Universal Geometric Graphs*

Fabrizio Frati[1], Michael Hoffmann[2], and Csaba D. Tóth[3,4]

[1] Roma Tre University, Rome, Italy
frati@dia.uniroma3.it
[2] ETH Zürich, Zürich, Switzerland
hoffmann@inf.ethz.ch
[3] California State University Northridge, Los Angeles, CA, USA
[4] Tufts University, Medford, MA, USA
csaba.toth@csun.edu

Abstract. We introduce and study the problem of constructing geometric graphs that have few vertices and edges and that are universal for planar graphs or for some sub-class of planar graphs; a geometric graph is *universal* for a class $\mathcal{H}$ of planar graphs if it contains an embedding, i.e., a crossing-free drawing, of every graph in $\mathcal{H}$.
Our main result is that there exists a geometric graph with n vertices and $O(n \log n)$ edges that is universal for n-vertex forests; this extends to the geometric setting a well-known graph-theoretic result by Chung and Graham, which states that there exists an n-vertex graph with $O(n \log n)$ edges that contains every n-vertex forest as a subgraph. Our $O(n \log n)$ bound on the number of edges is asymptotically optimal.
We also prove that, for every $h > 0$, every n-vertex convex geometric graph that is universal for the class of the n-vertex outerplanar graphs has $\Omega_h(n^{2-1/h})$ edges; this almost matches the trivial $O(n^2)$ upper bound given by the n-vertex complete convex geometric graph.
Finally, we prove that there is an n-vertex convex geometric graph with n vertices and $O(n \log n)$ edges that is universal for n-vertex caterpillars.

1 Introduction

A graph G is *universal* for a class $\mathcal{H}$ of graphs if G contains every graph in $\mathcal{H}$ as a subgraph. The study of universal graphs was initiated by Rado [20] in the 1960s. Obviously, the complete graph K_n is universal for any family $\mathcal{H}$ of n-vertex graphs. Research focused on finding the minimum size (i.e., number of edges) of universal graphs for various families of sparse graphs on n vertices. Babai et al. [3] proved that if $\mathcal{H}$ is the family of all graphs with m edges, then the size of a universal graph for $\mathcal{H}$ is in $\Omega(m^2/\log^2 m)$ and $O(m^2 \log\log m/\log m)$. Alon et al. [1,2] constructed a universal graph of optimal $\Theta(n^{2-2/k})$ size for n-vertex graphs with maximum degree k.

Significantly better bounds exist for minor-closed families. Babai et al. [3] proved that there exists a universal graph with $O(n^{3/2})$ edges for n-vertex planar

* Partially supported by the MSCA-RISE project "CONNECT" No 734922, the MIUR Project "AHeAD" under PRIN 20174LF3T8, and the NSF award DMS-1800734.

graphs. For bounded-degree planar graphs, Capalbo [10] constructed universal graphs of linear size, improving an earlier bound by Bhatt et al. [5], which extends to other families with bounded bisection width. Böttcher et al. [7,8] proved that every n-vertex graph with minimum degree $\Omega(n)$ is universal for n-vertex bounded-degree planar graphs. For n-vertex trees, Chung and Graham [12,13] constructed a universal graph of size $O(n \log n)$, and showed that this bound is asymptotically optimal apart from constant factors.

In this paper, we extend the concept of universality to geometric graphs. A *geometric graph* is a graph together with a straight-line drawing in the plane in which the vertices are distinct points and the edges are straight-line segments not containing any vertex in their interiors. We investigate the problem of constructing, for a given class $\mathcal{H}$ of planar graphs, a geometric graph with few vertices and edges that is *universal for* $\mathcal{H}$, that is, it contains an *embedding* of every graph in $\mathcal{H}$. For an (abstract) graph G_1 and a geometric graph G_2, an *embedding of* G_1 *onto* G_2 is an injective graph homomorphism $\varphi : V(G_1) \to V(G_2)$ such that (i) every edge $uv \in E(G_1)$ is mapped to a line segment $\varphi(u)\varphi(v) \in E(G_2)$; and (ii) every pair of edges $u_1v_2, u_2v_2 \in E(G_1)$ is mapped to a pair of noncrossing line segments $\varphi(u_1)\varphi(v_1)$ and $\varphi(u_2)\varphi(v_2)$ in the plane.

Previous research in the geometric setting was limited to finding the smallest *complete* geometric graph that is universal for the planar graphs on n vertices. The intersection pattern of the edges of a geometric graph is determined by the location of its vertices; hence universal complete geometric graphs are commonly referred to as *n-universal point sets*. De Fraysseix et al. [16] proved that the $2n \times n$ section of the integer lattice is an n-universal point set. Over the last 30 years, the upper bound on the size of an n-universal point set has been improved from $2n^2$ to $n^2/4 + O(n)$ [4]; the current best lower bound is $(1.293 - o(1))n$ [21] (based on stacked triangulations, i.e., maximal planar graphs of treewidth three; see also [11,19]). It is known that every set of n points in general position is universal for n-vertex outerplanar graphs [6,18]. An $O(n^{3/2} \log n)$ upper bound is known for n-vertex stacked triangulations [17].

Our Results. The results on universal point sets yield an upper bound of $O(n^4)$ for the size of a geometric graph that is universal for n-vertex planar graphs and $O(n^2)$ for n-vertex outerplanar graphs, including trees. We improve the upper bound for n-vertex trees to an optimal $O(n \log n)$, and show that the quadratic upper bound for outerplanar graphs is essentially tight for convex geometric graphs. More precisely, we prove the following results:

– For every $n \in \mathbb{N}$, there is a geometric graph G with n vertices and $O(n \log n)$ edges that is universal for forests with n vertices (Theorem 1 in Sect. 2). The $O(n \log n)$ bound is asymptotically optimal, even in the abstract setting, for caterpillars (a *caterpillar* is a tree such that the removal of its leaves results in a path, called *spine*), and if the universal graph is allowed to have more than n vertices [12, Theorem 1]. The proof of universality is constructive and yields a polynomial-time algorithm that embeds any n-vertex forest onto G.

