

The Weak Gap Property in Metric Spaces of Bounded Doubling Dimension*

Michiel Smid

School of Computer Science, Carleton University, Ottawa, Canada

Abstract. We introduce the weak gap property for directed graphs whose vertex set S is a metric space of size n . We prove that, if the doubling dimension of S is a constant, any directed graph satisfying the weak gap property has $O(n)$ edges and total weight $O(\log n) \cdot wt(MST(S))$, where $wt(MST(S))$ denotes the weight of a minimum spanning tree of S . We show that 2-optimal TSP tours and greedy spanners satisfy the weak gap property.

1 Introduction

Consider a directed graph $G = (S, E)$, where S is a set of n points in \mathbb{R}^d , and each edge (p, q) in E has a weight (or length) which is equal to the Euclidean distance $|pq|$ between the points p and q . We consider the problem of estimating the weight $wt(E)$ of the edge set E , which is defined to be the sum of the weights of the edges in E . Clearly, in order to obtain a non-trivial estimate, we need to make some assumptions about the edge set E .

Using geometric properties of \mathbb{R}^d , Chandra *et al.* [5] showed that the directed edge set of the greedy spanner algorithm (to be introduced later) can be partitioned into $O(1)$ subsets, such that each subset satisfies the *gap property*: For any two distinct edges (p, q) and (r, s) that are in the same subset, the distance $|pr|$ is at least proportional to the weight of the shorter of (p, q) and (r, s) . They proved that the gap property implies that the weight of the subset is $O(\log n)$ times the weight $wt(MST(S))$ of a minimum spanning tree of the point set S . As a result, the weight of the greedy spanner is $O(\log n) \cdot wt(MST(S))$. (A much more complicated analysis shows that, in fact, the weight of the greedy spanner is $O(wt(MST(S)))$.)

Later, Chandra *et al.* [6] showed that any tour T that is computed by the 2-opt heuristic (to be introduced later) for the traveling salesperson problem can be analyzed by the same approach: By again using geometric properties of \mathbb{R}^d , the directed edge set of T can be partitioned into $O(1)$ subsets, each of which satisfies the gap property. Thus, by the result in [5], the weight of T is $O(\log n) \cdot wt(MST(S))$. Since $wt(MST(S))$ is less than the minimum weight $wt(TSP(S))$ of any tour of S , it follows that $wt(T) = O(\log n) \cdot wt(TSP(S))$. Chandra *et al.* also showed that, for the case when $d \geq 2$, there exists a 2-opt tour having weight $\Omega(\log n / \log \log n) \cdot wt(TSP(S))$.

* This work was supported by NSERC.

Both results mentioned above use the fact that the input points are in Euclidean space \mathbb{R}^d , for some constant $d \geq 1$. This leads to the natural question whether these results hold in an arbitrary *metric space*. Recall that a set S , together with a distance function $|pq|$ for any two points p and q in S , is called a metric space, if for all p, q , and r in S ,

1. $|pp| = 0$,
2. $|pq| > 0$ if $p \neq q$,
3. $|pq| = |qp|$, and
4. $|pq| \leq |pr| + |rq|$.

The fourth property is called the *triangle inequality*.

The proof of Chandra *et al.* [5] of the fact that any directed edge set satisfying the gap property has weight $O(\log n) \cdot wt(MST(S))$ holds in any metric space. Narasimhan and Smid [13, Section 6.2] showed that this upper bound is tight, even in the one-dimensional Euclidean metric. On the other hand, in the metric space in which $|pq| = 1$ for all distinct points p and q , the weight of the greedy spanner is $\Theta(n) \cdot wt(MST(S))$. Chandra *et al.* [6] proved that, in any metric space, any tour that is computed by the 2-opt heuristic has weight $O(\sqrt{n}) \cdot wt(TSP(S))$; they also showed that this upper bound is tight. Thus, the weights of the outputs of the greedy spanner algorithm and the 2-opt heuristic behave very differently in Euclidean space \mathbb{R}^d than they do in a general metric space.

The analyses in [5, 6] for the Euclidean metric use the notion of angles. In particular, they use the fact that any set of vectors in \mathbb{R}^d , in which any two elements make an angle of at least θ , contains $O(1/\theta^{d-1})$ elements. This is basically a *packing argument*: The number of “large” objects that can be packed in another slightly “larger” object is bounded from above by a constant (which depends on the dimension d). In Euclidean space, the validity of such an argument follows from the fact that a “large” object has a “large” volume. In a general metric space, however, a packing argument cannot be applied.

In this paper, we consider the weights of the greedy spanner and 2-opt tours in metric spaces in which a packing argument is valid. Such metric spaces are called metric spaces of bounded doubling dimension. We will prove that, in such spaces, the weights of the greedy spanner and 2-opt tours are $O(\log n) \cdot wt(MST(S))$. We obtain these results by generalizing the gap property to the so-called weak gap property. We then show that any edge set satisfying the weak gap property has weight $O(\log n) \cdot wt(MST(S))$. Since both the greedy spanner and 2-opt tours satisfy the weak gap property, we obtain the same upper bounds on their weights.

