The Erdős-Sós Conjecture for Geometric Graphs[†]

Luis Barba[1,2] Ruy Fabila-Monroy[3‡] Dolores Lara[4†] Jesús Leaños[5]
Cynthia Rodríguez[6,7] Gelasio Salazar[8§] Francisco J. Zaragoza[6]

[1]*Boursier FRIA du FNRS, Département d'Informatique, ULB, Belgium*
[2]*School of Computer Science, Carleton University, Canada*
[3]*Departamento de Matemáticas, CINVESTAV, Mexico*
[4]*Departament de Matemàtica Aplicada II, Universitat Politènica de Catalunya, Barcelona, Spain.*
[5]*Unidad Académica de Matemáticas, UAZ. Mexico*
[6]*Departamento de Sistemas, UAM Azcapotzalco, Mexico*
[7]*Department of Combinatorics and Optimization, University of Waterloo, Canada*
[8]*Instituto de Física, UASLP, Mexico*

received 5th July 2012, *revised 19th January 2013*, *accepted 8th February 2013*.

Let $f(n,k)$ be the minimum number of edges that must be removed from some complete geometric graph G on n points, so that there exists a tree on k vertices that is no longer a planar subgraph of G. In this paper we show that $\left(\frac{1}{2}\right)\frac{n^2}{k-1} - \frac{n}{2} \leq f(n,k) \leq 2\frac{n(n-2)}{k-2}$. For the case when $k = n$, we show that $2 \leq f(n,n) \leq 3$. For the case when $k = n$ and G is a geometric graph on a set of points in convex position, we completely solve the problem and prove that at least three edges must be removed.

Keywords: extremal graph theory, geometric graph, spanning tree

1 Introduction

One of the most notorious problems in extremal graph theory is the Erdős-Sós Conjecture, which states that every simple graph with average degree greater than $k-2$ contains every tree on k vertices as a subgraph. This conjecture was recently proved true for all sufficiently large k (unpublished work of Ajtai, Komlós, Simonovits, and Szemerédi).

In this paper we investigate a variation of this conjecture in the setting of geometric graphs. Recall that a *geometric graph* G consists of a set S of points in the plane (these are the vertices of G), plus a set of straight line segments, each of which joins two points in S (these are the edges of G). In particular, any set S of points in the plane in *general position* (no three of its points are collinear) naturally induces a complete geometric graph. For brevity, we often refer to the edges of this graph simply as edges *of* S. If

[†]A previous version of this paper appeared in the conference proceedings of EUROCG'12 [2].
[‡]Partially supported by CONACYT of Mexico, Grant 153984.
[§]Partially supported by CONACYT of Mexico, Grant 106432.

S is in convex position then G is a *convex geometric graph*. A geometric graph is *planar* if no two of its edges cross each other. An *embedding* of an abstract graph H into a geometric graph G is an isomorphism from H to a planar geometric subgraph of G. For $r \geq 0$, an *r-edge* is an edge of G such that in one of the two open semi-planes defined by the line containing it, there are exactly r points of G. The convex hull of S is the intersection of all convex sets containing S. We will frequently need to refer to the vertices and edges at the boundary of the convex hull of S, which for brevity we will denote simply as *convex hull vertices* and *convex hull edges* of S.

In this paper all point sets are in general position and G is a complete geometric graph on n points. It is well known that for every integer $1 \leq k \leq n$, G contains every tree on k vertices as a planar subgraph [3]. Even more, it is possible to embed any such tree into G, when the image of a given vertex is prespecified [5].

Let F be a subset of edges of G, which we call *forbidden edges*. If T is a tree for which every embedding into G uses an edge of F, then we say that *F forbids T*. In this paper we study the question of what is the minimum size of F so that there is a tree on k vertices that is forbidden by F. Let $f(n, k)$ be the minimum of this number taken over all complete geometric graphs on n points. As $f(2, 2) = 1$, $f(3, 3) = 2$, $f(4, 3) = 3$, $f(4, 4) = 2$ and $f(n, 2) = \binom{n}{2}$, we assume through out the paper that $n \geq 5$ and $k \geq 3$.

We show the following bounds on $f(n, k)$.

