therefore the running time of Algorithm 3.1 in this case is $\mathcal{O}(\frac{n^2 \log^2 n}{(t-1)^{\mathcal{O}(\lambda)}})$.

Uniformly Distributed Point Set. Assume we have a set of n points uniformly distributed inside the unit square and let $L = \frac{1}{\sqrt{n}}$. Since the points are uniformly distributed, for each point p , the expected number of points inside the ball with center at p and radius L is $L^2 n$. So the expected number of pairs with distance less than L is $L^2 n^2 / 2 = \mathcal{O}(n)$. Therefore, the expected running time of Algorithm 3.1 in this case is $\mathcal{O}(\frac{n^2 \log^2 n}{(t-1)^{\mathcal{O}(\lambda)}})$.

4 An Improved Algorithm

To generalize the results of the previous section to bounded doubling dimension, we have to overcome obstacles. First, we need to speedup processing the pairs in the first set (i.e. the set E_0). The second problem is that if we decrease the number L to bound the number of pairs in the first interval, it increases the number of buckets which causes higher time complexity. We overcome these difficulties by modifying the previous algorithm in the following way.

- we partition the $\binom{n}{2}$ pairs into a linear number of buckets,

- we maintain a data structure for each point during the algorithm. When we need to update a point, instead of doing a single-shortest path computation from scratch, we use the data structure to update just the necessary part and use it to update the weight matrix.

First we claim that for updating the shortest path lengths with source at a point p , performing a length-bounded Dijkstra’s algorithm is sufficient. More precisely, if we are working on point pairs in a bucket with distances in $[L, 2L)$ and we need to update the shortest path lengths with source at p , it is sufficient to update the distance between p and all the points q such that $\mathbf{d}_G(p, q) < 2tL$, where G is the current graph. The reason is that if $\mathbf{d}_G(p, q) \geq 2tL$ then either the pair (p, q) is outside the current bucket or there is no t -path between p and q in G .

So, in the new algorithm, we maintain a data structure for each point, which is the same as the data structure used in the Dijkstra’s single-source shortest paths algorithm. When we perform a shortest path query with bound U , we execute Dijkstra’s algorithm but stop when the key of the element on the top of the priority queue (heap) is bigger than U . We also maintain a stack storing all changes that are made to the heap in this process, so that we can undo the procedure, if required.

Note that our graph is dynamic and we add edge(s) to the graph, in increasing order of length. We use the “undo facility” to fix the part of the execution of Dijkstra’s algorithm that is affected by the insertion of an edge. For the details of the algorithm, see Algorithm 4.1.

To complete the correctness proof, we show that in the algorithm, the shortest path queries are answered correctly. To do this, we show that the length-bounded Dijkstra’s algorithm on a subgraph of the greedy spanner works exactly same as Dijkstra’s algorithm on the greedy spanner.

Lemma 3. *Let G_e be the subgraph of the greedy spanner G that contains all the edges added to the graph up to the processing of the pair $e = (p, q)$ in the greedy algorithm. Let u be an arbitrary vertex of G . As long as the key of the element on the top of the heap in Dijkstra’s algorithm with source at u is less than $\mathbf{d}(p, q)$ the algorithm works the same on G and G_e .*

Proof. Let x be a vertex of the graph G and assume the shortest path between u and x passes through at least one edge in $E \setminus E_e$. Since the length of each edge in $E \setminus E_e$ is at least $\mathbf{d}(p, q)$, the key of the point x in the heap is at least $\mathbf{d}(p, q)$. This completes the proof. \square

4.1 Running Time

Now, we show that Algorithm 4.1 runs in $\mathcal{O}(n^2 \log n)$ time. To this end, we show that for each point $p \in V$ the overall time spent is proportional to running Dijkstra’s single-source shortest paths algorithm with source p .

Algorithm 4.1 basically performs Dijkstra’s algorithm with source at each point of the graph. The only differences are

Algorithm 4.3: DIJKSTRA-BOUNDED(G, s, L, PQ, τ)

Input: Graph G , a vertex s , a real number L and a priority queue PQ .
Output: Shortest path length between s and all other vertices in G which has length at most L .

- 1 **while** *The key of the element on the top of PQ is at most L* **do**
- 2 $u := \text{pop}$ the element with minimum key from PQ ;
- 3 $\text{weight}(s, u) := \text{weight}(u, s) :=$ the key of u ;
- 4 **foreach** *node v adjacent to u in G* **do**
- 5 **if** $\text{weight}(s, u) + \text{wt}(u, v) < \text{weight}(s, v)$ **then**
- 6 Decrease the key of v in PQ to $\text{weight}(s, u) + \text{wt}(u, v)$ and add all the changes in PQ to the stack τ ; /* To be used in the undo procedure. */

Algorithm 4.4: DIJKSTRA-UNDO(τ, PQ, L)

Input: a stack τ , a priority queue PQ and a real number L .