Thus, the contributions of this paper are twofold: First, by introducing the weak gap property, we obtain alternative proofs of known results for the Euclidean metric. Second, our analysis shows that these results in fact hold for any metric space whose doubling dimension is a constant.

2 The Doubling Dimension of Metric Spaces

Let S be a finite metric space. For any two points p and q in S , we denote their distance by $|pq|$. If p is a point in S and $R > 0$ is a real number, then the *ball* with center p and radius R is defined to be the set $\{q \in S : |pq| \leq R\}$.

We now define the notion of doubling dimension, which is due to Assouad [2]; see also Heinonen [11]:

Definition 1. *Let S be a finite metric space and let λ be the smallest integer such that the following is true: For each real number $R > 0$, every ball in S of radius R can be covered by at most λ balls of radius $R/2$. The doubling dimension of the metric space S is defined to be $\log \lambda$.*

It is not difficult to show that the doubling dimension of any finite set of points in the Euclidean metric space \mathbb{R}^d is $\Theta(d)$.

2.1 Non-Euclidean Spaces of Doubling Dimension 1

In this section, we give an example of a family of metric spaces having doubling dimension 1. As we will see, this family contains metric spaces whose properties are very different from Euclidean space \mathbb{R}^d for any constant d .

Let $S = \{p_1, p_2, \dots, p_n\}$ and let $0 < \epsilon < 1$ be a real number. For each i and j with $1 \leq i \leq j \leq n$, we define

$$|p_i p_j| = |p_j p_i| = \begin{cases} 0 & \text{if } i = j, \\ 4^j \text{ or } 4^j + \epsilon & \text{if } i < j. \end{cases}$$

This family of metric spaces occurs in Har-Peled and Mendel [10] (they use powers of 2 instead of powers of 4). For any three pairwise distinct indices i , j , and k with $i < j$, we have

$$|p_i p_j| \leq 4^j + \epsilon \leq 4^{\max(i,k)} + 4^{\max(k,j)} \leq |p_i p_k| + |p_k p_j|,$$

implying that this distance function satisfies the triangle inequality. It follows that S is a metric space.

Lemma 1. *The doubling dimension of the metric space S is equal to one.*

Proof. Let $R > 0$ be a real number and let B be a ball with radius R . We have to show that B can be covered by at most two balls of radius $R/2$.

First assume that $R < 4\epsilon$. Then $R < 4$. Since the minimum distance between any two distinct points of S is at least 16 , the ball B contains only its center. Therefore, B can be covered by one ball of radius $R/2$.

Now assume that $R \geq 4\epsilon$. Let j be the largest index such that $p_j \in B$. If $j = 1$, then B contains only one point, so that this ball can be covered by one ball of radius $R/2$. Assume that $j \geq 2$. We define B_1 to be the ball with center p_j and radius $R/2$, and define B_2 to be the ball with center p_{j-1} and radius $R/2$. We claim that $B \subseteq B_1 \cup B_2$.

We may assume that B contains more than one point, because otherwise, $B = B_1$. Let p_k be the center of B . Observe that $k \leq j$. First assume that $k < j$. Since $p_j \in B$, we have $|p_k p_j| \leq R$. On the other hand, we have $|p_k p_j| \geq 4^j$. Thus, $4^j \leq R$. If $k = j$, then let ℓ be any index less than k for which $p_\ell \in B$. In this case, we have $|p_k p_\ell| \leq R$ and $|p_k p_\ell| \geq 4^k = 4^j$. Thus, also in this case, we have $4^j \leq R$.

We are now ready to complete the proof of the claim that $B \subseteq B_1 \cup B_2$. Let p_i be any point in B . If $i \in \{j-1, j\}$, then obviously $p_i \in B_1 \cup B_2$. If $i \leq j-2$, then

$$|p_{j-1} p_i| \leq 4^{j-1} + \epsilon \leq R/4 + \epsilon \leq R/2$$

and, therefore, $p_i \in B_2$. \square

As Har-Peled and Mendel [10] show, this family of metric spaces can be used to show that the time complexity for solving the all-nearest-neighbors problem is $\Theta(n^2)$. Indeed, consider a metric space in the family such that, for each j with $2 \leq j \leq n$, there is exactly one index i_j with $i_j < j$ such that $|p_{i_j} p_j| = 4^j$, whereas for all indices i' with $i' < j$ and $i' \neq i_j$, we have $|p_{i'} p_j| = 4^j + \epsilon$. Then, the (unique) nearest neighbor of p_j is the point p_{i_j} . Thus, any algorithm that solves the all-nearest-neighbors problem must find all indices i_j . By an adversarial argument, it is easy to show that this takes $\Omega(n^2)$ time in the worst case.

Recall that the all-nearest-neighbors problem in Euclidean space \mathbb{R}^d , for a constant dimension d , can be solved in $O(n \log n)$ time; see Vaidya [16]. Thus, even though the metric space defined above has doubling dimension one, its behavior with respect to the all-nearest-neighbors problem is very different from the Euclidean metric in \mathbb{R}^d .