- For every $h \in \mathbb{N}$ and $n \geq 3h^2$, every n-vertex convex geometric graph that is universal for the family of n-vertex cycles with h disjoint chords has $\Omega_h(n^{2-1/h})$ edges (Theorem 2 in Sect. 3); this almost matches the trivial $O(n^2)$ bound, which hence cannot be improved by polynomial factors even for n-vertex outerplanar graphs of maximum degree three. For n-vertex cycles with 2 disjoint chords, there is an n-vertex convex geometric graph with $O(n^{3/2})$ edges (Theorem 3 in Sect. 3), which matches the lower bound above.
- For every $n \in \mathbb{N}$, a convex geometric graph with n vertices and $O(n \log n)$ edges exists that is universal for n-vertex caterpillars (Theorem 4 in Sect. 3).

A full version of the paper can be found in [14].

2 Universal Geometric Graphs for Forests

In this section, we prove the following theorem.

Theorem 1. *For every $n \in \mathbb{N}$, there exists a geometric graph G with n vertices and $O(n \log n)$ edges that is universal for forests with n vertices.*

Construction. We adapt a construction of Chung and Graham [13] to the geometric setting. For $n \in \mathbb{N}$, they construct an n-vertex graph G with $O(n \log n)$ edges that contains every n-vertex forest as a subgraph. We present this construction. For simplicity assume $n = 2^h - 1$ with $h \geq 2$. Let B be an n-vertex complete rooted ordered binary tree. A *level* is a set of vertices at the same distance from the root. The levels are labeled $1, \dots, h$, from the one of the root to the one of the leaves. A preorder traversal of B (visiting first the root, then recursively the vertices in its left subtree, and then recursively the vertices in its right subtree) determines a total order on the vertices, which also induces a total order on the vertices in each level. On each level, we call two consecutive elements in this order *level-neighbors*; in particular, siblings are level-neighbors. We denote by $B(v)$ the subtree of B rooted at a vertex v. The graph G contains B and three additional groups of edges (see Fig. 1): (E1) Every vertex v is adjacent to all vertices in $B(v)$; (E2) every vertex v with a level-neighbor u in B is adjacent to all vertices in $B(u)$; and (E3) every vertex v whose parent has a left level-neighbor p is adjacent to all vertices in $B(p)$.

Number of edges. The tree B has 2^{i-1} vertices on level i, for $i = 1, \dots, h$. A vertex v on level i has $2^{h-i+1} - 1$ descendants (including itself), and its at most two level-neighbors have the same number of descendants. In addition, the left level-neighbor of the parent of v (if present) has $2 \cdot (2^{h-i+1} - 1)$ descendants (excluding itself). Altogether v is adjacent to less than $5 \cdot 2^{h-i+1}$ vertices at the same or at lower levels of B. Hence, the number of edges in G is less than $5 \cdot \sum_{i=1}^{h} 2^{i-1} \cdot 2^{h-i+1} = 5 \cdot 2^h \cdot h = 5(n+1) \cdot \log_2(n+1) \in O(n \log n)$.

Chung and Graham [13] showed that G is universal for n-vertex forests.[5]

[5] The construction by Chung and Graham uses fewer edges: in the edge groups (E2) and (E3), they use siblings instead of level-neighbors. But we were unable to verify their proof with the smaller edge set, namely we do not see why the graph G_2 in [13, Fig. 7] is admissible. However, their proof works with the edge set we define here.

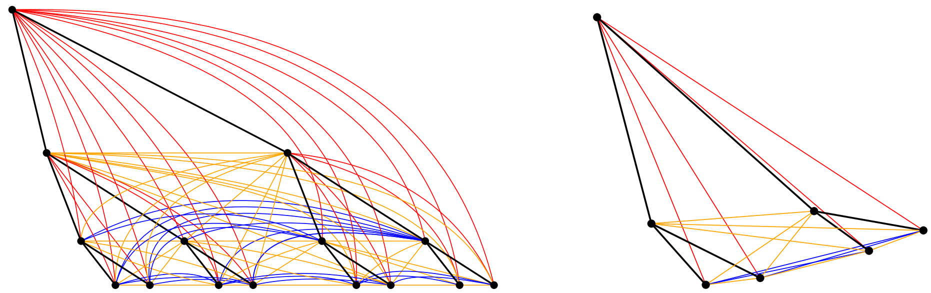

Fig. 1: A schematic drawing of the 15-vertex universal graph (left) and a geometric drawing of the 7-vertex universal graph. The edges of B are black; the edges of the groups (E1), (E2), and (E3) are red, orange, and blue, respectively. Edges in several groups have the color of the first group they belong to.

Geometric representation. We next describe how to embed the vertices of G into $\mathbb{R}^2$; see Fig. 2(left) for an illustration. First, the x-coordinates of the vertices are assigned in the order determined by a preorder traversal of B. For simplicity, let us take these x-coordinates to be $0, \dots, n-1$, so that the root of B is placed on the y-axis. The vertex of G with x-coordinate i is denoted by v_i.

The y-coordinates of the vertices are determined by a BFS traversal of B starting from the root, in which at every vertex the right sibling is visited before the left sibling. If a vertex u is visited before a vertex v by this traversal, then u gets a larger y-coordinate than v. The gap between two consecutive y-coordinates is chosen so that every vertex is above every line through two vertices with smaller y-coordinate; this implies that, for any vertex v, all vertices with larger y-coordinate than v, if any, see the vertices below v in the same circular order as v. The vertices of G are in general position, that is, no three are collinear.

