Topological graphs with no large grids

János Pach[1], Rom Pinchasi[2], Micha Sharir[3], Géza Tóth[4] *

[1] City College, CUNY and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; pach@cims.nyu.edu

[2] Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA; room@math.mit.edu

[3] School of Computer Science, Tel Aviv University, Tel Aviv 69 978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; michas@tau.ac.il

[4] Rényi Institute, Hungarian Academy of Sciences, Budapest, HUNGARY; geza@renyi.hu

Abstract. Let G be a *topological graph* with n vertices, i.e., a graph drawn in the plane with edges drawn as simple Jordan curves. It is shown that, for any constants k, l, there exists another constant $C(k,l)$, such that if G has at least $C(k,l)n$ edges, then it contains a $k \times l$-gridlike configuration, that is, it contains $k + l$ edges such that each of the first k edges crosses each of the last l edges. Moreover, one can require the first k edges to be incident to the same vertex.

Key words. Insert your keywords here.

1. Introduction

A *topological graph* is a graph drawn in the plane with no loops or parallel edges so that its vertices are represented by points, and its edges by Jordan curves connecting the corresponding points. The points (resp., curves) representing the vertices (resp., edges) of the abstract graph are also called vertices (resp., edges) of the resulting topological graph. If it is clear from the context, we will make no notational distinction between the vertices (edges) of the underlying abstract graph, and the points (curves) representing them in the plane. We assume that (i) the edges of a topological graph do not pass through any vertex, (ii) if two edges share an interior point, then at this point they properly cross, and (iii) no three edges cross at the same point. A topological graph is called *simple* if any pair of its edges have at most one point in common (either a common endpoint or a crossing).

It follows from Euler's Polyhedral Formula that every planar graph with n vertices has at most $3n - 6$ edges. Equivalently, every topological graph G with n vertices and more than $3n - 6$ edges has a pair of crossing edges. This simple statement was generalized in several directions.

* *János Pach and Micha Sharir has been supported by NSF Grants CCR-97-32101 and CCR-00-98246, and by a joint grant from the U.S.–Israel Binational Science Foundation. János Pach has also been supported by PSC-CUNY Research Award 63382-0032 and by OTKA T-032452. Micha Sharir has also been supported by a grant from the Israeli Academy of Sciences for a Center of Excellence in Geometric Computing at Tel Aviv University, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. Géza Tóth has been supported by OTKA-T-038397 and by an award from the New York University Research Challenge Fund.*

Pach and Tóth [PT97] proved that a *simple* topological graph of n vertices and more than $(r+2)(n-3)$ edges must have r edges that cross the same edge. This bound is tight for $r = 1, 2, 3$, but can be substantially improved for large values of r.

Agarwal et al. [AAPPS97] proved that for some $C > 0$, every *simple* topological graph with n vertices and more than Cn edges has *three* pairwise crossing edges. In [PRT02], it was shown, by a much shorter argument, that the same assertion is true for all (not necessarily simple) topological graphs. In [PRT03], this result was further strengthened: for every integer $r > 0$, there exists a constant $C_r > 0$, such that every topological graph with n vertices and more than $C_r n$ edges has $r + 2$ edges such that the first two cross each other and both of them cross the remaining r edges.

In this note, we establish another generalization. A set of $k + l$ edges of a topological graph is said to form a (k, l)-grid if each of the first k edges crosses all of the remaining l edges. If, in addition, the first k edges are incident to the same vertex, then the (k, l)-grid is called *radial.* If the last l edges are also incident to a vertex, then the grid is *biradial.*

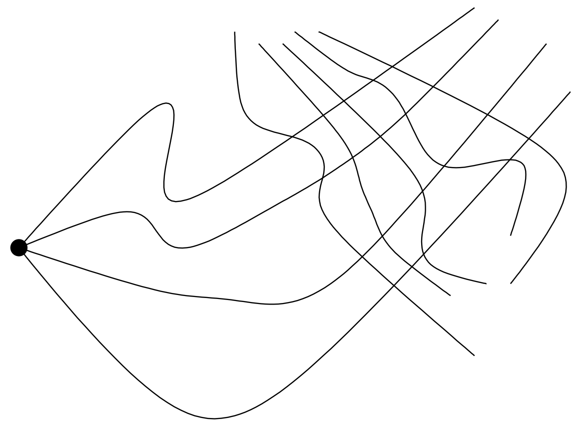

Figure 1. *A radial* $(4, 5)$*-grid.*

Theorem 1. *For any* $k, l \geq 1$*, every topological graph with* n *vertices and more than* $16 \cdot 24^{4^l} kn$ *edges contains a radial* (k, l)*-grid.*

If we assume that G is a *simple topological graph*, the proof becomes simpler and the bound becomes better.