Now consider the case when

$$|p_i p_j| = |p_j p_i| = \begin{cases} 0 & \text{if } i = j, \\ 4^j & \text{if } i = 1 \text{ and } j > 1, \\ 4^j + \epsilon & \text{otherwise.} \end{cases}$$

For each j with $2 \leq j \leq n$, p_1 is the nearest neighbor in S of the point p_j . Therefore, the all-nearest-neighbors graph of S is the star-graph consisting of all edges $\{p_1, p_j\}$, $2 \leq j \leq n$. In fact, the minimum spanning tree is also equal to this star-graph. Thus, in this metric space, both the all-nearest-neighbors graph and the minimum spanning tree have maximum degree $n-1$. It is well known that in Euclidean space \mathbb{R}^d , the maximum degree of both these graphs is bounded by a constant that depends only on d .

Since the minimum spanning tree contains the all-nearest-neighbors graph, the time complexity for computing the minimum spanning tree in a metric space of constant doubling dimension is $\Theta(n^2)$. On the other hand, in Euclidean space \mathbb{R}^d , the minimum spanning tree can be computed in $O(n \log n)$ time if $d = 2$ (see Preparata and Shamos [14, Section 6.1]) and $o(n^2)$ time if $d > 2$ (see Yao [17]).

2.2 The Packing Lemma

We now show that a packing argument can be applied in any metric space of constant doubling dimension. The following lemma states that a ball of radius R

cannot contain many points whose pairwise distances are at least proportional to R .

Lemma 2 (Packing Lemma). *Let S be a finite metric space with doubling dimension d , let $R > 0$ and $\alpha > 0$ be real numbers, let B be a ball in S of radius R , and let X be a subset of S , such that*

1. $X \subseteq B$ and
2. the distance between any two distinct points of X is at least αR .

Then, the number of points in the set X is at most

$$2^{d \cdot \max(0, 2 + \lceil \log(1/\alpha) \rceil)}.$$

Proof. If $\alpha > 2$, then the lemma holds because X contains at most one point. Assume that $0 < \alpha \leq 2$. Let $m = 2 + \lceil \log(1/\alpha) \rceil$. Then $m \geq 1$ and $m > 1 + \log(1/\alpha)$, implying that $R/2^{m-1} < \alpha R$. By repeatedly applying the definition of doubling dimension, the ball B can be covered by 2^{md} balls B_i ($1 \leq i \leq 2^{md}$) of radius $R/2^m$. Any two points in the same ball B_i have distance at most $2R/2^m < \alpha R$. Therefore, each ball B_i can contain only one point of X . Thus, X contains at most 2^{md} points. \square

3 The Weak Gap Property

We mentioned the gap property in Section 1. Here, we give a formal definition of this property, as well as a weaker version of it. The weak gap property will be the main focus of the rest of this paper.

Definition 2. *Let S be a finite metric space and let E be a set of directed edges whose endpoints are in S .*

1. For a real constant $w > 0$, we say that E satisfies the w -gap property, if for any two distinct edges (p, q) and (r, s) in E ,

$$|pr| \geq w \cdot \min(|pq|, |rs|).$$

2. For a real constant $w > 0$, we say that E satisfies the weak w -gap property, if for any two distinct edges (p, q) and (r, s) in E ,

$$|pr| \geq w \cdot \min(|pq|, |rs|)$$

or

$$|qs| \geq w \cdot \min(|pq|, |rs|).$$

Chandra *et al.* [5] proved that, in any metric space S and for any constant $w > 0$, any set of directed edges satisfying the w -gap property has weight $O(\log n) \cdot wt(TSP(S))$, where n is the number of points in S . The example in Narasimhan and Smid [13, Section 6.2] shows that this upper bound is tight, even in one-dimensional Euclidean space.

In general, there is no non-trivial upper bound on the weight of an edge set satisfying the weak gap property. Indeed, consider the metric space S on n points in which $|pq| = 1$ for all $p \neq q$, and take for E the edge set of the complete graph, by giving each edge an arbitrary direction. Then E satisfies the weak 1-gap property, $wt(E) = \Theta(n) \cdot wt(TSP(S))$, and E contains $\Theta(n^2)$ edges. The main result of this paper is a proof of the claims that, if the doubling dimension of the metric space is a constant, (i) the upper bound of $O(\log n) \cdot wt(TSP(S))$ does hold, and (ii) E contains $O(n)$ edges.

Theorem 1 (Gap Theorem). *Let S be a metric space with n points, let the doubling dimension of S be a constant, let $w > 0$ be a real constant, and let E be a set of directed edges whose endpoints are in S . If E satisfies the weak w -gap property, then*

1. *the total weight $wt(E)$ of all edges in E satisfies*

$$wt(E) = O(\log n) \cdot wt(TSP(S)),$$

2. *E contains $O(n)$ edges.*

Recall that $wt(TSP(S))$ and $wt(MST(S))$ differ by a factor of at most 2. Therefore, the Gap Theorem is still valid if we replace $TSP(S)$ by $MST(S)$.

Before we turn to the proof of the Gap Theorem, we consider two algorithms whose outputs satisfy the weak gap property.