Theorem 1.1
$$\left(\frac{1}{2}\right)\frac{n^2}{k-1} - \frac{n}{2} \leq f(n,k) \leq 2\frac{n(n-2)}{k-2}$$

Theorem 1.2
$$2 \leq f(n,n) \leq 3$$

In the case when G is a convex complete geometric graph, we show that the minimum number of edges needed to forbid a tree on n vertices is three.

An equivalent formulation of the problem studied in this paper is to ask how many edges must be removed from G so that it no longer contains *every* planar subtree on k vertices. A related problem is to ask how many edges must be removed from G so that it no longer contains *any* planar subtree on k vertices. For the case of $k = n$, in [6], it is proved that if any $n - 2$ edges are removed from G, it still contains a planar spanning subtree. Note that if the $n - 1$ edges incident to any vertex of G are removed, then G no longer contains a spanning subtree. In general, for $2 \leq k \leq n - 1$, in [1], it is proved that if any set of $\left\lceil \frac{n(n-k+1)}{2} \right\rceil - 1$ edges are removed from G, it still contains a planar subtree on k vertices. In the same paper it is also shown that this bound is tight—a geometric graph on n vertices and a subset of $\left\lceil \frac{n(n-k+1)}{2} \right\rceil$ of its edges are shown, so that when these edges are removed, every planar subtree has at most $k - 1$ vertices. In [4] the authors study the seemingly unrelated (non-geometric) problem of packing two trees into planar(i) (abstract) graphs. That is, given two trees on n vertices, the authors consider the question of when it is possible to find a planar graph having both of them as spanning trees and in which the trees are edge disjoint. However, although theirs is a combinatorial question rather than

(i) A *planar* (abstract) graph is an *abstract* graph that can be embedded in the plane; the embedding may not be unique.

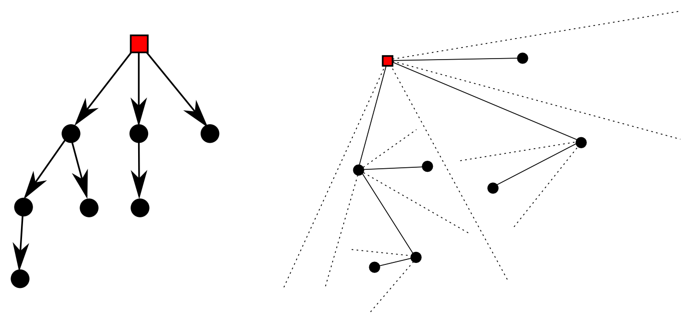

Fig. 1: An embedding of a tree using the algorithm.

geometric, their Theorem 2.1 implies our Lemma 2.2. We provide a self contained proof of Lemma 2.2 for completeness.

A previous version of this paper appeared in the conference proceedings of EUROCG'12 [2].

2 Spanning Trees

In this section we consider the case when $k = n$. Let T be a tree on n vertices. Consider the following algorithm to embed T into G. Choose a vertex v of T and root T at v. For every vertex of T choose an arbitrary order of its children. Suppose that the neighbors of v are $u_1, \dots, u_m$, and let $n_1, \dots, n_m$ be the number of nodes in their corresponding subtrees. Choose a convex hull point p of G and embed v into p. Sort the remaining points of G counter-clockwise by angle around p. A *wedge* is a region of the plane bounded by two infinite rays sharing a common apex. Choose $m + 1$ rays centered at p so that the wedge between two consecutive rays is convex and between the i-th ray and the $(i+1)$-th ray there are exactly n_i points of G. Let S_i be this set of points. A convex hull vertex q of S_i is *visible* from p if the line segment with endpoints p and q intersects the convex hull of S only at q. For each u_i choose a convex hull vertex of S_i visible from p and embed u_i into this point. Recursively embed the subtrees rooted at each u_i into S_i. Note that this algorithm provides an embedding of T into G. We will use this embedding frequently throughout the paper. See Figure 1.

For every integer $n \geq 2$ we define a tree T_n as follows: If $n = 2$, then T_n consists of only one edge; if n is odd, then T_n is constructed by subdividing once every edge of a star on $\frac{n+1}{2}$ vertices; if n is even and greater than 2, then T_n is constructed by subdividing an edge of T_{n-1}. See Figure 2.

