Spanners in l_1

J. Cáceres[1], C.I. Grima[2], A. Márquez[2] and A. Moreno-González[2]

[1] Departamento de Estadística y Matemática Aplicada, Universidad de Almería (Spain).
jcaceres@ualm.es
[2] Departamento de Matemática Aplicada I, Universidad de Sevilla (Spain).
grima@cica.es, almar@cica.es, amoreno@euler.fie.us.es

Abstract

In this work, three problems arising in geometric network design theory are considered using the l_1 metric. We first study the values of θ for which the θ-Yao graph (using the metric above) contains the MST of a collection of sites in the plane, and additionally, we consider the same problem with other l_p metrics. Secondly, we give upper bounds for the dilation of θ-Yao in the l_1 metric. And finally, we study the size of a graph with dilation 1 in the l_1 metric.
Key words: Spanner, MST, dilation, complete geometric graph, Minkowski metrics.

1 Introduction

The quality of a network interconnecting points can be measured in different ways. Typically some minimal conditions are imposed to the network; for instance, it is usually desired that the Minimum Spanning Tree (MST) must be contained in the network. But when one tries to design a good network for connecting components of a VLSI circuit such that uses little surface area on the chip, draws little power and propagates signals quickly, it could be interesting to find a sparse graph which approximates shortest paths between all pairs of vertices. Those graphs are called *spanners* and they have been extensively studied [1] [2] [3] [4] [7]. More precisely, given a set of points S in the plane, the *dilation* of a subgraph of the complete geometric graph is the largest ratio between the length of the shortest path from a pair of points of S to the distance of those points in the plane. A graph with dilation t is called a *(t)-spanner* of S. But, although the metric that reflexes the distance between components in an electronic circuit is the l_1 metric, all cited works are focused on the Euclidean metric. We try, in this work, to study some of the first questions that arise in the study of spanners but with the l_1 metric (obtaining some results for the l_∞ metric as well). Remind that the l_1 and l_∞ metrics are the Manhattan and the Supreme metrics, respectively; that is, given two points $A = (a_1, a_2)$, $B = (b_1, b_2) \in \boldsymbol{R}^2$, $d_{l_1}(A, B) = |b_1 - a_1| + |b_2 - a_2|$ and $d_{l_\infty}(A, B) = max\{|b_1 - a_1|, |b_2 - a_2|\}$.

It is possible to find sparse graphs approximating the complete Euclidean graph arbitrary closely. Thus, Keil [6] showed that a class of graphs called Yao graphs produces graphs with dilation arbitrary closed to 1, with $O(n)$ edges and that they can be constructed in time $O(n \log n)$. Thus, the first question, treated in the next section, will be to study whether a Yao graph contains the MST in both the l_1 or the l_∞ metrics. Secondly, we will study the dilation of those graphs. And, finally, we will see that in the l_1 metric graphs with dilation 1 have much less edges that in the Euclidean distance.

2 Θ-Yao graphs and MST of a collection of sites in the l_1 metric.

As it was pointed out in the Introduction, in this section we study which are the Yao graphs that contain the MST of a collection of sites in the metric l_1. First of all, we will give the θ-Yao graph construction in any metric. Let S be a collection of sites in R^2. We partition the space around each point into wedges with a given fixed opening angle, θ, and connect the point to the nearest neighbor in each wedge with the given metric. The graph obtained is called the θ-Yao graph of S.

In this work, we extend this definition and we will call (α, θ)-Yao graph to the θ-Yao graph constructed by placing the borders of the first wedge forming an angle α with the abscissae axis, as we can see in Figure 1.

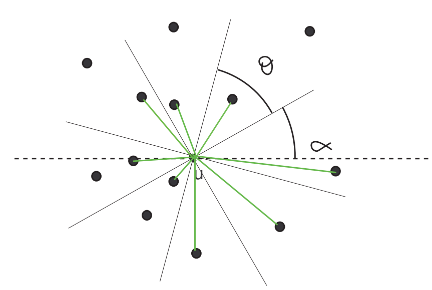

Figure 1: Construction of a (α, θ)-Yao graph.

With this new concept we want to know the values of α and θ for which the (α, θ)-Yao graph contains the MST of a collection of sites. This problem was studied by Yao [8] for the Euclidean metric and he proved that any $\pi/3$-Yao graph of a set of sites contains the MST of the sites. In this work we give similar results for the metrics l_1 and l_∞.