Theorem 2. *For any* $k, l \geq 1$*, every simple topological graph of* n *vertices and more than* $8 \cdot 24^l kn$ *edges contains a radial* (k, l)*-grid.*

In the special case when the vertices of the graph are in convex position, the edges are drawn by straight-line segments, $k = 2$, and $l = 1$, Theorem 2 has been established by Brass et al. [BKV03].

It follows from the results in [PPTT02] that, for any fixed k and l, the maximum number of edges of an x-monotone topological graph with n vertices that does not contain any *biradial* (k, l)-grid is $O(n \log n)$. We do not know whether this bound can be replaced by $O(n)$. We cannot decide either if, instead of requiring that the grid be radial or biradial, one can assume that all endpoints of the participating edges are distinct.

It is an easy corollary of Theorem 2 that for any $k, l \geq 1$, there exists a constant $c = c_{kl}$ such that every simple topological graph with n vertices and more than cn edges contains a radial (k, l)-grid whose first k edges cross the remaining edges in the same order. This means that their arrangement really looks like a grid with quadrilateral cells.

2. Proof of Theorem 1

Let G be a topological graph with $n > 5$ vertices, containing no radial (k,l)-grid for some $k, l > 0$. Assume without loss of generality that G is connected (as an abstract graph), otherwise we can argue separately for its connected components and complete the proof by induction.

Let us redraw G, if necessary, so that the resulting topological graph $\tilde{G}$ satisfies the following two conditions:

(i) If two edges of $\tilde{G}$ cross each other, then the corresponding edges also cross in G;
(ii) $\tilde{G}$ has the minimum number of crossing points among all drawings with property (i).

It follows from property (i) that $\tilde{G}$ cannot contain a radial (k,l)-grid, because the corresponding edges would form a radial (k,l)-grid in G.

Obviously, no edge of $\tilde{G}$ intersects itself, otherwise we could reduce the number of crossings by removing the loop. Suppose that $\tilde{G}$ has two distinct edges, e and f, that meet at least twice (including their common endpoint, in the case they have one). A simply connected region whose boundary is composed of an arc of e and an arc of f is called a *lens.*

Claim 2.1. *Every lens in $\tilde{G}$ has a vertex in its interior.*

Proof. Suppose, for a contradiction, that there is a lens ℓ that contains no vertex of $\tilde{G}$ in its interior. Consider a *minimal* lens $\ell' \subseteq \ell$, by containment. Notice that by swapping the two sides of ℓ', we could reduce the number of crossings without creating any new pair of crossing edges, contradicting property (ii) above. □

The following property is a direct consequence of a result of Schaefer and Stefankovič [SS01].

Claim 2.2. *For any edge e of $\tilde{G}$ and for any $m > 0$, every set of 2^m consecutive crossings along e involves at least m distinct edges other than e.* □

Since the abstract underlying graph of G and $\tilde{G}$ is connected, we can choose a sequence of edges $e_1, e_2, \ldots, e_{n-1} \in E(\tilde{G})$ such that $e_1, e_2, \ldots, e_i$ form a tree T_i, for every $1 \le i \le n-1$. In particular, $e_1, e_2, \ldots, e_{n-1}$ form a spanning tree T_{n-1} of $\tilde{G}$.

As in [PRT02], we first construct a sequence of crossing-free topological graphs (trees), $\tilde{T}_1$, $\tilde{T}_2$, $\ldots$, $\tilde{T}_{n-1}$, as follows. Let $\tilde{T}_1$ be defined as a topological graph of two vertices, consisting of the single edge e_1. Suppose that $\tilde{T}_i$ has already been defined for some $1 \le i < n-1$, and let v denote the endpoint of e_{i+1} that does not belong to T_i. Then we define $\tilde{T}_{i+1}$ as follows. Add to $\tilde{T}_i$ the piece of e_{i+1} between v and its first crossing with $\tilde{T}_i$. More precisely, follow the edge e_{i+1} from v up to the point v' where it hits $\tilde{T}_i$ for the first time, and denote this piece of e_{i+1} by $\tilde{e}_{i+1}$. If v' is a vertex of $\tilde{T}_i$, then add v and $\tilde{e}_{i+1}$ to $\tilde{T}_i$ and let $\tilde{T}_{i+1}$ be the resulting topological graph. If v' is in the interior of an edge e of $\tilde{T}_i$, then introduce a new vertex at v'. It divides e into two edges, e' and e''. Add both of them to $\tilde{T}_i$, and delete e. Also add v and $\tilde{e}_{i+1}$, and let $\tilde{T}_{i+1}$ be the resulting topological graph.