3.1 The 2-Opt Heuristic for the Traveling Salesperson Problem

Let S be an arbitrary finite metric space. The *2-opt heuristic* is a well-known approach for heuristically solving the traveling salesperson problem; it was introduced by Lin [12]. The algorithm starts with an arbitrary initial tour along the points of S . Then it repeatedly tries to improve the current tour by making small local changes.

Let T be the current (directed) tour, and assume that T contains two distinct edges (p, q) and (r, s) such that

$$|pr| + |qs| < |pq| + |rs|.$$

Then the algorithm replaces (p, q) and (r, s) by the edges (p, r) and (q, s) , and reverses the direction of the edges on the path from q to r . These replacements result in a shorter (directed) tour. The algorithm continues making these replacements until the current tour T has the property that

$$|pq| + |rs| \leq |pr| + |qs| \tag{1}$$

for any two distinct edges (p, q) and (r, s) of T . A tour having this property is called a *2-optimal tour*.

If S is an arbitrary metric space, then the weight of any 2-optimal tour is $O(\sqrt{n}) \cdot wt(TSP(S))$; see Chandra *et al.* [6]. These authors also showed that

this upper bound is tight. On the other hand, Chandra *et al.* [6] used the gap property to show that, in the Euclidean metric \mathbb{R}^d , any 2-optimal tour has weight $O(\log n) \cdot wt(TSP(S))$.

We will use the Gap Theorem to prove that, if the doubling dimension of the metric space is constant, any 2-optimal tour has weight $O(\log n) \cdot wt(TSP(S))$.

Lemma 3. *Let T be a (directed) 2-optimal tour of the points in S . Then the edge set of T satisfies the weak 1-gap property.*

Proof. Let (p, q) and (r, s) be two distinct edges of T . We may assume without loss of generality that $|pq| \leq |rs|$. Thus, we have to show that $|pr| \geq |pq|$ or $|qs| \geq |pq|$. If $|pr| \geq |pq|$, then we are done. Assume that $|pr| < |pq|$. By (1), we have $|pq| + |rs| \leq |pr| + |qs|$. Combining this with our assumption that $|pr| < |pq|$, we obtain

$$|pq| + |rs| \leq |pr| + |qs| \leq |pq| + |qs|,$$

which implies that $|qs| \geq |rs| \geq |pq|$. □

Lemma 3 and the Gap Theorem imply the following result.

Theorem 2. *Let S be a metric space with n points and constant doubling dimension. The 2-opt heuristic computes a tour along the points of S , whose weight is $O(\log n) \cdot wt(TSP(S))$.*

3.2 The Greedy Spanner Algorithm

Given a metric space S consisting of n points and a real constant $t > 1$, an undirected graph $G = (S, E)$ is called a t -spanner for S , if the following is true: For any two points p and q in S , there exists a path in G between p and q whose weight is at most $t|pq|$. Any such path is called a t -spanner path between p and q .

Althöfer *et al.* [1] introduced the following simple greedy algorithm for computing such a spanner (according to them, this algorithm was discovered independently by Bern in 1989):

Algorithm GREEDYSPANNER(S, t):
 sort the $\binom{n}{2}$ pairs of distinct points in non-decreasing order of their distances and store them in a list L ;
 $E = \emptyset$;
 $G = (S, E)$;
for each $\{p, q\} \in L$ (* in sorted order *)
do $\delta =$ weight of a shortest path in G between p and q ;
 if $\delta > t|pq|$
 then $E = E \cup \{\{p, q\}\}$;
 $G = (S, E)$
 endif
endfor;
 output the graph G

It is obvious that this algorithm computes a t -spanner for S . If $t < 2$ and the metric space S has the property that $|pq| = 1$ for all distinct points p and q , then the greedy spanner contains $\binom{n}{2}$ edges and its total edge weight is $\Theta(n) \cdot wt(MST(S))$. On the other hand, Soares [15] proved that, in the Euclidean metric \mathbb{R}^d , the greedy spanner has bounded degree and, thus, contains only $O(n)$ edges. Again in \mathbb{R}^d , Chandra *et al.* [5] used the gap property to show that the weight of the greedy spanner is $O(\log n) \cdot wt(MST(S))$. In fact, the latter bound was improved to $O(wt(MST(S)))$ by Das *et al.* [7, 8]. (Chapter 14 in [13] contains a complete proof of this claim.)

It is not difficult to show that the greedy spanner contains the all-nearest-neighbors graph. Therefore, by the results in Section 2.1, the maximum degree of the greedy spanner can be as large as $n - 1$, even when the doubling dimension is equal to one.

We will use the Gap Theorem to prove that, if the doubling dimension of the metric space is constant, the greedy spanner has weight $O(\log n) \cdot wt(MST(S))$ and contains $O(n)$ edges.

Consider the t -spanner G that is computed by the greedy algorithm. Let G^* be the directed graph obtained by giving each edge of G an arbitrary direction. Thus, each edge $\{p, q\}$ of G appears in G^* either as (p, q) or as (q, p) .