Theorem 1 *Let S be a set of sites in $\boldsymbol{R}^2$. Any $(\alpha, \pi/4)$-Yao graph of S contains the MST of the sites in the l_1 metric.*

Proof:
Without loss of generality, we can suppose that $\alpha \in [0, \pi/4)$.

We know that the MST of a set of sites S can be built incrementally by adding the shortest edge joining S_1 and S_2 not explored yet, which also maintains the acyclicity, (where S_1 is the subset of sites that have already been taken and $S_2 = S - S_1$).

Suppose then that in a step of this algorithm, we have to take the edge $uw, u \in S_1, w \in S_2$ and that this edge is not an edge of the $(\alpha, \pi/4)$-Yao graph of S with the l_1 metric. In this case, if we place the wedges in u, there is an edge uv, shorter than uw, that have been selected before.

Then, we only have to prove that $d(u,w) \geq d(w,v)$ in the l_1 metric. The worst case occurs when uv and uw are similar in length but widely separated in angle, as we see in Figure 2.

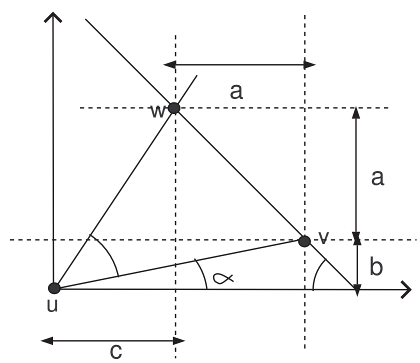

Figure 2: The worst case for the proof of $d(u,w) \geq d(w,v)$, (see text).

In Figure 2 we can also see that

$$\tan\alpha = \frac{b}{a+c}, \tan(\alpha + \pi/4) = \frac{a+b}{c}.$$

But, by other side we have that

$$\tan(\alpha + \pi/4) = \frac{\sin(\alpha + \pi/4)}{\cos(\alpha + \pi/4)} = \frac{1 + \tan\alpha}{1 - \tan\alpha}.$$

So, we get

$$\frac{a+b}{c} = \frac{1 + \frac{b}{a+c}}{1 - \frac{b}{a+c}},$$

and, simplifying we have $a^2 = b^2 + c^2$. Now, it is easy to see that $a \leq b + c$, so the result holds. □

The bound obtained in Theorem 1 is tight, and so given $\theta > \pi/4$ it is possible to find a θ-Yao graph for a set of sites that does not contain the MST.

Theorem 2 *Let S be a set of sites in $\boldsymbol{R}^2$. The $(\pi/4, \pi/2)$-Yao graph of S contains the MST of the sites in the l_1 metric, but there exists collections of sites S such the $(0, \pi/2)$-Yao graph does not contain the MST of S in that metric.*

Proof: Firstly, we prove that the $(\pi/4, \pi/2)$-Yao graph contains the MST of any set of sites. The proof is similar to that of Theorem 1, so we only have to prove that the edge vw is shorter than uw in the l_1 metric, being w a site in the wedge where v is, (see Figure 3).

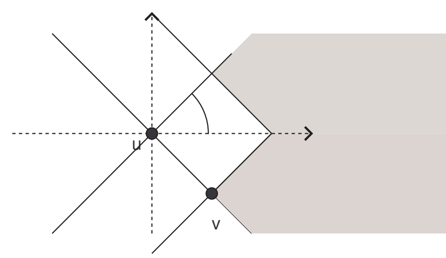

Figure 3: w is a site in the colored region.

But, as we can see in Figure 3, any site in the border of the disc with center u and radio $d(u,v)$ is at the same distance from u than from v. So, it is trivial to see that the distance between w and u is not smaller than the distance between w and v. So, the $(\pi/4, \pi/2)$-Yao graph contains the MST of any set of sites.

Now, we give an example of a collection of sites S such the $(0, \pi/2)$-Yao graph does not contain the MST of S in the l_1 metric. We consider the set $S = \{x = (0,0), y = (1'25, -0'25), z = (2,1), u = (0'5, 1'5), v = (1, 1'25)$. Then, in Figure 4 we see that the $(0, \pi/2)$-Yao graph does not contain the MST of S.