Our figures display the vertices of B in the correct x- and y-order, but—with the exception of Fig. 1(right)—they are not to scale. The y-coordinates in our construction are rapidly increasing (similarly to [9,17]). For this reason, in our figures we draw the edges in B as straight-line segments and all other edges as Jordan arcs. We have the following property (refer to Fig. 2(right)).

Observation 1 *If $ab, cd \in E(G)$ are such that (1) a has larger y-coordinate than b, c, and d, and (2) b has smaller or larger x-coordinate than both c and d, then ab and cd do not cross.*

Intervals. A geometric graph is *plane* if it contains no crossings. For every interval $[i, j] \subseteq [0, n-1]$ we define $G[i, j]$ as the subgraph of G induced by the vertices with x-coordinate in $[i, j]$. Then $G[i, j]$ is an *interval* of G. The *length* of $G[i, j]$ is defined as $|G[i, j]| = j - i + 1$, which is the number of vertices in $G[i, j]$. If I is an interval of integers, then we denote by $G(I)$ the corresponding interval of G. For a subset $U \subset V(G)$, we denote by $G[U]$ the subgraph of G induced by U. We will show eventually that every tree on h vertices admits an embedding onto every interval of length h of G. We now present the following lemma.

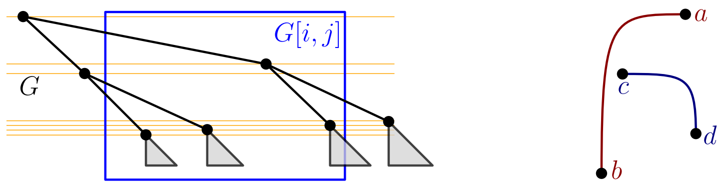

Fig. 2: Illustration for the assignment of x- and y-coordinates to the vertices of G, and for the definition of interval (left). Illustration for Observation 1(right).

Lemma 1. *Every interval $G[i,j]$ of G on at least two vertices contains two plane spanning stars, centered at the highest vertex v_k and at the second highest vertex v_s of $G[i,j]$. If $k < j$, then $G[i,j]$ contains a plane spanning star centered at the highest vertex of $G[k+1,j]$ (which may or may not be v_s).*

Our upcoming recursive algorithm sometimes embeds a subtree of T onto an induced subgraph of G that is "almost" an interval, in the sense that it can be obtained from an interval of G by deleting its highest vertex or by replacing its highest vertex with a vertex that does not belong to the interval.

We first prove that the "structure" of an interval without its highest vertex is similar to that of an interval. Let U and W be two subsets of $V(G)$ with $h = |U| = |W|$. Let $u_1, \dots, u_h$ and $w_1, \dots, w_h$ be the vertices of U and W, respectively, ordered by increasing x-coordinates. We say that $G[U]$ and $G[W]$ are *crossing-isomorphic* if: (C1) for any $p, q \in \{1, \dots, h\}$, the edge $u_p u_q$ belongs to $G[U]$ if and only if the edge $w_p w_q$ belongs to $G[W]$; (C2) for any $p, q, r, s \in \{1, \dots, h\}$ such that the edges $u_p u_q$ and $u_r u_s$ belong to $G[U]$, the edges $u_p u_q$ and $u_r u_s$ cross if and only if the edges $w_p w_q$ and $w_r w_s$ cross; and (C3) if u_i is the highest vertex of $G[U]$, for some $i \in \{1, \dots, h\}$, then w_i is the highest vertex of $G[W]$. The graph isomorphism given by $\lambda(u_i) = w_i$, for all $i = 1, \dots, n$, is a *crossing-isomorphism*. Clearly, the inverse of a crossing-isomorphism is also a crossing-isomorphism. We have the following.

Lemma 2. *Let v_k be the highest vertex in an interval $G[i,j]$, and assume that $G[i,j]$ contains neither the right child of v_k nor any descendant of the left child of its left sibling (if it exists). Then $G[i,j] - v_k$ is crossing-isomorphic to some interval $G(I)$ of G; the interval I can be computed in $O(1)$ time.*

We now present our tools for embedding trees onto "almost" intervals.

Lemma 3. *Let v_k be the highest vertex in an interval $G[i,j]$ with $h+1$ vertices. Suppose that there is a crossing-isomorphism λ from $G[i,j] - v_k$ to some interval $G(I)$ of G with h vertices. Further, suppose that a tree T with h vertices admits an embedding φ onto $G(I)$. Then $\varphi' = \lambda^{-1} \circ \varphi$ is an embedding of T onto $G[i,j] - v_k$, and if a is the vertex of T such that $\varphi(a)$ is the highest vertex of $G(I)$, then $\varphi'(a)$ is the highest vertex of $G[i,j] - v_k$.*

Lemma 4. *Let $G[i,j]$ be an interval of G with h vertices and let v_k be its highest vertex. Let v_x be a vertex of G that is higher than all vertices in $G[i,j] - v_k$ and that does not belong to $G[i,j]$. Suppose that a tree T with h vertices admits an embedding φ onto $G[i,j]$. Let a be the vertex of T such that $\varphi(a) = v_k$; further, let $\varphi'(a) = v_x$ and $\varphi'(b) = \varphi(b)$ for every vertex b of T other than a. Then φ' is an embedding of T onto $G[i,j] - v_k + v_x$.*

The following lemma is a variant of the (unique) lemma in [13].

Lemma 5. *Given a rooted tree T on $m \geq 2$ vertices and an integer s, with $1 \leq s \leq m$, there is a vertex c of T such that $|V(T(c))| \geq s$ but $|V(T(d))| \leq s-1$, for all children d of c. Such a vertex c can be computed in time $O(m)$.*

Proof strategy. Given a tree T on h vertices and an interval $G[i,j]$ of length h, we describe a recursive algorithm that constructs an embedding φ of T onto $G[i,j]$. For a subtree T' of T, we denote by $\varphi(T')$ the image of φ restricted to T'. A step of the algorithm explicitly embeds some vertices; the remaining vertices form subtrees that are recursively embedded onto pairwise disjoint subintervals of $G[i,j]$. We insist that in every subtree at most two vertices, called *portals*, are adjacent to vertices not in the subtree. We also ensure that whenever a subtree is embedded onto a subinterval, the vertices not in the subtree that connect to the portals of that subtree are embedded above the subinterval.