Lemma 4. *Let w be a real number with $0 < w < 1 - 1/t$ and let $w' = \min(w, 1 - w - 1/t)$. Then the set of directed edges of G^* satisfies the weak w' -gap property.*

Proof. Let (p, q) and (r, s) be two distinct edges of G^* . We may assume without loss of generality that algorithm GREEDYSPANNER(S, t) considers the pair $\{p, q\}$ before $\{r, s\}$. Thus, we have $|pq| \leq |rs|$. The lemma will follow from the claim that $|pr| \geq w|pq|$ or $|qs| \geq (1 - w - 1/t)|rs|$.

Assume that $|pr| < w|pq|$ and $|qs| < (1 - w - 1/t)|rs|$. Then both $|pr|$ and $|qs|$ are less than $|rs|$. Consider the iteration in which the algorithm adds the edge $\{r, s\}$ to the graph G . Assuming that $p \neq r$ and $q \neq s$, the algorithm has already considered the pairs $\{p, r\}$, $\{p, q\}$, and $\{q, s\}$. Thus, at the start of this iteration, the graph G contains (i) a t -spanner path between p and r , (ii) the edge $\{p, q\}$, and (iii) a t -spanner path between q and s . Obviously, (i) also holds if $p = r$ and (iii) also holds if $q = s$. Since

$$\begin{aligned} t|pr| + |pq| + t|qs| &\leq tw|pq| + |pq| + t(1 - w - 1/t)|rs| \\ &\leq (tw + 1 + t(1 - w - 1/t))|rs| \\ &= t|rs|, \end{aligned}$$

the graph G contains a t -spanner path between r and s and, therefore, the algorithm does not add the edge $\{r, s\}$ to G . This is a contradiction. \square

If we choose $w = \frac{1}{2}(1 - 1/t)$ in Lemma 4, then the Gap Theorem implies the following result.

Theorem 3. *Let S be a metric space with n points and constant doubling dimension, and let $t > 1$ be a real constant. Algorithm GREEDYSPANNER(S, t) computes a t -spanner for S having $O(n)$ edges and total edge weight $O(\log n) \cdot wt(MST(S))$.*

4 Proof of the Gap Theorem

In this section, we present a proof of the Gap Theorem. This proof will consist of the following four steps:

1. We start by showing that, in any metric space, if a directed edge set E with vertex set S satisfies the gap property and all edges in E have approximately the same weight, then the total weight of E is $O(\text{wt}(TSP(S)))$.
2. We then show that, again in any metric space, the total weight of any set E of edges with vertex set S , such that all edges in E are “short”, is $O(\text{wt}(TSP(S)))$.
3. Next, we show that, in any metric space of constant doubling dimension, and for any directed edge set E with vertex set S , such that E satisfies the weak gap property and all edges in E are “long”, the following holds: First, the edge set E can be partitioned into $O(\log n)$ subsets, such that within each subset, edges have approximately the same weight. Then, we show how to further partition each subset into $O(1)$ subsets, each one satisfying the gap property. The analysis in the first step then shows that the total weight of E is $O(\log n) \cdot \text{wt}(TSP(S))$.
4. In the final step, we use the well-separated pair decomposition to show that, in any metric space of constant doubling dimension, any directed edge set satisfying the weak gap property contains $O(n)$ edges.

4.1 Edges of Similar Weights Satisfying the Gap Property

We start by considering directed edges having approximately the same weights and that satisfy the gap property (as opposed to the weak gap property). The proof of the following lemma is a simple modification of a proof technique introduced in Chandra *et al.* [5].

Lemma 5. *Let S be a metric space, let $w > 0$ be a real constant, and let E be a set of directed edges whose endpoints are in S and that satisfy the w -gap property. Assume that $1/2 \leq |pq|/|rs| \leq 2$ for any two edges (p, q) and (r, s) in E . Then, the total weight $\text{wt}(E)$ of all edges in E satisfies*

$$\text{wt}(E) \leq \frac{2}{w} \cdot \text{wt}(TSP(S)).$$

Proof. For any (directed) edge (p, q) of E , we call p the *source* of this edge. First observe that, by the definition of the w -gap property, each point of S can be the source of only one edge in E . Consider the traveling salesperson tour of S . By walking along this tour (starting at some arbitrary point), we visit the sources of the edges in E in some order. We number the edges of E as e_0, e_1, \dots, e_{m-1} , as given by this order. For each i with $0 \leq i < m$, we write the edge e_i as $e_i = (p_i, q_i)$ and define T_i to be the portion of the tour that starts at p_i and ends at p_{i+1} (where indices are read modulo m). By the triangle inequality, we have

$$|p_i p_{i+1}| \leq \text{wt}(T_i),$$

whereas, by the w -gap property and our assumption that edges in E differ in weight by a factor of at most two,

$$|p_i p_{i+1}| \geq w \cdot \min(|p_i q_i|, |p_{i+1} q_{i+1}|) \geq (w/2)|p_i q_i|.$$

Thus, we have

$$|p_i q_i| \leq (2/w) \cdot wt(T_i),$$

which implies that

$$wt(E) = \sum_{i=0}^{m-1} |p_i q_i| \leq (2/w) \sum_{i=0}^{m-1} wt(T_i) = (2/w) \cdot wt(TSP(S)).$$