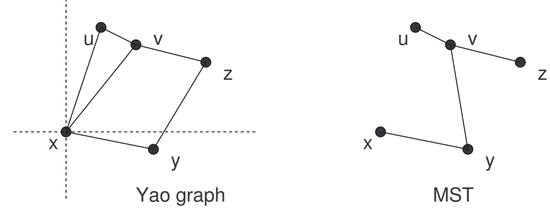

Figure 4: vy is an edge of the MST of S but it is not an edge of the $(0, \pi/2)$-Yao graph of S.

□

Our next step is to study the same problem for the l_∞ metric. Firstly, we consider the problem of finding an angle θ that satisfies that any (α, θ)-Yao graph of a set of sites contains the MST of the sites. Here, the result is similar to that given for the l_1 metric.

Theorem 3 *Let S be a set of sites in $\boldsymbol{R}^2$. Any $(\alpha, \pi/4)$-Yao graph of S contains the MST of the sites in the l_∞ metric.*

Proof: The proof of this result is similar to those of Theorems 1 and 2. As in Theorem 1, we only prove the result for $\alpha \in [0, \pi/4)$. We have to prove that the edge uw is longer than vw with the l_∞ metric, (see Figure 5).

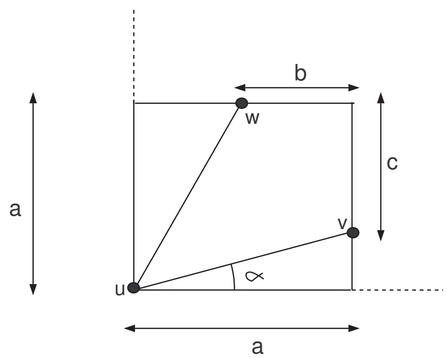

Figure 5: uw is not shorter than vw.

But, as we can see in Figure 5, it is trivial that $b, c < a$, so the result holds. □

Secondly, we study what happens for the $\pi/2$-Yao graphs of a collection of sites and we see that the situation is completely different.

Theorem 4 *Let S be a set of sites in $\boldsymbol{R}^2$. The $(0, \pi/2)$-Yao graph of S contains the MST of the sites in the l_∞ metric, but there exists collections of sites S such the $(\pi/4, \pi/2)$-Yao graph does not contain the MST of S in that metric.*

Proof: Firstly, we prove that the $(0, \pi/2)$-Yao graph of aany set of sites S contains its MST. The proof of this result is similar to that of Theorem 3, so we only have to prove that the edge uw is not shorter than the edge vw in the l_∞ metric, where w is a site in the wedge where v is, (see Figure 6).

But, it is easy to see that any site in the border of the disc with center v and radio $d(u, v)$ is nearer to v than to u, so the result holds.

Now, we give an example of a collection of sites S such the $(\pi/4, \pi/2)$-Yao graph does not contain the MST of S in the l_∞ metric. We consider the set $S = \{x = (0,0), y = (1'5, 1), z = (1,3), u = (-0'25, 2'25), v = (-1, 2)$. Then, in Figure 7 we see that the $(\pi/4, \pi/2)$-Yao graph does not contain the MST of S.

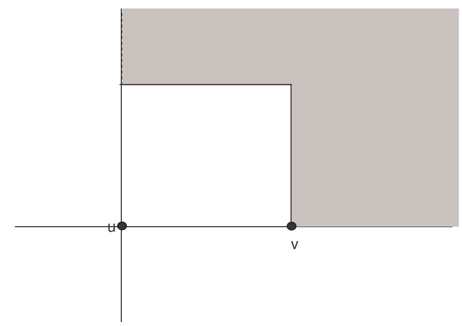

Figure 6: w lies in the colored region.

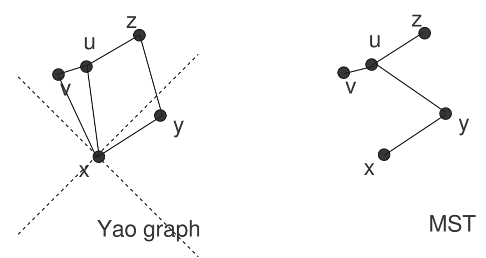

Figure 7: uy is an edge of the MST of S but it is not an edge of the $(\pi/4, \pi/2)$-Yao graph of S.