- 1 **while** *the key of the element on the top of τ is at most L* **do**
- 2 Pop the top element from τ ;
- 3 Undo the changes on PQ based on the information in the element;

length of the new edge and redo the execution until the key of the point on the top of the heap is more than $2tL_i$. But the length of the new edge is at least L_i which means that the undo process never goes further when the key on the top of the heap is less than L_i . We say that a vertex q is in the undo area of E_i if $L_i \leq \mathbf{d}_G(p, q) \leq 2tL_i$. The claim is that each point q appears in the undo area for a constant number of sets. To prove this let q be in the undo area of E_i . This means

$$L_i \leq \mathbf{d}_G(p, q) \leq 2tL_i. \quad (1)$$

We show that q can not be in the undo area of E_j if $j > i + \log t$. Let G' be the graph when we process the pairs in E_j . Since we add edge(s) to the graph, we know that for each pair (p, q) of points $\mathbf{d}_{G'}(p, q) \leq \mathbf{d}_G(p, q)$. Since for each i , $L_{i+1} > 2L_i$ we have $L_{i+k} > 2^k L_i$, for each $k > 0$. Using Equation (1), we have

$$\mathbf{d}_{G'}(p, q) \leq \mathbf{d}_G(p, q) \leq 2tL_i < \frac{2t}{2^{\log t}} L_j \leq L_j,$$

which means q is not in the undo area of E_j .

On the other hand, by Lemma 2, we update the shortest paths with source p at most $\mathcal{O}(\frac{1}{(t-1)^{\mathcal{O}(\lambda)}})$ times in each set E_i . This means in total we need to recompute the shortest path between p and any point q at most $\mathcal{O}(\frac{\log t}{(t-1)^{\mathcal{O}(\lambda)}})$ times. Therefore, the whole process for a fixed point p takes $\mathcal{O}(\frac{\log t}{(t-1)^{\mathcal{O}(\lambda)}} n \log n)$ time. So we have the following theorem.

Theorem 2. *For each point set V of n points from a metric space with doubling dimension λ , we can compute the greedy t -spanner of V in $\mathcal{O}(\frac{\log t}{(t-1)^{\mathcal{O}(\lambda)}} n^2 \log n)$ time using $\mathcal{O}(n^2)$ space.*

5 Conclusion

In this paper we presented an algorithm which, given a set of n points from a metric space with bounded doubling dimension, computes the greedy spanner of the point set in $\mathcal{O}(n^2 \log n)$ time. In the greedy spanner, every point is connected to its nearest neighbor. Therefore, we can compute all nearest neighbors of the input point set using the greedy spanner in $\mathcal{O}(n)$ time. Har-Peled and Mendel [8] showed that the all nearest neighbor problem has an $\Omega(n^2)$ lower bound for metric spaces with bounded doubling dimension. This implies that computing the greedy t -spanner also has an $\Omega(n^2)$ lower bound. There is a logarithmic gap between the running time of the greedy algorithm presented in this paper and the lower bound. Another interesting problem is finding a $o(mn^2)$ time algorithm for constructing the greedy graph on a set of n points from a metric space, where m is the number of edges in the greedy graph. In the special case where the points are a subset of \mathbb{R}^2 , is there a $o(n^2)$ algorithm for constructing the greedy graph?

References

- [1] I. Althöfer, G. Das, D. P. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete and Computational Geometry*, 9(1):81–100, 1993.
- [2] P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k -nearest-neighbors and n -body potential fields. *Journal of the ACM*, 42:67–90, 1995.
- [3] L. P. Chew. There is a planar graph almost as good as the complete graph. In *SCG'86: Proceedings of the 2nd Annual ACM Symposium on Computational Geometry*, pages 169–177, 1986.
- [4] G. Das and G. Narasimhan. A fast algorithm for constructing sparse Euclidean spanners. *Int. J. of Computational Geometry & Applications*, 7:297–315, 1997.
- [5] M. Farshi and J. Gudmundsson. Experimental study of geometric t -spanners. In *ESA'05*, volume 3669 of *LNCS*, pages 556–567. Springer-Verlag, 2005.
- [6] M. Farshi and J. Gudmundsson. Experimental study of geometric t -spanners: A running time comparison. In *WEA'07*, volume 4525 of *LNCS*, pages 270–284. Springer-Verlag, 2007.
- [7] S. Har-Peled. A simple proof?, 2006. <http://valis.cs.uiuc.edu/blog/?p=362>.
- [8] S. Har-Peled and M. Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM Journal on Computing*, 35(5):1148–1184, 2006.
- [9] J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete and Computational Geometry*, 7:13–28, 1992.
- [10] G. Narasimhan and M. Smid. *Geometric spanner networks*. Cambridge University Press, 2007.
- [11] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, Philadelphia, PA, 2000.
- [12] D. Peleg and A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13:99–116, 1989.
- [13] D. Russel and L. J. Guibas. Exploring protein folding trajectories using geometric spanners. *Pacific Symposium on Biocomputing*, pages 40–51, 2005.