□

4.2 The Weight of Short Edges

Lemma 6. *Let S be a metric space with n points, let E be a set of directed edges whose endpoints are in S , let D be the weight of a longest edge in E , and let E' be the subset of E consisting of all edges having weight at most D/n^2 . Then $wt(E') \leq wt(TSP(S))$.*

Proof. The lemma follows from the observations that E' contains at most $\binom{n}{2} \leq n^2$ edges and $wt(TSP(S)) \geq D$. □

4.3 Long Edges Satisfying the Weak Gap Property

Let S be a metric space with n points and constant doubling dimension d , let $w > 0$ be a real constant, and let E be a set of directed edges whose endpoints are in S and that satisfy the weak w -gap property.

Let D be the weight of a longest edge in E , let E' be the subset of E consisting of all edges having weight at most D/n^2 , and let $E'' = E \setminus E'$. Thus, E'' consists of all edges in E having weight more than D/n^2 . For each j with $0 \leq j \leq \lfloor 2 \log n \rfloor$, we define

$$E_j = \{e \in E'' : D/2^{j+1} < wt(e) \leq D/2^j\}.$$

Thus, the edge sets E_j form a partition of E'' .

We fix an index j with $0 \leq j \leq \lfloor 2 \log n \rfloor$ and analyze the total weight of all edges in E_j . Even though edges in E_j differ in weight by a factor of at most two, Lemma 5 cannot be applied to E_j , because we only know that E_j satisfies the weak gap property. Our approach will be to further partition the set E_j into $O(1)$ subsets, each of which satisfies the w -gap property. Then, Lemma 5 can be applied to each of these subsets.

Let $L = D/2^{j+1}$, so that

$$E_j = \{e \in E'' : L < wt(e) \leq 2L\}.$$

We define an undirected graph H with vertex set E_j , in which any two distinct vertices (p, q) and (r, s) are connected by an edge if and only if

$$|pr| < w \cdot \min(|pq|, |rs|).$$

Lemma 7. *The maximum degree of the graph H is $O(1)$.*

Proof. Consider any vertex (p, q) of H . We have to prove an upper bound on the number of elements (r, s) in E_j for which $(p, q) \neq (r, s)$ and $|pr| < w \cdot \min(|pq|, |rs|)$.

Let B be the ball with center p and radius $2wL$. For any edge (r, s) in E_j with $|pr| < w \cdot \min(|pq|, |rs|)$, we have $|pr| \leq 2wL$ and, thus, $r \in B$.

By applying the definition of doubling dimension twice, the ball B can be covered by 2^{2d} balls B_i ($1 \leq i \leq 2^{2d}$) of radius $wL/2$. For each i with $1 \leq i \leq 2^{2d}$, we define

$$E_j^i = \{(r, s) \in E_j : (r, s) \neq (p, q), |pr| < w \cdot \min(|pq|, |rs|), r \in B_i\}.$$

Then

$$\{(r, s) \in E_j : (p, q) \neq (r, s), |pr| < w \cdot \min(|pq|, |rs|)\} = \bigcup_{i=1}^{2^{2d}} E_j^i. \quad (2)$$

Since the degree of (p, q) in H is equal to the size of the set on the left-hand side in (2), we need an upper bound on $\sum_{i=1}^{2^{2d}} |E_j^i|$.

Consider a fixed value of i with $1 \leq i \leq 2^{2d}$. If (r, s) is an edge in E_j^i , then

$$|ps| \leq |pr| + |rs| \leq w \cdot \min(|pq|, |rs|) + |rs| \leq 2wL + 2L = 2(1+w)L.$$

If (r, s) and (r', s') are two distinct edges in E_j^i , then r and r' are both in B_i and thus

$$|rr'| \leq wL < w \cdot \min(|rs|, |r's'|).$$

Therefore, by the weak w -gap property, we have

$$|ss'| \geq w \cdot \min(|rs|, |r's'|) \geq wL.$$

In particular, $s \neq s'$. Thus, if we define X to be the set of the sinks s of all edges (r, s) in E_j^i , then (i) X contains the same number of elements as E_j^i , (ii) all points of X are contained in the ball with center p and radius $2(1+w)L$, and (iii) the distance between any two distinct points of X is at least wL . By applying Lemma 2 with $\alpha = w/(2(1+w))$, it follows that

$$|E_j^i| \leq 2^{d \cdot \max(0, 2 + \log(1/\alpha))} = 2^{d \cdot \max(0, 3 + \log(1+1/w))}.$$

Thus, we have shown that the degree of the vertex (p, q) in H is at most

$$\sum_{i=1}^{2^{2d}} |E_j^i| \leq 2^{d(2 + \max(0, 3 + \log(1+1/w)))}.$$

Since d and w are constants, the proof is complete. \square

Let m be the maximum degree of any vertex in H . We color the vertices of H (i.e., the elements of E_j) using $m + 1$ colors, such that any two adjacent vertices have distinct colors. For each k with $0 \leq k \leq m$, let

$$E_{jk} = \{e \in E_j : e \text{ has color } k\}.$$

The subsets E_{jk} , $0 \leq k \leq m$, partition the set E_j .