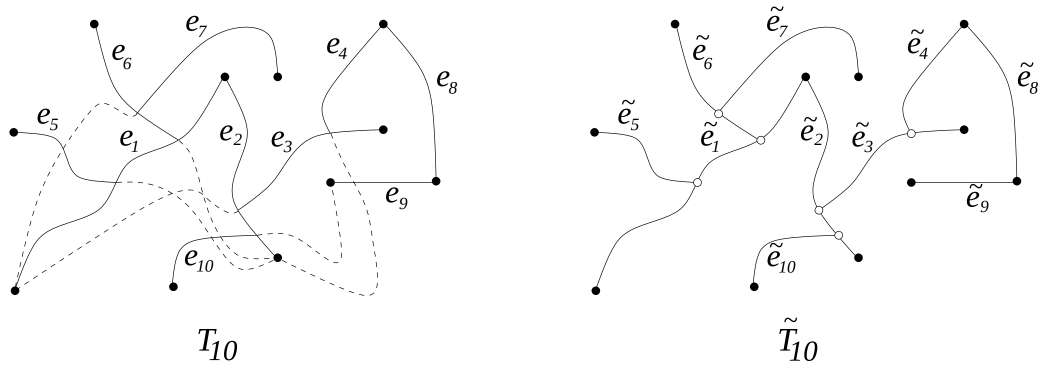

Figure 2. *Constructing* $\tilde{T}_{10}$ *from* T_{10}

After completing $n-2$ steps, we obtain a topological tree $\tilde{T} := \tilde{T}_{n-1}$, which (1) is crossing-free, (2) has fewer than $2n$ vertices, (3) contains each vertex of $\tilde{G}$, and (4) has the property that each of its edges is either a *full edge*, or a *piece of an edge* of $\tilde{G}$.

Next, we recursively construct another sequence of connected, crossing-free topological graphs, $\tilde{H}_1, \tilde{H}_2, \ldots \tilde{H}_m = \tilde{H}$. For every i, let $\tilde{E}_i$ consist of all edges of $\tilde{G}$ that have a piece which is an edge of $\tilde{H}_i$. We will refer to these edges as *used edges* at stage i.

Set $\tilde{H}_1 := \tilde{T}$ and $\tilde{E}_1 := \{e_1, e_2, \ldots, e_{n-1}\}$. Suppose that we have already constructed $\tilde{H}_i$ and $\tilde{E}_i$ for some $i > 0$.