For a point p, we denote $Q^+(p) = \{q \in \mathbb{R}^2 : x(p) < x(q) \text{ and } y(p) < y(q)\}$ and $Q^-(p) = \{q \in \mathbb{R}^2 : x(q) < x(p) \text{ and } y(p) < y(q)\}$.

We inductively prove the following lemma, which immediately implies Theorem 1 with $G[i,j] = G[0, n-1]$ and a portal a chosen arbitrarily.

Lemma 6. *We are given a tree T on h vertices, an interval $G[i,j]$ of length h, and either (1) a single portal a in T, or (2) two distinct portals a and b in T. Then there exists an embedding φ of T onto $G[i,j]$ with the following properties:*

1. *If only one portal is given, then*
   - *(a) $\varphi(a)$ is the highest vertex in $G[i,j]$; and*
   - *(b) if $\deg_T(a) = 1$ and a' is the unique neighbor of a in T, then $Q^-(\varphi(a'))$ does not intersect any vertex or edge of the embedding $\varphi(T(a'))$.*
2. *If two distinct portals are given, then*
   - *(a) $\varphi(a)$ is to the left of $\varphi(b)$;*
   - *(b) $Q^-(\varphi(a))$ does not intersect any edge or vertex of $\varphi(T)$; and*
   - *(c) $Q^+(\varphi(b))$ does not intersect any edge or vertex of $\varphi(T)$.*

Proof sketch: We proceed by induction on h. In the base case $h = 1$, hence T has one vertex, which must be the portal a, and the map $\varphi(a) = v_i$ maps a to the highest vertex of $G[i,i]$. For the induction step we assume that $h \geq 2$.

Case 1: There is only one portal a. Let v_k denote the highest vertex in $G[i,j]$. We need to find an embedding of T onto $G[i,j]$ where $\varphi(a) = v_k$. Consider T to be rooted at a. We distinguish two cases depending on the degree of a in T.

Case 1.1: $\deg_T(a) \geq 2$. Assume that a has t children $a_1, \ldots, a_t$. Refer to Fig. 3. Partition the integers $[i,j] \setminus \{k\}$ into t contiguous subsets $I_1, \ldots, I_t$ such

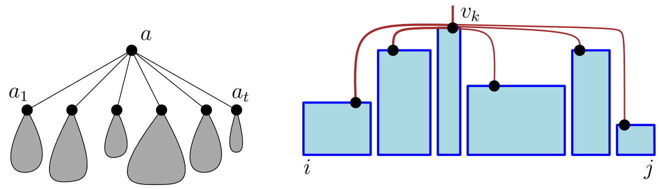

Fig. 3: Case 1.1: Tree T (left) and its embedding onto $G[i, j]$ (right).

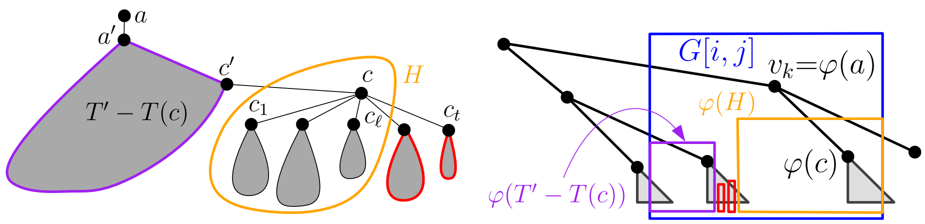

Fig. 4: Case 1.2.4. Tree T (left) and its embedding onto $G[i, j]$ (right).

that $|I_x| = |V(T(a_x))|$, for $x = 1, \ldots, t$. W.l.o.g. assume that I_q contains $k-1$ or $k+1$, and so $I_q \cup \{k\}$ is an interval of integers. By induction, there is an embedding φ_x of $T(a_x)$ onto $G(I_x)$ such that $\varphi_x(a_x)$ is the highest vertex of $G(I_x)$, for all $x \neq q$, and there is an embedding φ_q of $T - \bigcup_{x \neq q} T(a_x)$ onto $G(I_q \cup \{k\})$ such that $\varphi_q(a) = v_k$. Then the combination of these embeddings is an embedding φ of T onto $G[i, j]$ satisfying Properties 1(a) and 1(b). In particular, the edges $\varphi(a)\varphi(a_x)$ are in $G[i, j]$ by Lemma 1, and an edge $\varphi(a)\varphi(a_x)$ does not cross $\varphi(T(a_y))$, where $y \neq x$, by Observation 1. ◀

Case 1.2: $\deg_T(a) = 1$. Let a' be the neighbor of a in T and let $T' = T(a')$.

Case 1.2.1: $k = j$. Set $\varphi(a) = v_k$ and recursively embed T' onto $G[i, k-1]$ with a single portal a', which is mapped to the highest vertex in $G[i, k-1]$. Clearly, φ is an embedding of T onto $G[i, j]$ satisfying Properties 1(a) and 1(b), since the edge $\varphi(a)\varphi(a')$ is above, and hence does not cross, $\varphi(T')$. ◀