Lemma 8. *For each k with $0 \leq k \leq m$, the set E_{jk} satisfies the w -gap property.*

Proof. Let (p, q) and (r, s) be two distinct edges in E_{jk} . Since these two edges have the same color, they are not connected by an edge in H . Then, the definition of H implies that $|pr| \geq w \cdot \min(|pq|, |rs|)$. \square

We are now able to complete the proof of the first claim in the Gap Theorem: We started by partitioning the edge set E into E' and E'' . Lemma 6 gives an upper bound on $wt(E')$. Then, we partitioned E'' into $O(\log n)$ subsets E_j , and further partitioned each E_j into $m + 1 = O(1)$ subsets E_{jk} . By Lemmas 5 and 8, the total weight of all edges in each subset E_{jk} is $O(wt(TSP(S)))$.

4.4 The Number of Edges

In this section, we prove the second claim in the Gap Theorem. Our proof will use the well-separated pair decomposition of Callahan and Kosaraju [4].

Let S be a metric space with n points. For any two non-empty subsets A and B of S , we define their distance $|AB|$ as

$$|AB| = \min\{|pq| : p \in A, q \in B\}$$

and the diameter $diam(A)$ of A as

$$diam(A) = \max\{|pq| : p \in A, q \in A\}.$$

For a real number $c > 1$, called the *separation ratio*, we say that A and B are *well-separated*, if

$$|AB| \geq c \cdot \max(diam(A), diam(B)).$$

Thus, if c is large, then (i) all distances between points in A and points in B are approximately equal and (ii) distances within A (or B) are much smaller than distances between points in A and points in B .

Definition 3. *A well-separated pair decomposition for S is a sequence*

$$\{A_1, B_1\}, \{A_2, B_2\}, \dots, \{A_m, B_m\}$$

of pairs of non-empty subsets of S , for some integer m , such that

1. *for each i with $1 \leq i \leq m$, A_i and B_i are well-separated, and*
2. *for any two distinct points p and q of S , there is exactly one index i such that*

- (a) $p \in A_i$ and $q \in B_i$, or
- (b) $p \in B_i$ and $q \in A_i$.

The integer m is called the size of the well-separated pair decomposition.

Callahan and Kosaraju [4] showed that, in the Euclidean metric \mathbb{R}^d , a well-separated pair decomposition of size $m = O(n)$ exists, and can in fact be computed in $O(n \log n)$ time. Har-Peled and Mendel [10] generalized this result to metric spaces of constant doubling dimension:

Theorem 4. *Let S be a metric space with n points and constant doubling dimension, and let $c > 1$ be a real constant. There exists a randomized algorithm that constructs, in $O(n \log n)$ expected time, a well-separated pair decomposition for S , consisting of $O(n)$ pairs.*

The following lemma completes the proof of the Gap Theorem:

Lemma 9. *Let S be a metric space with n points and constant doubling dimension, let $w > 0$ be a real constant, and let E be a set of directed edges whose endpoints are in S and that satisfy the weak w -gap property. Then, E contains $O(n)$ edges.*

Proof. Choose the separation ratio c to be larger than $1/w$. Consider a well-separated pair decomposition $\{A_i, B_i\}$, $1 \leq i \leq m$, where $m = O(n)$.

We will prove below that for each i with $1 \leq i \leq m$, the set E contains (i) at most one edge (p, q) with $p \in A_i$ and $q \in B_i$, and (ii) at most one edge (p, q) with $p \in B_i$ and $q \in A_i$. Therefore, E contains at most $2m = O(n)$ edges.

Assume that (i) is not true. Then, E contains two distinct edges (p, q) and (r, s) with $p, r \in A_i$ and $q, s \in B_i$. We may assume without loss of generality that $|pq| \leq |rs|$. Since A_i and B_i are well-separated, we have

$$|pr| \leq \text{diam}(A_i) \leq \frac{1}{c} \cdot |A_i B_i| \leq \frac{1}{c} \cdot |pq| < w|pq|$$

and, by a symmetric argument,

$$|qs| < w|pq|.$$

This contradicts the fact that (p, q) and (r, s) satisfy the weak w -gap property. \square

5 Final Remarks

5.1 Open Problems

We have introduced the weak gap property as an alternative to the gap property of Chandra *et al.* [5]. We have shown that, in any metric space whose doubling dimension is constant, any set of edges satisfying the weak gap property has $O(n)$ elements and total weight $O(\log n) \cdot \text{wt}(TSP(S))$ or, equivalently, $O(\log n) \cdot$

$wt(MST(S))$. The example in Narasimhan and Smid [13, Section 6.2] shows that this upper bound is tight, even in the one-dimensional Euclidean metric.

We have shown that both 2-optimal tours and greedy spanners satisfy the weak gap property. Thus, in case the doubling dimension is constant, their total weight is within a factor of $O(\log n)$ of the minimum possible weight. These results lead to the following open problems:

Problem 1. Chandra *et al.* [6] showed that, in the Euclidean plane \mathbb{R}^2 , there exists a 2-optimal tour whose length is $\Omega(\log n / \log \log n) \cdot wt(TSP(S))$.