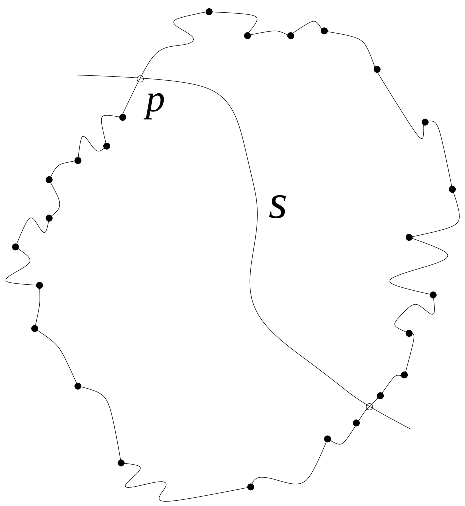

Figure 3. *A proper cut.*

Let E_i be the set of *unused edges* of $\tilde{G}$ at this stage, that is, $E_i = E(\tilde{G}) \setminus \tilde{E}_i$. Every edge $e \in E_i$ may cross $\tilde{H}_i$ at several points. These crossing points divide e into several pieces, called *segments*. Let S denote the set of all segments over all edges $e \in E_i$. To form $\tilde{H}_{i+1}$, we add one segment $s \in S$ to $\tilde{H}_i$ as an edge. Each endpoint of s is either a vertex of $\tilde{H}_i$, or a point p on an edge e of $\tilde{H}_i$. In the latter case, p divides e into two segments, e' and e''. Add p to $\tilde{H}_i$ as a vertex, and replace e with two edges, e' and e''. Denote the resulting topological graph by $\tilde{H}_i(s)$. The segment s divides one of the cells of $\tilde{H}_i$ into two cells. If each of these cells has at least 8 sides, including those along s, then we say that s is a *proper cut* of $\tilde{H}_i$ (see Figure 3). Note that if we encounter both sides of an edge while walking around the boundary of a cell, then this edge contributes 2 to the number of sides of the cell.

We distinguish two cases.

Case 1. *There exists no $s \in S$ which is a proper cut of $\tilde{H}_i$.*
Set $m := i$, $\tilde{H} := \tilde{H}_i$, $\tilde{E} := \tilde{E}_i$, and the procedure terminates.

Case 2. *There exists $s \in S$ which is a proper cut of $\tilde{H}_i$.*
Suppose that s is a segment of an unused edge $e \in E_i$ at stage i. Set $\tilde{H}_{i+1} := \tilde{H}_i(s)$, $\tilde{E}_{i+1} := \tilde{E}_i \cup \{e\}$, and proceed further to the next step.

Since $\tilde{G}$ has finitely many edges, this procedure will terminate. Let $\tilde{E}$ and E denote the set of used and unused edges after the last stage, resp., i.e., $\tilde{E} = \tilde{E}_m$ and $E = E_m$.

Claim 2.3. *$\tilde{H}$ has fewer than $8n$ edges. Consequently, the number of used edges of $\tilde{G}$ after the last stage satisfies $|\tilde{E}| < 8n$.*

Proof. Let ϵ_i (resp., δ_i) be the number of edges (resp., cells) in $\tilde{H}_i$. We know that $\epsilon_1 < 2n$ and $\delta_1 = 1$. For every i, we have $\epsilon_{i+1} \leq \epsilon_i + 3$ and $\delta_{i+1} = \delta_i + 1$, so that $\epsilon_m < 3m + 2n - 3$, $\delta_m = m$. On the other hand, each cell has at least 8 edges. This holds for $\tilde{H}_1 = \tilde{T}$, provided that $n \geq 5$, and, by construction, it remains true as segments are added. So we have $\epsilon_m \geq 4\delta_m$. It follows that $3m + 2n - 3 > \epsilon_m \geq 4\delta_m = 4m$. Thus, we obtain $2n - 3 > m$. This implies that for the number of edges of $\tilde{H}$ we have $\epsilon_m < 8n$.

The second statement follows from the fact that every used edge of $\tilde{G}$ has a segment which is an edge of $\tilde{H}$. □

It remains to estimate the number of unused edges of $\tilde{G}$ at the final stage. For any vertex v of $\tilde{G}$, let $\deg(v)$ be the number of unused edges of $\tilde{G}$ incident to v, and let $\deg_{\tilde{H}}(v)$ be the degree of v in $\tilde{H}$. Consider the cyclic order of the unused edges of $\tilde{G}$ and the edges of $\tilde{H}$ incident to v, in a small neighborhood of v. We show that between any fixed pair of consecutive edges e' and e'' of $\tilde{H}$, there are at most $(2k-2)24^{4^l}$ unused edges of $\tilde{G}$ in this sequence. Let $E(e', e'')$ be the set of all such unused edges. An edge $e \in E(e', e'')$ is called *long* or *short*, according to whether it has altogether at least 4^l or fewer than 4^l crossings with edges of $\tilde{H}$, respectively.

In what follows, we would like to distinguish between the two sides of the same edge. To this end, orient the edges of $\tilde{H}$ arbitrarily, and orient every element of $E(e', e'')$ away from v. If two oriented edges, e and f, cross at a point p, we say that at this point e *crosses f from left to right* (or, equivalently, *f crosses e from right to left*), if the direction of e at p can be obtained from the direction of f at p by a clockwise turn through an angle less than π.