□

3 Dilation in Θ-Yao graphs.

In this section we study the dilation in the (α, θ)-Yao graphs of a collection of sites. As we did in the previous section, we give a similar result to one given by Keil [6] for the Euclidean metric. In this way, we have found upper bounds for the dilation in the (α, θ)-Yao graphs of a set of sites in the l_1 metric.

Theorem 5 *Let S be a set of sites in $\boldsymbol{R}^2$. The (α, θ)-Yao graph of the sites has dilation arbitrary closed to 1 when θ tends to 0.*

Proof: To find a path in this graph from u to v, one at each step determines the wedge containing v and moves along a graph edge to the nearest vertex, w, in that wedge. The worst case for the algorithm occurs when uv and uw are similar in length but widely separated in angle, as we see in Figure 8, but with properties of angles and triangles we can bound the dilation, as follows.

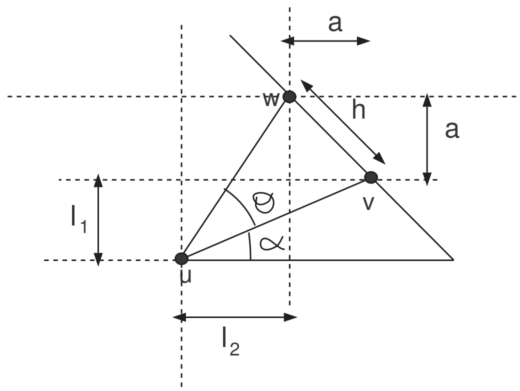

Figure 8: The worst case for the dilation.

We denote by $d_{(\alpha,\theta)}(S)$ the dilation of the (α,θ)-Yao graph of the set S. We know that

$$d_{(\alpha,\theta)}(S) = \frac{|uv| + |vw|}{|uw|}.$$

But, as we can see in Figure 8, $|uv| = |uw| = a + l_1 + l_2$ and $|vw| = 2a$, so we have that

$$d_{(\alpha,\theta)}(S) = 1 + \frac{2a}{l_1 + l_2 + a}.$$

On the other hand, we have

$$\frac{\sin\theta}{h} = \frac{\sin(\pi - \pi/4 - \alpha - \theta)}{d_e(u,v)} \geq \frac{\sin(\pi - \pi/4 - \alpha - \theta)}{|uv|}.$$

So, as $a \leq h$, we get that

$$d_{(\alpha,\theta)}(S) \leq 1 + 2\frac{\sin\theta}{\sin(\pi - \pi/4 - \alpha - \theta)},$$

that tends to 1 when θ tends to 0. □

With this result, we provide a way to construct spanners with dilation as closed to 1 as we want.

4 Graphs with dilation 1 in the l_1 metric.

As we mentioned in the Introduction, in this section we see that in the l_1 metric graphs with dilation 1 have much less edges that in the Euclidean distance. We also give three different algorithms for constructing these graphs.

Let S be a set of n sites in $\boldsymbol{R}^2$. We denote by M_n a graph of minimal size of S with dilation 1. In the Euclidean metric, except if the sites are on a straight line, M_n is the complete geometric graph of the sites K_n. That is, M_n has $\frac{n(n-1)}{2}$ edges.

With the l_1 metric this result can be improved by virtue of a result by Erdös and Szekeres [5]: In any sequence of $pq + 1$ integers, there exists an increasing subsequence of length p or a decreasing subsequence of length q.

We order the sites by their first coordinates and then, we can use the result by Erdös and Szekeres taking the second coordinates as a sequence, obtaining three different cases:

- If $(\lfloor n \rfloor + 1)\lfloor n \rfloor + 1 \leq n$, then there exist an increasing subsequence of length $\lfloor n \rfloor + 1$ and a decreasing subsequence of the same length.
- If $\lfloor n \rfloor \lfloor n \rfloor + 1 \leq n$, then there exists a decreasing or increasing subsequence of length $\lfloor n \rfloor + 1$.
- In other case, there exists a decreasing subsequence of length $\lfloor n \rfloor$ and an increasing subsequence of the same length.

In these cases, all the edges that form the complete geometric graph of the subsequences are not needed in M_n, except the ones that join the correlative sites. In Figure 9 we can see an example with 7 sites, where we have two subsequences of length 3. The edges that we save in each subsequence are marked.