Case 1.2.2: $k = i$. This case is symmetric to Case 1.2.1. ◀

Case 1.2.3: $i < k < j$ and the left sibling v_ℓ of v_k exists and is in $G[i, j]$. It follows that $\ell = i$, as if $\ell > i$, then $v_{\ell-1}$, which is the parent of v_ℓ and v_k, would be a vertex in $G[i, j]$ higher than v_k. By construction, v_i is the second highest vertex in $G[i, j]$. Recursively construct an embedding ψ of T' onto $G[i+1, j]$ with a single portal a'. By Property 1(a), we have $\psi(a') = v_k$. By Lemma 4, there exists an embedding φ of T' onto $G[i+1, j] - v_k + v_i = G[i, j] - v_k$ in which $\varphi(a') = v_i$ (hence φ satisfies Property 1(b)). Finally, set $\varphi(a) = v_k$ (hence φ satisfies Property 1(a)). As in Case 1.2.1, the edge $\varphi(a)\varphi(a') = v_k v_i$ does not cross $\varphi(T')$, hence φ is an embedding of T onto $G[i, j]$. ◀

Case 1.2.4: $i < k < j$, the left sibling of v_k does not exist or is not in $G[i, j]$, and the right child of v_k is not in $G[i, j]$. Refer to Fig. 4. By construction, the left child of v_k is v_{k+1}, which is in $G[i, j]$. By the assumptions of this case, v_{k+1} is the second highest vertex in $G[i, j]$. Set $s = j-k+1$; then $s < h$, given that $k > i$. By

Lemma 5, there is a vertex c in T' such that $|V(T'(c))| \geq s$ but $|V(T'(d))| \leq s-1$ for all children d of c. Label the children of c as $c_1, \ldots, c_t$ in an arbitrary order and let $\ell \in [1, t]$ be the smallest index such that $1 + \sum_{x=1}^{\ell} |V(T(c_i))| \geq s$. Since $|V(T(c_\ell))| \leq s-1$, we have $s \leq 1 + \sum_{x=1}^{\ell} |V(T(c_i))| \leq 2s-2$.

Let c' be the parent of c in T. Let H denote the subtree of T induced by c and by $V(T(c_1)), \ldots, V(T(c_\ell))$, and let $m = |V(H)|$. By the above inequalities, we have $s \leq m \leq 2s-2$. On the one hand, $j-k+1 \leq m$ implies that the subinterval $G[j-m, j]$ contains v_k, and so v_k is the highest vertex in $G[j-m, j]$. On the other hand, the interval $G[j-m, k-1]$ contains $k-1-j+m+1 = m-s+1 \leq s-1$ vertices, given that $m \leq 2s-2$; however, since the right child of v_k is not in $G[i, j]$, we know that the size of a subtree of B rooted at any vertex at the level below v_k is larger than or equal to $s-1$. It follows that $G[j-m, j]$ does not contain any descendants of the left child of the left sibling of v_k (if it exists). By Lemma 2, $G[j-m, j] - v_k$ is crossing-isomorphic to an interval $G(I)$ of size m.

Recursively embed H onto $G(I)$ with one portal c. By Lemma 3, there exists an embedding φ of H onto $G[j-m, j] - v_k$ such that $\varphi(c) = v_{k+1}$. Set $\varphi(a) = v_k$. If c has more than ℓ children, then embed the subtrees $T(c_{\ell+1}), \ldots, T(c_t)$ onto subintervals to the left of $G[j-m, j]$, with single portals $c_{\ell+1}, \ldots, c_t$, respectively. Finally, by induction, we can embed $T' - T(c)$ onto the remaining subinterval of $G[i, j]$ with two portals a' and c'. Then φ is an embedding of T onto $G[i, j]$ satisfying Properties 1(a) and 1(b). In particular, the edge $\varphi(c)\varphi(c')$ does not cross $\varphi(T' - T(c))$, since this satisfies Property 2(c) (note that $\varphi(c)$ is in $Q^+(\varphi(c'))$).◀

Case 1.2.5: $i < k < j$, the left sibling of v_k does not exist or is not in $G[i, j]$, and the right child v_r of v_k is in $G[i, j]$. By assumption, we have $k+1 < r \leq j$; further, the second highest vertex in $G[i, j]$ is v_r. Set $s = j - r + 1$. Lemma 5 yields a vertex c in T' such that $|V(T(c))| \geq s$ but $|V(T(d))| \leq s-1$ for all children d of c. Let $T(c)$ be the subtree of T rooted at c, set $m = |V(T(c))|$, and label the children of c by $c_1, \ldots, c_t$ in an arbitrary order. Let c' be the parent of c and denote by $T_c(c')$ the subtree of T induced by c' and $V(T(c))$.

Case 1.2.5.1: $m \leq j-k-1$. Then the interval $[j-m, j]$ contains r but does not contain k, hence v_r is the highest vertex in $G[j-m, m]$. By induction, there is an embedding ψ_1 of $T' - T(c)$ onto $G[i, j-m-1]$ with two portals a' and c'. By Lemma 4, there is an embedding φ of $T' - T(c)$ onto $G[i, j-m-1] - v_k + v_r$. Set $\varphi(a) = v_k$. Again by induction, there is an embedding ψ_2 of $T_c(c')$ onto $G[j-m, j]$ with a single portal c'. Let $\varphi(T(c)) = \psi_2(T(c))$ and note that $\varphi(c')$ may be different from $\psi_2(c') = v_r$. This completes the definition of $\varphi(T)$, which is an embedding of T onto $G[i, j]$ satisfying Properties 1(a) and 1(b). In particular, we argue that the edge $\varphi(c)\varphi(c')$ is in $G[i, j]$. Let $v_p = \varphi(c')$ and $v_q = \varphi(c)$, and note that $p < q$ or $p = r$. In the latter case, $\varphi(c)\varphi(c')$ exists as $\psi_2(c') = v_r$ and the edge cc' belongs to $T_c(c')$; hence, assume that $p < q$. By Property 2(c) of ψ_1, we have that $Q^+(\psi_1(c'))$ does not intersect $\psi_1(T' - T(c))$, hence $k < p$, as otherwise v_k would be in $Q^+(\psi_1(c'))$; hence, v_p is the highest vertex in $G[p, j-m-1]$. Further, by Property 1(b) of ψ_2, we have that $Q^-(\psi_2(c))$ does not intersect $\psi_2(T(c))$; hence, v_q is either the highest or the second highest vertex in $G[j-m, q]$ (as v_r might belong to such an interval). Overall, one of v_p or v_q is the highest or