For any long edge $e = \overrightarrow{vw} \in E(e', e'')$, consider the first 4^l crossings along e with the edges of $\tilde{H}$, as we move from v towards w, and let $\langle h_1, h_2, \ldots, h_{4^l} \rangle$ be the list of the corresponding edges of $\tilde{H}$. For any $1 \leq i \leq 4^l$, write $t_i = h_i^-$ (resp., let $t_i = h_i^+$) if e crosses h_i from left to right (resp., from right to left). Finally, define the *type* of $e = \overrightarrow{vw}$ as the ordered list (sequence) $T(e) = \langle t_1, t_2, \ldots, t_{4^l} \rangle$.

The type of a short edge $e = vw$ is defined similarly, except that now the list $T(e)$ might be shorter (because we do not have 4^l crossings), and we add w to the list as a last element.

Claim 2.4. *The number of types is at most 24^{4^l}.*

Proof. Since there exists no proper cut of $\tilde{H}$, the two endpoints of each segment of an edge $e = \overrightarrow{vw} \in E(e', e'')$ are "close" to each other in the sense that their distance along the boundary of the corresponding cell of $\tilde{H}$ is at most 6. More precisely, if we fix the first $i < 4^l$ elements of the sequence $T(e)$, there are at most 24 possibilities how to select the $(i+1)$-st element: in both directions we can choose one of the edges supporting the 6 closest sides of the cell, or one of the 6 closest vertices as w (if $e = vw$ is short). Thus, the total number of possible types is at most 24^{4^l}. □

Claim 2.5. *For any fixed type, the number of edges in $E(e', e'')$ is at most $2k-2$.*

Proof. The statement is trivial if the type corresponds to a short edge. Indeed, if two short edges belonging to $E(e', e'')$ have the same type, then both of their endpoints coincide, which is impossible, because our graph has no parallel edges.

Suppose now, in order to obtain a contradiction, that there are $2k-1$ long edges, $e_1, e_2, \ldots, e_{2k-1} \in E(e', e'')$, which have the same type $(t_1, \ldots, t_{4^l})$, where, for each j, $t_j = h_j^+$ or h_j^- for some edge h_j of $\tilde{H}$.

For $1 \le i \le 2k-1$ and $1 < j \le 4^l$, let e_i^j denote the segment of e_i between its $(j-1)$-st and j-th crossing with the edges of $\tilde{H}$. Let e_i^1 stand for for the segment of e_i between v and the first crossing of e_i with an edge of $\tilde{H}$, as we move away from v. Furthermore, let $\bar{e}_i = \cup_{j=1}^{4^l} e_i^j$, that is, the piece of e_i between v and its 4^l-th crossing with the edges of $\tilde{H}$.

Clearly, for any fixed $1 \le j \le 4^l$, the segments e_i^j, for $1 \le i \le 2k-1$, are pairwise non-crossing. Indeed, otherwise, let j be the smallest number such that e_i^j and $e_{i'}^j$ have a crossing p, for some $i \ne i'$. Then the lens enclosed by the pieces of $\bar{e}_i$ and $\bar{e}_{i'}$ between v and p would be vertex-free, contradicting Claim 2.1. Therefore, the curves $\bar{e}_1, \bar{e}_2, \ldots, \bar{e}_{2k-1}$ "run parallel", in the same order. More precisely, if $\bar{e}_1, \bar{e}_2, \ldots, \bar{e}_{2k-1}$ emanate from v in this clockwise order, then for each $1 \le j \le 4^l$, the segments $e_1^j, e_2^j, \ldots, e_{2k-1}^j$ meet the edge h_j in the same order. See Figure 4.

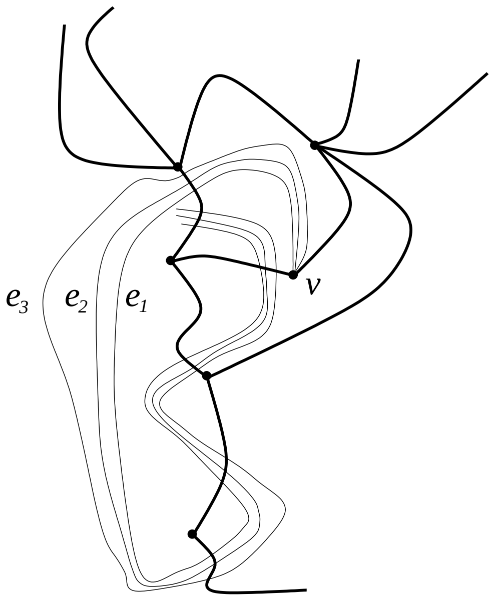

Figure 4. *The edges e_1, e_2, e_3 run parallel.*

It follows from the definition that $\bar{e}_k$ participates in at least 4^l crossings. By Claim 2.2, there are at least $2l$ *distinct* edges of $\tilde{G}$ that cross $\bar{e}_k$. Fix any one of them, e, and denote by x one of its intersection points with $\bar{e}_k$.