We prove the lower bound of $f(n, n) \geq 2$ of Theorem 1.2.

Theorem 2.1 *If G has only one forbidden edge, then any tree on n vertices can be embedded into G, without using the forbidden edge.*

Proof: Let e be the forbidden edge of G. Let T be a tree on n vertices. Choose a root for T. Sort the children of each node of T, by increasing size of their corresponding subtrees. Embed T into G with the embedding algorithm, choosing at all times the *rightmost* point (the first point when sorting clockwise around the root) as the root of the next subtree. Suppose that e is used in this embedding. Let $e := pq$ so that u is embedded into p and v is embedded into q (note that u is the parent of v in T).

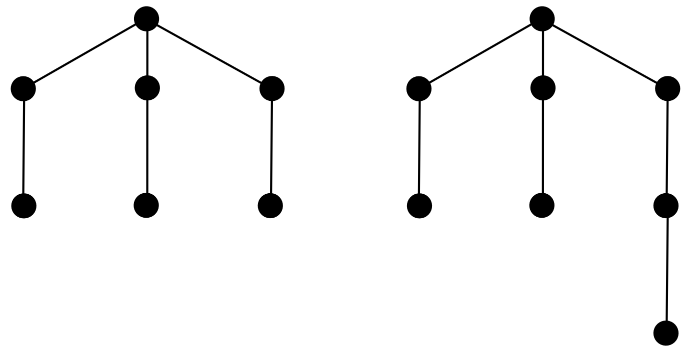

Fig. 2: T_7 and T_8.

Suppose that the subtree rooted at v has $m \geq 2$ nodes. In the algorithm, we embedded this subtree into a set of exactly m points. We chose a convex hull point (q), of this set visible from p to embed v. In this case we may choose another convex hull point visible from p to embed v and continue with the algorithm. Note that pq is no longer used in the final embedding.

Suppose that v is a leaf, and that v has a sibling v' whose subtree has at least two nodes. Then we may interchange v and v' in the order of the children of u, so that e is no longer used in the embedding, or if it is, then v' is embedded into q, but then we proceed as above.

Suppose that v is a leaf, has at least one sibling and all its siblings are leaves. The subtree rooted at u is a star. We choose a point distinct from p and q in the point set where this subtree is embedded, and embed u into this point. Afterward we join it to the remaining points. This produces an embedding that avoids e.

Assume then, that v is a leaf and that it has no siblings. We distinguish the following cases:

1. u **has no siblings.** In this case, the subtree rooted at the parent of u is a path of length two. If u has no grandparent then $n = 3$ and T can be trivially embedded into G without using e. Suppose u has a grandparent. In this case there are only four vertices to consider: v, u, the parent of u and the grandparent of u. We keep the current location of the grandparent of u, and change the points into which the remaining vertices are embedded. This can always be done so that e is not used in the embedding. All possible cases are shown in Figure 3.

2. u **has a sibling** u' **whose subtree is not an edge.** We may change the order of the siblings of u, with respect to their parent, so that the subtree rooted at u' will be embedded into the point set containing p and q. In the initial order—increasing by size of their corresponding subtrees—u' is after u. We may assume that in the new ordering, the order of the siblings of u before it, stays the same. Therefore p is the rightmost point of the set into which the subtree rooted at u' will be embedded. Embed u' into p. Either we find an embedding not using e, or this embedding falls into one of the cases considered before.