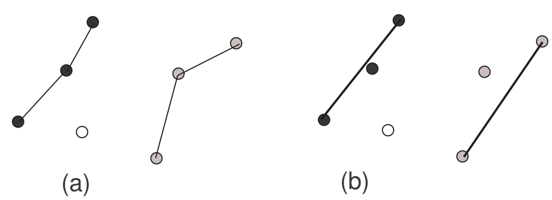

Figure 9: (a) Two subsequences of length 3; (b) edges saved in each subsequence.

Now, we can use the result again with the sites that have not been taken in the decreasing or increasing subsequences, so we obtain different subsequences of different size in which we can erase edges. We can even take a site in each subsequence and use the result by Erdös and Szekeres again.

Then, with this method, we have a way to approximate the number of edges of $K_n - M_n$. In fact we have got a function that produces this number and we have compared it with other functions obtaining that that function is in $O(n^{3/2})$.

Now, we present a result in which we give an upper bound for the size of $K_n - M_n$.

Theorem 6 *Let S be a collection of n sites in $\mathbf{R}^2$. With the l_1 metric, $|K_n - M_n| \in O(n^{3/2})$.*

Proof: Let S be a set of sites in $\mathbf{R}^2$ and let $L_1 \dots L_k$ its convex layers. In any of these layers we have four different decreasing or increasing chains of sites, (see Figure 10).

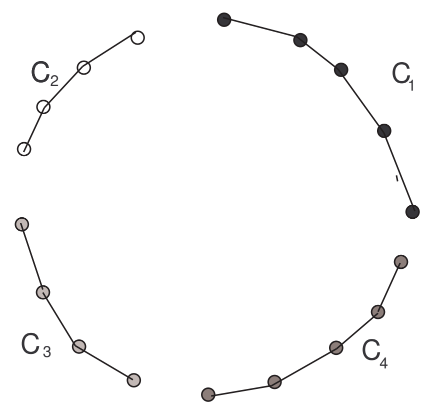

Figure 10: An example of the four chains in a convex layer.

In Figure 10 we can also see that in each chain we only need the edges joining correlative sites, that is, in each chain C_i we do not use

$$\frac{|C_i|(|C_i|-1)}{2} - (|C_i|-1)$$

edges.

Given any convex layer L_i, the worst case is when we have $|L_i|/4$ sites in each chain. In this case, we do not use

$$\frac{\frac{|L_i|}{4}(\frac{|L_i|}{4}-1)}{2} - (\frac{|L_i|}{4}-1)$$

edges in each chain, so the number of edges not needed in a lawyer is four times the previous one.

Then, as we have k convex layers, it is trivial to see that the total number of edges that we do not use is

$$\sum_{i=1}^{k} \frac{\frac{|L_i|}{4}(\frac{|L_i|}{4}-1)}{2} - (\frac{|L_i|}{4}-1).$$

Now, if we study the previous expression, we see that the worst case is when $k = \sqrt{n}$ and there are $\sqrt{n}$ sites in each layer. So, we get that at least

$$\sqrt{n}(\frac{n}{8} - \frac{3\sqrt{n}}{2} + 4)$$

edges are not needed in M_n.

On the other side, we can consider a site in the last convex layer and then partition the space around this point into four wedges with borders parallel to the axis. Then, we do not need the edges joining sites in the first wedge and the third one and edges joining sites in the second wedges and the fourth one. This happens because we have a path between these sites, (see Figure 11).

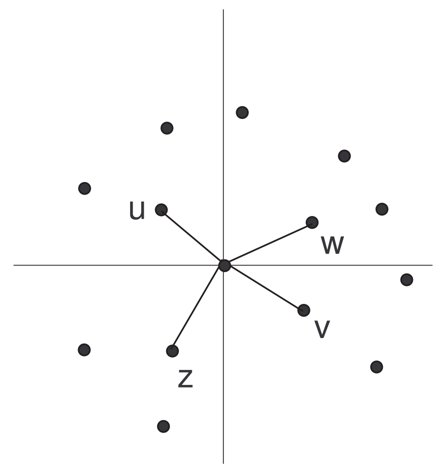

Figure 11: We do not need the edges uv and wz.

Now, in the last convex layer there are $\sqrt{n}$ sites and if we study the position of the rest of sites we obtain that the worst case is when we have $\sqrt{n}/2$ sites in the first and second wedges and $(n-2\sqrt{n})/2$ in the third and fourth wedges. So, we save $2\frac{\sqrt{n}}{2}(\frac{n-2\sqrt{n}}{2})$ edges.