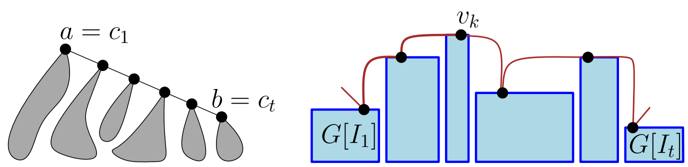

Fig. 5: Case 2. Tree T (left) and its embedding onto $G[i, j]$ (right).

the second highest vertex in $G[p, q]$. By Lemma 1, $G[p, q]$ contains a star centered at v_p or v_q, and so it contains the edge $v_p v_q = \varphi(c')\varphi(c)$.

Case 1.2.5.2: $j - k - 1 < m$. Then $[j - m, j]$ contains both k and r. Partition $[j - m, j] \setminus \{k, r\}$ into t contiguous subsets $I_1, \ldots, I_t$ such that $|I_x| = |V(T(c_x))|$, for $x = 1, \ldots, t$. W.l.o.g. assume that I_q contains $r - 1$. Let $\mathcal{I}(c)$ be the collection of the sets $I_q \cup \{v_r\}$ and I_x, for $x \in [1, t] \setminus \{q\}$. At least $t - 1$ of these sets are intervals, and at most one of them, say I_p, is an interval minus its highest element. Since every tree $T(c_i)$ has at most $s - 1$ vertices, I_p has at most s elements. Since $s = j - r - 1$, we have $|I_p| \leq |V(B(v_r))|$, hence I_p contains neither the right child of v_k nor any descendant of its left sibling. By Lemma 2, $G(I_p)$ is crossing-isomorphic to an interval. By Lemma 3, we can embed $T(c_p)$ onto $G(I_p)$. We also recursively embed $T(c_x)$ onto $G[I_x]$ for all $x \in [1, t] \setminus \{p, q\}$ and $T(c) - \bigcup_{x \neq q} T_x$ onto $G(I_q \cup \{u\})$. Embed a at v_k. Finally, embed $T' - T(c)$ onto $G[i, j - m - 1]$ with portals a' and c'. The combination of these embeddings is an embedding φ of T onto $G[i, j]$ satisfying Properties 1(a) and 1(b). ◀

Case 2: Two portals a and b; refer to Fig. 5. Let $P = (a = c_1, \ldots, c_t = b)$ be the path between a and b in T, where $t \geq 2$. The deletion of the edges in P splits T into t trees rooted at $c_1, \ldots, c_t$. Partition $[i, j]$ into t subintervals $I_1, \ldots, I_t$ such that $|I_x| = |V(T(c_x))|$, for $x = 1, \ldots, t$. For $x = 1, \ldots, t$, recursively embed $T(c_x)$ onto $G(I_x)$ with one portal c_x. The combination of these embeddings is an embedding φ of T onto $G[i, j]$ satisfying Properties 2(a), 2(b), and 2(c). □

3 Convex Geometric Graphs

Every graph embedded onto a *convex* geometric graph is outerplanar. Clearly, an n-vertex complete convex geometric graph has $O(n^2)$ edges and is universal for the n-vertex outerplanar graphs. We show that this trivial bound is almost tight. For $h \geq 0$ and $n \geq 2h + 2$, let $\mathcal{O}_h(n)$ be the family of n-vertex outerplanar graphs consisting of a spanning cycle plus h pairwise disjoint chords.

Theorem 2. *For every positive integer h and $n \geq 3h^2$, every convex geometric graph C on n vertices that is universal for $\mathcal{O}_h(n)$ has $\Omega_h(n^{2-1/h})$ edges.*

Proof: Denote by ∂C the outer (spanning) cycle of C. The *length* of a chord uv of ∂C is the length of a shortest path between u and v along ∂C. For $k \geq 2$, denote by E_k the set of length-k chords in C, and let $m \in \{2, \ldots, \lfloor n/(3h) \rfloor\}$ be an integer such that $|E_m| = \min\{|E_2|, \ldots, |E_{\lfloor n/(3h) \rfloor}|\}$.

Let $\mathcal{L}$ be the set of *labeled* n-vertex outerplanar graphs that consist of a spanning cycle $(v_0, \dots, v_{n-1})$ plus h pairwise-disjoint chords of length m such that one chord is $v_0 v_m$ and all h chords have both vertices on the path $P = (v_0, \dots, v_{\lfloor n/3 \rfloor + hm-1})$. Every graph $G \in \mathcal{L}$ has a unique spanning cycle H, which is embedded onto ∂C. Since they all have the same length, the h chords of H have a well-defined cyclic order along H. A *gap* of G is a path between two consecutive chords along H. Note that G has h gaps. The length of P is $\lfloor n/3 \rfloor + hm - 1 \le \lfloor n/3 \rfloor + h \cdot \lfloor n/(3h) \rfloor - 1 < 2\lfloor n/3 \rfloor$, hence the length of the gap between the last and the first chords is more than $n - 2\lfloor n/3 \rfloor = \lceil n/3 \rceil$. This is the longest gap, as the lengths of the other gaps sum up to at most $\lfloor n/3 \rfloor$.

Let $\mathcal{U}$ denote the subset of *unlabeled* graphs in $\mathcal{O}_h(n)$ that correspond to some labeled graph in $\mathcal{L}$. Each graph in $\mathcal{L}$ is determined by the lengths of its $h-1$ shortest gaps. The sum of these lengths is an integer between $h-1$ and $(\lfloor n/3 \rfloor + hm - 1) - hm < \lfloor n/3 \rfloor$. The number of compositions of $\lfloor n/3 \rfloor$ into h positive integers (i.e., $h-1$ lengths and a remainder) is $\binom{\lfloor n/3 \rfloor}{h-1} \in \Theta_h(n^{h-1})$. Each unlabeled graph in $\mathcal{U}$ corresponds to at most two labeled graphs in $\mathcal{L}$, since any graph automorphism setwise fixes the unique spanning cycle as well as the longest gap. Hence, $|\mathcal{U}| \in \Theta(|\mathcal{L}|) \subseteq \Theta_h(n^{h-1})$.