It is sufficient to show that e must cross either every element of $\{\bar{e}_1, \bar{e}_2, \ldots, \bar{e}_k\}$ or every element of $\{\bar{e}_k, \bar{e}_{k+1}, \ldots, \bar{e}_{2k-1}\}$. Indeed, if this holds, then at least half of the $2l$ distinct edges of $\tilde{G}$ that cross $\bar{e}_k$ will also cross every element in one of the above sets. Thus, $\tilde{G}$ will contain a radial (k,l)-grid, contradicting our assumptions.

Suppose, in order to obtain a contradiction, that there are two arcs, $\bar{e}_a$ and $\bar{e}_b$, $1 \leq a < k < b \leq 2k-1$, that do not cross e. For any $1 < j \leq 4^l$, let R^j denote the region bounded by e_a^j, e_b^j, and the portions of h_{j-1} and h_j between their intersection points with e_a^j and e_b^j.

By definition and by the previous observations, R^j fully contains e_k^j, but it has no vertex of $\tilde{H}$ in its interior. Let $j \leq k$ be the *minimum* integer such that e has a point in R^j. Since e cannot end in R^j, it must meet one of its sides, belonging e_a^j, e_b^j, h_{j-1}, or h_j. (We set $h_0 := v$. Note also that R^1 has only three sides.) By definition, e cannot meet e_a^j and e_b^j. It is also impossible that e enters and leaves R^j through its side belonging to h_j, because then, within R^j it would form an empty lens with h_j, contradicting Claim 2.1. Thus, if $j > 1$, then e must leave R^j through its side belonging to h_{j-1}. This implies that e must have a point in R^{j-1} contradicting the minimality of j. On the other hand, if $j = 1$, then e must end at $h_0 = v$. However, in this case, the portions of e and $\bar{e}_k$ between v and x would form an empty lens within the region $R^1 \cup R^2 \cup \cdots \cup R^{4^l}$. This contradiction completes the proof of Claim 2.5. $\square$

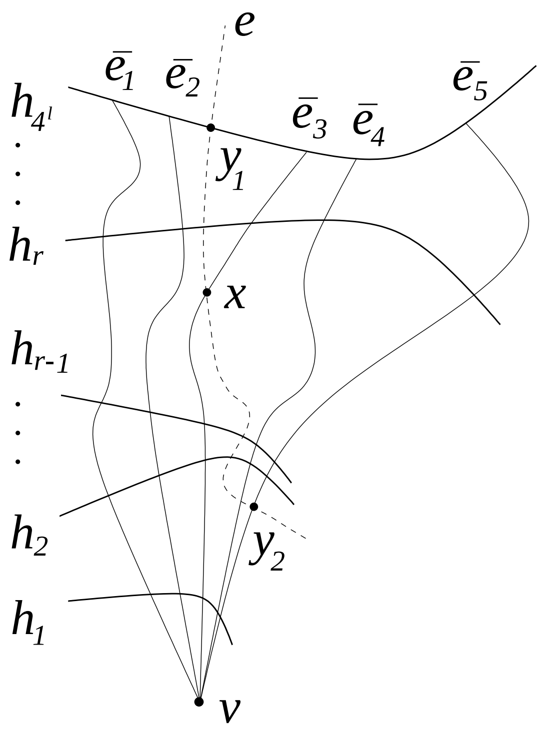

Figure 5. *Any edge of $\tilde{G}$ that crosses $\bar{e}_3$ must also cross either $\bar{e}_1$ and $\bar{e}_2$, or $\bar{e}_4$ and $\bar{e}_5$.*

Now we return to the proof of Theorem 1. Recall that $\tilde{E}$ and E denote the sets of used and unused edges of $\tilde{G}$ at the last stage, respectively. Thus, combining the last three claims, we obtain

$$|E(\tilde{G})| = |E| + |\tilde{E}| \leq |E| + |E(\tilde{H})| \leq \left(24^{4^l}(2k-2)+1\right)|E(\tilde{H})| \leq \left(24^{4^l}(2k-2)+1\right)8n$$

$$< 16 \cdot 24^{4^l} kn,$$

as required.