3. u **has at least one sibling, and the subtree at every sibling of** u **is an edge.**

   Suppose that u has no grandparent; then T is equal to T_n and n is odd. Let w be the parent of u. Embed w into p. Let $p_1, \ldots, p_{n-1}$ be the points of G different from p sorted counter-clockwise

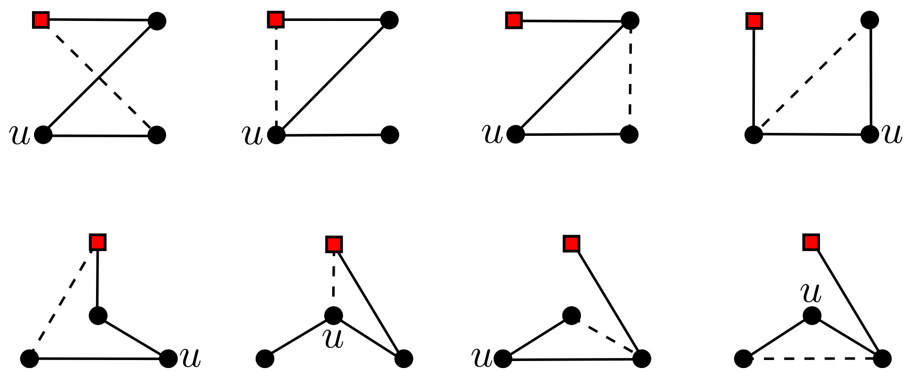

Fig. 3: The embedding of a path of length three. The grandparent of u is highlighted and the forbidden edge is dashed.

by angle around p; choose p_1 so that the angle between two consecutive points is less than π. Let $u_1, \ldots, u_{(n-1)/2}$ be the neighbors of w. Embed each u_i into p_{2i-1} and its child into p_{2i}. If q equals p_{2j-1} for some j then embed u_j into p_{2j} and its child into p_{2j-1}. This embedding avoids e.

Suppose that w is the grandparent of u and let p' be the point into which w is embedded. Let S be the point set into which the subtree rooted at the parent of u is embedded. Note that S has an odd number of points. We replace the embedding as follows. Sort S counter-clockwise by angle around p'. Call a point *even* if it has an even number of points before it in this ordering. Call a point *odd* if it has an odd number of points before it in this ordering. If e is incident to an odd point, then we embed the parent of u into this point. The remaining subtree rooted at u can be embedded without using e. If the endpoints of e are both even, between them there is an odd point. We embed the parent of u into this point. The remaining vertices can be embedded without using e (see Figure 4).

□

The upper bound of $f(n, n) \leq 3$ of Theorem 1.2 follows directly from Lemma 3.2. Now we prove in Lemma 2.2 and Theorem 2.3, that if G is a convex geometric graph, at least three edges are needed to forbid some tree on n vertices.

Lemma 2.2 *Let T be a tree on n vertices. If G is a convex geometric graph, then T can be embedded into G using less than $\frac{n}{2}$ convex hull edges of G.*

Proof: If T is a star, then any embedding of T into G uses only two convex hull edges. If T is a path then it can be embedded into G using at most two convex hull edges. Therefore, we may assume that T is neither a star nor a path.

Since T is not a path, it has a vertex of degree at least three. Choose this vertex as the root. Since T is not a star, the root has a child whose subtree has at least two nodes. Order the children of T so that this node is first. Embed T into G with the embedding algorithm.

Let u and v be vertices of T, so that u is the parent of v. Suppose that the subtree rooted at v has at least two nodes. Then in the embedding algorithm we have at least two choices to embed v once the ordering

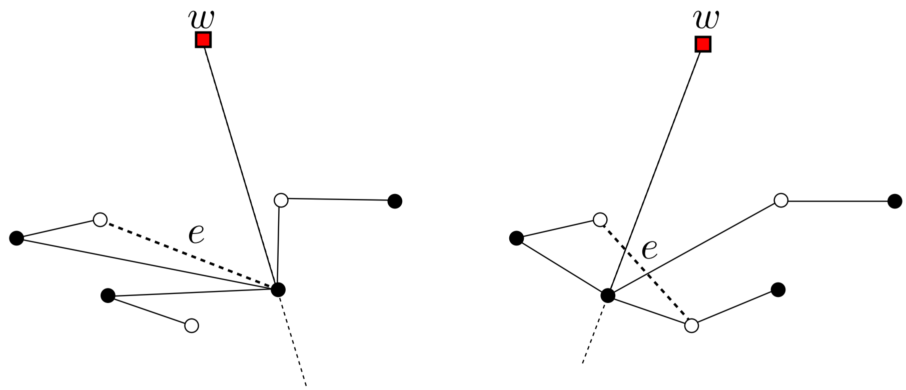

Fig. 4: The two sub-cases, when u has a grandparent w, and all the subtrees of its children are edges. Odd points are painted in black and even points in white. The forbidden edges are dashed.

of the children of u has been chosen. At least one of the choices is such that uv is not embedded into a convex hull edge. Therefore, we may assume that the embedding is such that each convex hull edge used, is incident to a leaf.