Since C is universal for $\mathcal{O}_h(n)$ and $\mathcal{U} \subset \mathcal{O}_h(n)$, every graph G in $\mathcal{U}$ embeds onto C. Since every embedding of G maps the spanning cycle of G onto ∂C and the h chords of G into a subset of E_m, we have that C contains at most $\binom{|E_m|}{h} \le |E_m|^h$ graphs in $\mathcal{U}$. The combination of the lower and upper bounds for $|\mathcal{U}|$ yields $|E_m|^h \in \Omega_h(n^{h-1})$, hence $|E_m| \in \Omega_h(n^{1-1/h})$. Overall, the number of edges in C is at least $\sum_{i=1}^{\lfloor n/(3h) \rfloor} |E_i| \ge \lfloor n/(3h) \rfloor \cdot |E_m| \in \Omega_h(n^{2-1/h})$. □

For the case $h = 2$, the lower bound of Theorem 2 is the best possible.

Theorem 3. *For every $n \in \mathbb{N}$, there exists a convex geometric graph C with n vertices and $O(n^{3/2})$ edges that is universal for $\mathcal{O}_2(n)$.*

Proof: The vertices $v_0, \dots, v_{n-1}$ of C form a convex n-gon and the edges of this spanning cycle are in C. Let $S = \{0, \dots, \lfloor \sqrt{n} \rfloor - 1\} \cup \{i \lfloor \sqrt{n} \rfloor : 1 \le i \le \lfloor \sqrt{n} \rfloor\}$ and add a star centered at v_s, for every $s \in S$, to C. Clearly, C contains $O(n^{3/2})$ edges. Moreover, for every $d \in \{1, \dots, \lfloor n/2 \rfloor\}$ there exist $a, b \in S$ so that $b - a = d$. For any $G \in \mathcal{O}_2(n)$, let $a, b \in S$ so that the distance along the outer cycle between the two closest vertices of the two chords of G is $b - a$. As C contains stars centered at both v_a and v_b, the graph G embeds onto C. □

We next construct a convex geometric graph G with n vertices and $O(n \log n)$ edges that is universal for n-vertex caterpillars; this bound is asymptotically optimal [12]. In order to construct G, we define a sequence π_n of n integers. Let $\pi_1 = (1)$. For every integer m of the form $m = 2^h - 1$, where $h \ge 2$, let $\pi_m = \pi_{(m-1)/2}(m)\pi_{(m-1)/2}$. For any $n \in \mathbb{N}$, the sequence π_n consists of the first n integers in π_m, where $m \ge n$ and $m = 2^h - 1$, for some $h \ge 1$. For example, $\pi_{10} = (1, 3, 1, 7, 1, 3, 1, 15, 1, 3)$. Let $\pi_n(i)$ be the ith term of π_n.

Property 1 ([15]). For every $n \in \mathbb{N}$ and for every x with $1 \le x \le n$, the maximum of any x consecutive elements in π_n is at least x.

The graph G has vertices $v_1, \dots, v_n$, placed in counterclockwise order along a circle c. Further, for $i = 1, \dots, n$, we have that G contains edges connecting v_i to the $\pi_n(i)$ vertices preceding v_i and to the $\pi_n(i)$ vertices following v_i along c.

Theorem 4. *For every $n \in \mathbb{N}$, there exists a convex geometric graph G with n vertices and $O(n \log n)$ edges that is universal for n-vertex caterpillars.*

Proof sketch: The number of edges of G is at most twice the sum of the integers in π_n; the latter is less than or equal to the sum of the integers in π_m, where $m < 2n$ and $m = 2^h - 1$, for some integer $h \geq 1$. Further, π_m is easily shown to be equal to $(h-1) \cdot 2^h + 1 \in O(n \log n)$.

Let C be an n-vertex caterpillar and let $(u_1, \dots, u_s)$ be the spine of C. For $i = 1, \dots, s$, let S_i be the star composed of u_i and its adjacent leaves; let $n_i = |V(S_i)|$. Let $m_1 = 0$; for $i = 2, \dots, s$, let $m_i = \sum_{j=1}^{i-1} n_j$. For $i = 1, \dots, s$, we embed S_i onto the subgraph G_i of G induced by the vertices $v_{m_i+1}, v_{m_i+2}, \dots, v_{m_i+n_i}$: This is done by embedding u_i at the vertex v_{x_i} of G_i whose degree (in G) is maximum, and by embedding the leaves of S_i at the remaining vertices of G_i. By Property 1, we have that v_{x_i} is adjacent in G to the n_i vertices preceding it and to the n_i vertices following it along c, hence it is adjacent to all other vertices of G_i; thus, the above embedding of S_i onto G_i is valid. The arguments showing that the edge $v_{x_i}v_{x_{i+1}}$ belongs to G for all $i = 1, \dots, s-1$ are analogous. The proof is concluded by observing that the edges of the spine $(u_1, \dots, u_s)$ do not cross each other, since the vertices $u_1, \dots, u_s$ appear in this order along c. $\square$

4 Conclusions and Open Problems

In this paper we introduced and studied the problem of constructing geometric graphs with few vertices and edges that are universal for families of planar graphs. Our research raises several challenging problems.

What is the minimum number of edges of an n-vertex convex geometric graph that is universal for n-vertex trees? We proved that the answer is in $O(n \log n)$ if convexity is not required, or if caterpillars, rather than trees, are considered, while it is close to $\Omega(n^2)$ if outerplanar graphs, rather than trees, are considered.