3. Proof of Theorem 2

We only sketch the proof, since it is a specialization of the previous one. Let G be a simple connected topological graph with n vertices that does not contain a radial (k,l)-grid.

In contrast to the previous proof, now we do not have to build another topological graph, $\tilde{G}$, since G is simple, i.e., any pair of edges of G cross at most once. Thus, we can follow the previous argument with $\tilde{G} = G$.

Construct a topological graph $\tilde{H}$ in exactly the same way as in the previous proof, and define the sets of used and unused edges as before. Obviously, Claim 2.3 remains true.

Now we estimate the number of unused edges of G. For any vertex v of G, let $\deg(v)$ denote the number of unused edges of G incident to v, and let $\deg_{\tilde{H}}(v)$ denote the degree of v in $\tilde{H}$. As before, consider a small neighborhood of v, and list the unused edges of G and the edges of $\tilde{H}$ incident to v in the order as they emanate from v. Now we can show that between any two consecutive edges e' and e'' of $\tilde{H}$, there are fewer than $24^l k$ unused edges of G: Let $E(e',e'')$ denote the set of all such edges. For each $e \in E(e',e'')$, consider the first l crossings of e with the edges of $\tilde{H}$, and let $T(e)$ denote the sequence of the corresponding edges of $\tilde{H}$. As before, $T(e)$ is called the *type* of e. As before, edges with fewer crossings with $\tilde{H}$ have shorter types; they list all these crossings and terminate at the other endpoint of the edge. Instead of Claims 2.4 and 2.5, now we obtain

Claim 3.1. *The number of different types is at most* 24^l. □

Claim 3.2. *For any fixed type,* $E(e',e'')$ *has at most* $k-1$ *elements.*

Proof. The case of short edges is argued as above. Suppose, for the sake of contradiction, that there are k elements, $e_1, e_2, \ldots, e_k$, incident to v, that "run parallel" for at least l steps, i.e., we have $T = T(e_1) = T(e_2) = \ldots = T(e_k)$. Since G is *simple*, the l edges of G corresponding to T are all distinct. Thus, we have found a radial (k,l)-grid in G, the desired contradiction. □

Now the proof can be completed in exactly the same way as before.

References

[AAPPS97] P. K. Agarwal, B. Aronov, J. Pach, R. Pollack, and M. Sharir, Quasi-planar graphs have a linear number of edges, *Combinatorica* **17** (1997), 1–9.

[BKV03] P. Braß, G. Károlyi, and P. Valtr, A Turán-type extremal theory of convex geometric graphs, in: *Discrete and Computational Geometry – The Goodman-Pollack Festschrift* (B. Aronov et al., eds.), Algorithms and Combinatorics **25**, Springer Verlag, Berlin, 2003, 275-300.

[P99] J. Pach, Geometric graph theory, in: *Surveys in Combinatorics, 1999* (J. D. Lamb and D. A. Preece, eds.), London Mathematical Society Lecture Notes **267**, Cambridge University Press, Cambridge, 1999, 167–200.

[PRT02] J. Pach, R. Radoičić, and G. Tóth, Relaxing planarity for topological graphs, in: *Finite and Infinite Combinatorics* (E. Győri et al., eds), Bolyai Society Lecture Notes, to appear.

[PRT03] J. Pach, R. Radoičić, and G. Tóth, On quasiplanar graphs, in preparation.

[PPTT02] J. Pach, R. Pinchasi, G. Tardos, and G. Tóth, Geometric graphs with no self-intersecting path of length three, in: Graph Drawing (M. T. Goodrich, S. G. Kobourov, eds.), Lecture Notes in Computer Science **2528**, Springer-Verlag, Berlin, 2002, 295–311.

[PT97] J. Pach and G. Tóth, Graphs drawn with few crossings per edge, *Combinatorica* **17** (1997), 427–439.

[SS01] M. Schaefer and D. Stefankovič, Decidability of string graphs, in: *Proceedings of the 33rd Annual Symposium on the Theory of Computing (STOC 2001)*, 2001, 241–246.

Received: December 22, 2002
Final version received: December 22, 2002