Note that every vertex of T, distinct from the root, is incident to at most one convex hull edge in the embedding. Since the first child of the root is not a leaf, no convex hull edge is used to embed this child. Only in the embedding of the last child of the root a convex hull may have been used. Therefore every vertex of T is incident to at most one convex hull edge. Thus the set of convex hull edges used in the embedding is a matching. Therefore at most $n/2$ convex hull edges are used in the embedding.

Suppose that exactly $n/2$ convex hull edges are used. One of these edges must be incident to the root. Since the root was chosen of degree at least three it has a child which is not a leaf nor the first child; we place this vertex last in the ordering of the children of the root. The leaf adjacent to the root can no longer be a convex hull edge and the embedding uses less than $n/2$ convex hull edges. □

Theorem 2.3 *If G is a convex geometric graph and has at most two forbidden edges, then any tree on n vertices can be embedded into G, without using a forbidden edge.*

Proof: Let f_0 be an embedding given by Lemma 2.2, of T into G. For $0 \leq i \leq n$, let f_i be the embedding produced by rotating f_0, i places to the right. Assume that in each of these rotations at least one forbidden edge is used, as otherwise we are done. Let $e_1, \ldots, e_m$ be the edges of T that are mapped to a forbidden edge in some rotation. Assume that the two forbidden edges are an l-edge and an r-edge respectively.

Suppose that $l \neq r$. Then, each edge of T can be embedded into a forbidden edge at most once in all of the n rotations. Thus $m \geq n$. This is a contradiction, since T has $n-1$ edges.

Suppose that $l = r$. Then, each of the e_i is mapped twice to a forbidden edge. Thus $m \geq n/2$. By Lemma 2.2, f_0 uses less than $n/2$ convex hull edges. Therefore, $l = r > 0$. But a set of $n/2$ or more r-edges, with $r > 0$, must contain a pair of edges that cross. And we are done, since f_0 is an embedding. □

3 Bounds on $f(n, k)$

In this section we prove Theorem 1.1. First we show the upper bound.

Lemma 3.1 *If T_n is embedded into G_n then every edge incident to a leaf of T_n must be embedded into a convex hull edge.*

Proof: Let $e := uv$ be an edge of T_n incident to leaf. Suppose that u is the leaf vertex. Then v is of degree two. Suppose that e is not embedded into a convex hull edge of G. Then e divides $S \setminus \{u, v\}$ into two non-empty subsets S_1 and S_2, so that S_1 lies on the opposite side of S_2 with respect to e. Assume that the parent of v is embedded into S_1. Then no vertex of T_n can be embedded into S_2 without crossing e. Therefore e must be a convex hull edge of G. □

Lemma 3.2 *If G is a convex geometric graph, then forbidding three consecutive convex hull edges of G forbids the embedding of T_n.*

Proof: Recall that T_n comes from subdividing a star, let v be the non leaf vertex of this star. Let p_1p_2, p_2p_3, p_3p_4 be the forbidden edges, in clockwise order around the convex hull of G. Note that by Lemma 3.1, in any embedding of T_n into G, an edge incident to a leaf of T_n, must be embedded into a convex hull edge. Neither the leaves of T_n nor its neighbors can be embedded into p_2 or p_3, without using a forbidden edge. Thus, v must be embedded into p_2 or p_3. Without loss of generality assume that v is embedded into p_2. But then, the embedding must use p_2p_3 or p_3p_4 □

Lemma 3.3 *If G is a convex geometric graph, then forbidding the convex hull edges incident to any three vertices p_1, p_2 and p_3 of G, forbids the embedding of T_n.*

Proof: Note that by Lemma 3.1, neither a leaf of T_n, nor its neighbor can be embedded into p_1, p_2 or p_3, without using a forbidden edge. But at most two points do not fall into this category. □