- Is it true that, in the Euclidean metric in \mathbb{R}^d , the weight of any 2-optimal tour is $O(\log n / \log \log n) \cdot wt(TSP(S))$?
- Does there exist a metric space S of constant doubling dimension, such that S contains a 2-optimal tour whose length is $\Omega(\log n) \cdot wt(TSP(S))$?

Problem 2. Das *et al.* [7, 8] showed that, in the Euclidean metric in \mathbb{R}^d , the weight of the greedy spanner is $O(wt(MST(S)))$; see Chapter 14 in [13] for a complete proof.

- Is it true that, in any metric space whose doubling dimension is constant, the weight of the greedy spanner is $O(wt(MST(S)))$?

It follows from the proofs of Lemmas 3 and 4 that the (directed) edges of any 2-optimal tour and the greedy spanner satisfy the following slightly stronger property: For any two distinct edges (p, q) and (r, s) ,

$$|pr| \geq w \cdot \min(|pq|, |rs|)$$

or

$$|qs| \geq w \cdot \max(|pq|, |rs|).$$

It is not clear, however, whether this is of any help to solve Problems 1 and 2.

Problem 3. As we have seen in this paper, edge sets in metric spaces of bounded doubling dimension can be analyzed using a packing argument.

- Find other classes of metric spaces, in which non-trivial upper bounds on the weight of 2-optimal tours, greedy spanners, or other interesting graphs, can be obtained.

5.2 Further Reading

A detailed analysis of tours produced by the 2-opt heuristic, both for general metric spaces and Euclidean spaces, can be found in Chandra *et al.* [6].

Althöfer *et al.* [1] contains an analysis of the greedy spanner for general metric spaces. Recently, Bose *et al.* [3] have shown that, in case the doubling dimension is constant, the greedy spanner can be computed in $O(n^2 \log n)$ time. Observe that, since the greedy spanner contains the all-nearest-neighbors graph, the results in Section 2.1 imply that the greedy spanner cannot be computed in

subquadratic time. We remark, however, that, in the Euclidean metric in \mathbb{R}^d , an approximation of the greedy spanner can be computed in $O(n \log n)$ time; see Gudmundsson *et al.* [9]. A detailed description of their algorithm, as well as many other algorithms for constructing spanners, can be found in Narasimhan and Smid [13].

Acknowledgements: The author thanks Hubert Chan, Anupam Gupta, and Giri Narasimhan for helpful discussions during the workshop *Geometric Networks and Metric Space Embeddings*, which was held in Dagstuhl (Germany) in November 2006.

References

1. I. Althöfer, G. Das, D. P. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993.
2. P. Assouad. Plongements lipschitziens dans \mathbb{R}^N . *Bulletin de la Société Mathématique de France*, 111:429–448, 1983.
3. P. Bose, P. Carmi, M. Farshi, A. Maheshwari, and M. Smid. Computing the greedy spanner in near-quadratic time. To appear in *Algorithmica*.
4. 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.
5. B. Chandra, G. Das, G. Narasimhan, and J. Soares. New sparseness results on graph spanners. *International Journal of Computational Geometry & Applications*, 5:125–144, 1995.
6. B. Chandra, H. Karloff, and C. Tovey. New results on the old k -opt algorithm for the traveling salesman problem. *SIAM Journal on Computing*, 28:1998–2029, 1999.
7. G. Das, P. Heffernan, and G. Narasimhan. Optimally sparse spanners in 3-dimensional Euclidean space. In *Proceedings of the 9th ACM Symposium on Computational Geometry*, pages 53–62, 1993.
8. G. Das, G. Narasimhan, and J. Salowe. A new way to weigh malnourished Euclidean graphs. In *Proceedings of the 6th ACM-SIAM Symposium on Discrete Algorithms*, pages 215–222, 1995.
9. J. Gudmundsson, C. Levcopoulos, and G. Narasimhan. Fast greedy algorithms for constructing sparse geometric spanners. *SIAM Journal on Computing*, 31:1479–1500, 2002.
10. S. Har-Peled and M. Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM Journal on Computing*, 35:1148–1184, 2006.
11. J. Heinonen. *Lectures on Analysis on Metric Spaces*. Springer-Verlag, Berlin, 2001.
12. S. Lin. Computer solutions of the traveling salesman problem. *Bell Systems Technical Journal*, 44:2245–2269, 1965.
13. G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, Cambridge, 2007.
14. F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, Berlin, 1988.
15. J. Soares. Approximating Euclidean distances by small degree graphs. *Discrete & Computational Geometry*, 11:213–233, 1994.

16. P. M. Vaidya. An $O(n \log n)$ algorithm for the all-nearest-neighbors problem. *Discrete & Computational Geometry*, 4:101–115, 1989.
17. A. C. Yao. On constructing minimum spanning trees in k -dimensional spaces and related problems. *SIAM Journal on Computing*, 11:721–736, 1982.