What is the minimum number of edges in a geometric graph that is universal for all n-vertex planar graphs? For abstract graphs, Babai et al. [3] constructed a universal graph with $O(n^{3/2})$ edges based on separators. Can such a construction be adapted to a geometric setting? The current best lower bound is $\Omega(n \log n)$, same as for trees [13], while the best upper bound is only $O(n^4)$.

Finally, the problems we studied in this paper can be posed for *topological* (multi-)graphs, as well, in which edges are represented by Jordan arcs.

Theorem 5. *For every $n \in \mathbb{N}$, there is a topological multigraph with n vertices and $O(n^3)$ edges that contains a planar drawing of every n-vertex planar graph.*

Theorem 6. *For every $n \in \mathbb{N}$, there is a topological multigraph with n vertices and $O(n^2)$ edges that contains a planar drawing of every n-vertex subhamiltonian planar graph.*

References

1. Alon, N., Capalbo, M.R.: Optimal universal graphs with deterministic embedding. In: Proc. 19th ACM-SIAM Sympos. Discrete Algorithms (SODA). pp. 373–378 (2008), https://dl.acm.org/doi/abs/10.5555/1347082.1347123
2. Alon, N., Capalbo, M.R., Kohayakawa, Y., Rödl, V., Rucinski, A., Szemerédi, E.: Universality and tolerance. In: Proc. 41st IEEE Sympos. Foundations of Computer Science (FOCS). pp. 14–21 (2000). https://doi.org/10.1109/SFCS.2000.892007
3. Babai, L., Chung, F.R.K., Erdős, P., Graham, R.L., Spencer, J.H.: On graphs which contain all sparse graphs. Annals Discrete Math. **12**, 21–26 (1982). https://doi.org/10.1016/S0304-0208(08)73486-8
4. Bannister, M.J., Cheng, Z., Devanny, W.E., Eppstein, D.: Superpatterns and universal point sets. J. Graph Algorithms Appl. **18**(2), 177–209 (2014). https://doi.org/10.7155/jgaa.00318
5. Bhatt, S.N., Chung, F.R.K., Leighton, F.T., Rosenberg, A.L.: Universal graphs for bounded-degree trees and planar graphs. SIAM J. Discrete Math. **2**(2), 145–155 (1989). https://doi.org/10.1137/0402014
6. Bose, P.: On embedding an outer-planar graph in a point set. Comput. Geom. **23**(3), 303–312 (2002). https://doi.org/10.1016/S0925-7721(01)00069-4
7. Böttcher, J., Pruessmann, K.P., Taraz, A., Würfl, A.: Bandwidth, expansion, treewidth, separators and universality for bounded-degree graphs. European Journal of Combinatorics **31**(5), 1217–1227 (2010). https://doi.org/10.1016/j.ejc.2009.10.010
8. Böttcher, J., Schacht, M., Taraz, A.: Proof of the bandwidth conjecture of Bollobás and Komlós. Mathematische Annalen **343**(1), 175–205 (2009). https://doi.org/10.1007/s00208-008-0268-6
9. Bukh, B., Matoušek, J., Nivasch, G.: Lower bounds for weak epsilon-nets and stair-convexity. Israel Journal of Mathematics **182**, 199–228 (2011). https://doi.org/10.1007/s11856-011-0029-1
10. Capalbo, M.R.: Small universal graphs for bounded-degree planar graphs. Combinatorica **22**(3), 345–359 (2002). https://doi.org/10.1007/s004930200017
11. Cardinal, J., Hoffmann, M., Kusters, V.: On universal point sets for planar graphs. J. Graph Algorithms Appl. **19**(1), 529–547 (2015). https://doi.org/10.7155/jgaa.00374
12. Chung, F.R.K., Graham, R.L.: On graphs which contain all small trees. J. Combin. Theory Ser. B **24**(1), 14–23 (1978). https://doi.org/10.1016/0095-8956(78)90072-2
13. Chung, F.R.K., Graham, R.L.: On universal graphs for spanning trees. J. London Math. Soc. **27**(2), 203–211 (1983). https://doi.org/10.1112/jlms/s2-27.2.203
14. Frati, F., Hoffmann, M., Tóth, C.D.: Universal geometric graphs. CoRR **abs/2006.11262** (2020), https://arxiv.org/abs/2006.11262
15. Frati, F., Patrignani, M., Roselli, V.: LR-drawings of ordered rooted binary trees and near-linear area drawings of outerplanar graphs. J. Comput. Syst. Sci. **107**, 28–53 (2020). https://doi.org/10.1016/j.jcss.2019.08.001
16. de Fraysseix, H., Pach, J., Pollack, R.: How to draw a planar graph on a grid. Combinatorica **10**(1), 41–51 (1990). https://doi.org/10.1007/BF02122694
17. Fulek, R., Tóth, C.D.: Universal point sets for planar three-trees. J. Discrete Algorithms **30**, 101–112 (2015). https://doi.org/10.1016/j.jda.2014.12.005
18. Gritzmann, P., Mohar, B., Pach, J., Pollack, R.: Embedding a planar triangulation with vertices at specified points. Amer. Math. Monthly **98**(2), 165–166 (1991)

19. Kurowski, M.: A 1.235 lower bound on the number of points needed to draw all n-vertex planar graphs. Inf. Process. Lett. **92**(2), 95–98 (2004). https://doi.org/10.1016/j.ipl.2004.06.009
20. Rado, R.: Universal graphs and universal functions. Acta Arithmetica **9**(4), 331–340 (1964). https://doi.org/10.4064/aa-9-4-331-340
21. Scheucher, M., Schrezenmaier, H., Steiner, R.: A note on universal point sets for planar graphs. J. Graph Algorithms Appl. **24**(3), 247–267 (2020). https://doi.org/10.7155/jgaa.00529