In conclusion, $|K_n - M_n|$ has at least $\frac{5}{8}n\sqrt{n} - \frac{5}{2}n + 4\sqrt{n}$ edges, so $|K_n - M_n| \in O(n^{3/2})$. $\square$

Corollary 7 *Let S be a set of sites in $\mathbf{R}^2$. With the l_∞ metric, $|K_n - M_n| \in O(n^{3/2})$.*

Proof: The proof of this result is based in the relation between the l_1 and the l_∞ metrics: if we consider a disc with center $u \in \mathbf{R}^2$ and radio r with the $l_1$metric and we rotate the plane an angle of $\pi/4$, we get the disc with center u and radio r in the l_∞ metric.

Then, to construct the graph $K_n - M_n$ of a set of sites S in the l_∞ metric, we only have to rotate the plane an angle of $\pi/4$, construct $K_n - M_n$ in the l_1 metric and rotate the plane again an angle of $-\pi/4$, as we can see in Figure 12.

$\square$

The question that arises now is to compare the two methods we have given to approximate the size of $|K_n -$

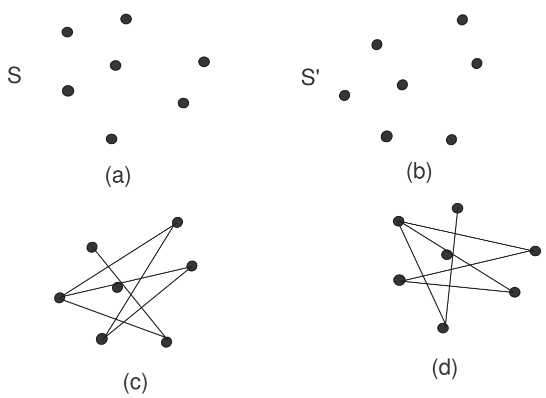

Figure 12: (a) A set of sites S; (b) S rotated an angle of $\pi/4$, S'; (c) the graph $K_n - M_n$ of S' in the l_1 metric; (d) the graph $K_n - M_n$ of S in the l_∞ metric.

$M_n|$. In this way, we introduce now some results for different sets of sites, as we can see in Figure 13, where n is the size of S, E_1 the edges we save using the result by Erdös and Szekeres and E_2 the edges we do not use with the method given in the proof of Theorem 6.

n	E_1	E_2
10^3	10.705	16.999
10^4	345.118	600.400
10^5	10.789.476	19.501.264
10^6	338.247.777	622.504.000
$2 \cdot 10^6$	954.163.715	1.762.505.656

Figure 13: Some results of the size of $K_n - M_n$.

4.1 Three algorithms to construct M_n.

As it was pointed out above, we give here three different algorithms for constructing the graph M_n of any set of sites in the plane. But, as $K_n - M_n$ and M_n are complementary graphs, these algorithms let us to construct the graph $K_n - M_n$, too. In this way, we present the first and the second algorithms to get $K_n - M_n$ and the third one to construct M_n.

We must say that these three algorithms let us to construct the graphs M_n and $K_n - M_n$ of a set of sites in the l_∞ metric. As we did in Corollary 7, we only have to rotate the plane an angle of $\pi/4$, construct the graph M_n or $K_n - M_n$ of the new set of sites and rotate the plane again an angle of $-\pi/4$.

4.1.1 The first algorithm.

The first algorithm runs in time $O(n^3)$ in the worst case, but we think that the average-case running time is much better. This algorithm is based in the following assert: Let S be a set of sites in $\boldsymbol{R}^2$. Then if we place the axis in $u \in S$ we save the edges that join sites in the first quadrant with sites in the third one. The same happens with the sites in the second and the fourth quadrant, as we saw in the proof of Theorem 6.

So, if we place the axis in all the sites of S, we only have to add the edges joining the sites in the position we said above for getting $K_n - M_n$. This is what the algorithm does. Let S be a set of n sites in $\boldsymbol{R}^2$.

Firstly, we order the sites by the second coordinate $p_1, p_2, \ldots, p_n$, that can be done in time $O(n \log n)$.

Secondly, we visit all the sites of S from p_1 to p_n. When we place the axis in a site p_j we consider two lists, l_T and l_B. In l_T we have the sites in the second and the first quadrant ordered by the first coordinate and separated by a pointer M and in l_B we have the points of the third and the fourth quadrant ordered by the first coordinate and separated by a pointer N. we can maintain the two lists in time $O(n \log n)$.