Lemma 3.4

$$f(n, k) \leq 2\frac{n(n-2)}{k-2}$$

Proof: Let G be a complete convex geometric graph. We forbid every r-edge of G for $r = 0, \ldots, \left\lceil 2\frac{n-2}{k-2} - 2 \right\rceil$. Note that, in total we are forbidding at most $n\left(\left\lceil 2\frac{n-2}{k-2} - 2 \right\rceil + 1\right) \leq 2\frac{n(n-2)}{k-2}$ edges. As every subset of points of G is in convex position, it suffices to show that every induced subgraph H of G on k vertices is in one of the two configurations of Lemmas 3.2 and 3.3.

Assume then, that H does not contain three consecutive forbidden edges in its convex hull nor three vertices, each with its two convex hull edges forbidden. H has at most two (non-adjacent) pairs of consecutive forbidden edges in its convex hull. Therefore every forbidden edge of H in its convex hull—with the exception of at most two—must be preceded by an ℓ-edge (of G), with $\ell > \left\lceil 2\frac{n-2}{k-2} - 2 \right\rceil$. The number of these ℓ-edges contained in H is at least $\frac{k-2}{2}$. The points separated by these edges amount to more than $\frac{k-2}{2}\left\lceil 2\frac{n-2}{k-2} - 2 \right\rceil \geq n - k$ points of G. This is a contradiction, since together with the k points of H this is strictly more than n. □

Now, we show the lower bound of Theorem 1.1.

Lemma 3.5
$$f(n,k) \geq \left(\frac{1}{2}\right)\frac{n^2}{k-1} - \frac{n}{2}$$

Proof: Let F be a set of edges whose removal from G forbids some k-tree. Let $H := G \setminus F$. Note that H contains no complete K_k as a subgraph, otherwise any k-tree can be embedded into this subgraph. By Turán's Theorem [7], H cannot contain more than $\left(\frac{k-2}{k-1}\right)\frac{n^2}{2}$ edges. Thus F must have size at least

$$\binom{n}{2} - \left(\frac{k-2}{k-1}\right)\frac{n^2}{2} = \left(\frac{1}{2}\right)\frac{n^2}{k-1} - \frac{n}{2}$$

□

Acknowledgments

Part of this work was done at the "First Workshop in Combinatorial Optimization at Cinvestav". It was continued during a visit of L. Barba, R. Fabila-Monroy, J. Leaños and G. Salazar to Abacus research center[(ii)].

References

[1] O. Aichholzer, S. Cabello, R. Fabila-Monroy, D. Flores-Peñaloza, T. Hackl, C. Huemer, F. Hurtado, and D. R. Wood. Edge-removal and non-crossing configurations in geometric graphs. *Discrete Math. Theor. Comput. Sci.*, 12(1):75–86, 2010.

[2] L. F. Barba, R. Fabila-Monroy, D. Lara, J. Leaños, C. Rodríguez, G. Salazar, and F. Zaragoza. The Erdős-Sós conjecture for geometric graphs. In *Proc. 28th European Worskhops in Computational Geometry, EUROCG '12*, Asisi, Italy, 2012.

[3] P. Bose, M. McAllister, and J. Snoeyink. Optimal algorithms to embed trees in a point set. *J. Graph Algorithms Appl.*, 1(2):15 pp. (electronic), 1997.

[4] A. García, C. Hernando, F. Hurtado, M. Noy, and J. Tejel. Packing trees into planar graphs. *J. Graph Theory*, 40(3):172–181, 2002.

[5] Y. Ikebe, M. A. Perles, A. Tamura, and S. Tokunaga. The rooted tree embedding problem into points in the plane. *Discrete Comput. Geom.*, 11(1):51–63, 1994.

[6] G. Károlyi, J. Pach, G. Tóth, and P. Valtr. Ramsey-type results for geometric graphs. II. *Discrete Comput. Geom.*, 20(3):375–388, 1998.

[7] P. Turán. On an extremal problem in graph theory. *Mat. Fiz. Lapok*, 48:436–452, 1941.

(ii) ABACUS, CONACyT grant EDOMEX-2011-C01-165873