In the first step, we have all the sites in l_T with the pointer M in the place of p_1 and in l_B we only have the pointer N. Then for $k = 1, \ldots, n$ the lists change as follows. In l_T we put M in the place of p_k. In l_B we add p_{k-1} in the place of N and we put N in the place of p_k. In Figure 14 we can see an example for a set of 7 sites.

At last we have to add the edges joining the sites on the left of M with the sites on the right of N and the sites on the right of M with the sites on the left of N that have not been considered yet. Each step can be done in quadratic time, so the whole algorithm runs in time $O(n^3)$.

4.1.2 The second algorithm.

The second algorithm we present to construct $K_n - M_n$ runs in time $O(n^2 \log n)$ and is based in the following result: given two sites u and v in the plane, we save the edge uv if there is a site, different from u and v, in the rectangle that u and v form, (see Figure 15).

Then, let S be a set of sites in $\boldsymbol{R}^2$. We consider a pair of sites u and v, we check if there is another site in the rectangle that they form and in affirmative case, we add the edge uv. Now, to check if there is any site in a rectangle takes time $O(\log n + k)$, where k is the number of points inside the rectangle. But we stop when we find one site, so each step of the algorithm can

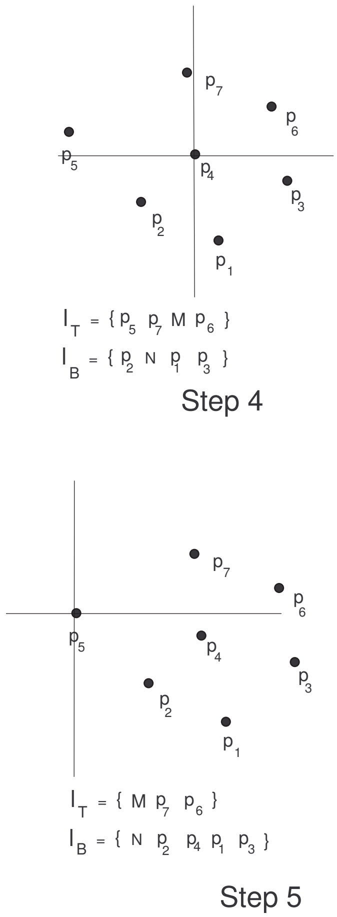

Figure 14: Two steps of the first algorithm.

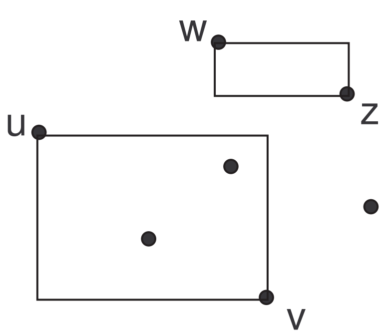

Figure 15: We save the edge uv, but not the wz.

be done in $O(\log n)$. As we have n^2 pairs of sites, the whole algorithm runs in time $O(n^2 \log n)$.

4.1.3 The third algorithm.

Here, we present a third algorithm which constructs the graph M_n of a point set in the plane. This algorithm runs in optimal time $O(n^2)$ in the worst case but, we think that the average-case running time is worst than the obtained by the two algorithms given above.

Without loss of generality, suppose that the points $\{p_1, p_2, \ldots, p_n\}$ have been ordered from left to right (that can be done in $O(n \log n)$). For simplicity, we split the algorithm in two steps but it is not needed they are implemented separately.

The first step consists of constructing a binary tree T with $\{p_1, p_2, \ldots, p_n\}$ as its vertices which records the order of the points from up to down. Let p_1 be the root of the tree then, one of its descendant subtrees contains all the points which are higher than p_1 and the other subtree contains the points which are lower. The tree can be constructed simply by inserting the points successively in order and every insertion takes time at most $O(\log n)$ (see Figure 16 as an example), so the whole step can be done in $O(n \log n)$. We will referred the subtrees of every non-leaf node as its *high* and *low subtrees*.

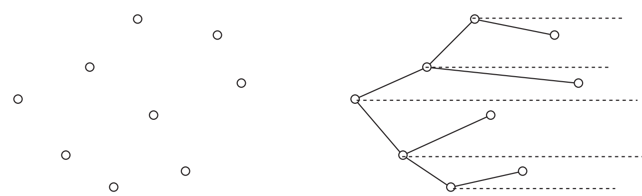

Figure 16: An example of tree

As it was said above, the second step can be done simultaneously with the first one. Consider a point p_j which have been just inserted in the tree. Now, the goal is to find the previous points which are joined with p_j in the graph M_n. Clearly, p_j is joined in M_n with its parent and with the parent of its parent if and only if p_j is a low son of a high son or vice versa. It is not difficult to check this claim and that no other ancestor of p_j is joined with it in M_n.

Now, for every ancestor p_i of p_j we will make some operations. For the sake of simplicity, let us suppose that p_j is containing in the low subtree of p_i as you can see in Figure 17 (the other case is treated in a symmetric way). Next we will find the highest leaf of the low subtree and the lowest one of the high subtree and call them p_k and p_l respectively. If $p_k \neq p_j$, we set the variable r as k, otherwise $r := 0$.

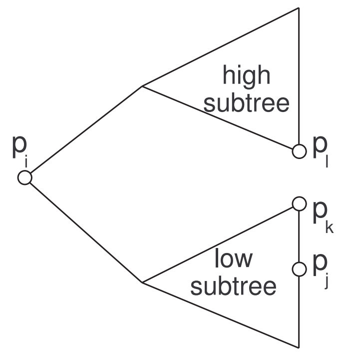

Figure 17: For every ancestor p_i of p_j, this is one of the two possible situations.

Finally, we will explore the nodes of the the high subtree, beginning with p_l in such a way that a node is visited after all of its its descendants have been visited. For every such node p_m, then:

- If p_m is a leaf and $m > r$, then the edge $p_j p_m$ belongs to M_n. Also, we set $r := m$.
- If p_m is not a leaf and it has no low subtree then $p_j p_m$ is an edge of the graph M_n

Adding the edges of the tree T to the final result, we get the graph M_n. Since in the worst case, every point is needed to be checked with all the previous points in the order, the whole algorithm runs in time $O(n^2)$ but this is optimal.

5 Conclusions and open problems

In this work we have proved that any $(\alpha, \pi/4)$-Yao graph contains the MST of a collection of sites in the l_1 and l_∞ metrics and we have studied what happens with some $(\alpha, \pi/2)$-Yao graphs. It could be interesting to study the rest of cases and try to generalize these results to other l_p metrics. Other question related with Yao graphs is to find upper bounds for the dilation in those metrics, as we have done for the l_1.

We also have studied graphs with dilation 1 in the l_1 metric and the number of edges that we do not need to construct them. We have obtained that $|K_n - M_n| \in O(n^{3/2})$, so the open question is to find better upper bounds for the size of $K_n - M_n$. In fact, we have found some particular cases, for instance if the sites are in convex position, where $|K_n - M_n| \in O(n^2)$.

References

[1] L.P. Chew. *There are planar graphs almost as good as the complete graph.* J. Comput. System. Sci., 39 (1989), pp. 205–219

[2] G. Das and D. Joseph. *Which triangulations approximate the complete graph?* Proc. Int. Symp. Optimal Algorithms, Springer LNCS 401 (1989), pp. 168–192

[3] G. Das and G. Narasimhan. *A fast algorithm for constructing sparse Euclidean spanners.* Proc. 10th ACM Symp. Comp. Geom., (1994), pp. 132–139

[4] D. Epstein. *Spanning Trees and Spanners.* Handbook of Computational Geometry. Edited by J.-R. Sack and J. Urrutia. Elsevier Science B. V. (1999), pp. 425–461

[5] P. Erdös and A. Szekeres. *A combinatorial problem in geometry.* Compositio Mathematica, vol. 2 (1935) 463-470.

[6] J. M. Keil. *Approximating the complete Euclidean graph.* Proc. 1st. Scand. Worksh. Algorithm Theory, Springer LNCS 318 (1988), pp. 208–213.

[7] C. Levcopoulos and A. Lingas. *There are planar graphs almost as good as the complete graphs and as short as minimum spanning trees.* Proc. Int. Symp. Optimal Algorithms, Springer LNCS 401 (1989), pp. 9–13

[8] A. C. Yao. *On constructing minimum spanning trees in k-dimensional space and related problems.* SIAM J. Comput. (1982), no. 11, pp. 721–736.