

Crossing stars in topological graphs*

Gábor Tardos[†]

Géza Tóth[‡]

December 18, 2008

Abstract

Let G be a graph without loops or multiple edges drawn in the plane. It is shown that, for any k , if G has at least $C_k n$ edges and n vertices, then it contains three sets of k edges, such that every edge in any of the sets crosses all edges in the other two sets. Furthermore, two of the three sets can be chosen such that all k edges in the set have a common vertex.

1 Introduction

A *topological graph* is a graph drawn in the plane with no loops or multiple edges so that its vertices are represented by points, and its edges by Jordan curves connecting the corresponding points. We do not distinguish these points and curves of the topological graph from the vertices and edges of the underlying abstract graph they represent. We assume that (i) the edges of a topological graph do not pass through any vertex, (ii) two edges share a finite number of interior points and they properly cross each other, and (iii) no three edges cross at the same point. Conditions (ii) and (iii) are simplifying assumptions only, graph drawings violating them can be modified to satisfy them without effecting which pairs of edges cross. 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 is well known 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.

Pach et al. [PT97], [PRTT04] proved that a topological graph of n vertices and more than $(k + 2)(n - 2)$ edges must have k edges that cross the same edge. This bound is tight for $k = 1, 2, 3$, but can be substantially improved for large values of k .

For $k \geq 2$, let $f_k(n)$ (resp. $f_k^s(n)$) be the maximum number of edges of a topological graph (resp. simple topological graph) on n vertices and no k pairwise crossing edges. Agarwal et al. [AAPPS97] proved (for simple topological graphs) and then, with a shorter and more general argument, Pach et al. [PRT03] proved that for some $c > 0$, every topological graph with n vertices and more than cn edges

*A preliminary version appeared in [TT05]

[†]School of Computing Science, Simon Fraser University and Rényi Institute, Hungarian Academy of Sciences, Budapest; tardos@cs.sfu.ca. Supported by the NSERC grant 611470 and the Hungarian Foundation for Scientific Research Grant (OTKA) Nos. T037846, T046234, AT048826, and NK62321.

[‡]Rényi Institute, Hungarian Academy of Sciences, Budapest; geza@renyi.hu. Supported by OTKA T038397 and T046246.

has *three* pairwise crossing edges. That is, $f_3^s(n) \leq f_3(n) \leq cn$. Very recently, Ackerman and Tardos [AT07] proved that $7n - O(1) \leq f_3(n) \leq 8n$, and $f_3^s(n) = 6.5n + \Theta(1)$. Moreover, Ackerman [A06] managed to prove that $f_4(n) \leq 36n$. For $m \geq 5$ the best known upper bounds are $f_m^s(n) \leq cn \log^{2m-8}$ and $f_m(n) \leq c'n \log^{4m-16} n$ (see [PSS96], [PT05], [A06]), while the best known lower bounds are all linear functions of n and they are conjectured to be much closer to the truth.

Conjecture 1. *For every $k \geq 3$ there is a $c_k > 0$ such that every topological graph with n vertices and $C_k n$ edges contains k pairwise crossing edges.*

In [PRT04], the results about *three* pairwise crossing edges were further generalized: for every integer $k > 0$, there exists a constant $c_k > 0$, such that every topological graph with n vertices and more than $c_k n$ edges has $k + 2$ edges such that the first two cross each other and both of them cross the remaining k edges (see Fig. 1a).

In [PPST05] another generalization was shown. For any k and l there is a constant $c_{k,l}$ with the following property. Every topological graph with n vertices and more than $c_{k,l} n$ edges has $k + l$ edges such that the first k have a common vertex, and each of them cross all of the remaining l edges (see Fig. 1b).




Figure 1: A topological graph without either configuration has only a linear number of edges.

In this paper we prove a common generalization of the above results.

Let k be a positive integer. The edges $A \cup B \cup X$ of a topological graph form a *k-star grid* if A is a set of k edges incident to a common endpoint x , B is a set of k edges incident to a common endpoint y and any edge from A crosses any edge from B , furthermore X also contains k edges and any edge in X crosses all edges in $A \cup B$. See Figure 2. In this definition we allow the case $x = y$ and we also allow the edges of X to be incident to x or y . These pathological cases are not possible in a simple topological graph.

Theorem 1. *For any $k \geq 1$, there is a constant C_k such that every topological graph with n vertices and at least $C_k n$ edges contains a k -star grid.*

We did not attempt to optimize our proof for C_k but note that this proof gives C_k that is triply exponential in k . The condition that all edges in the set A (resp. B) have a common endpoint is essential; our proof does not work if we want to have *independent* edges in the set A (or B). The situation is very similar with the previous results [PRT04], [PPST05], and we cannot even prove the following well-established conjecture:

Conjecture 2. For every $k \geq 2$ there is a $C_k > 0$ such that every topological graph with n vertices and $C_k n$ edges contains $k + 2$ independent edges such that the first two of them cross each of the last k .




Figure 2: A 4-star grid.

2 Proof of the Theorem

The proof of Theorem 1 is rather technical and consists of several steps. We give an overview first and indicate which steps of the proofs can be eliminated if we only consider simple topological graphs. Note that we do not strive for absolute preciosity in this overview. The reader finds the precise definitions later in the proof.

For the proof we fix k , take an arbitrary topological graph F . We let $C = |E(F)|/|V(F)|$. Our goal is to prove that if C is large enough (as a function of k), then we find a k -star grid in F . This clearly establishes Theorem 1.

First we take a *densest subgraph* F_0 of F and concentrate on F_0 only.

Next we *redraw* F_0 , i. e., we take another topological graph G_0 which has the same underlying abstract graph as F_0 but eliminates certain unnecessary crossings. This step of the proof is not needed if F is a simple topological graph, i. e., we may take $G_0 = F_0$.

We then use *subdivisions*, i. e., we introduce vertices at certain edge-crossings. We obtain a subdivision G_1 of G_0 with a crossing-free spanning tree T . This step is taken from [PRT04] and [PPST05].

We further subdivide G_1 to obtain G_2 and its crossing-free spanning subgraph H with no *proper cut*. This means that any two consecutive crossing points of any edge e in $G_2 \setminus H$ with H are with “close-by” edges of H . This step is taken from [PPST05]. In this and the previous step we make sure that the size (number of vertices) of the graph increases by a constant factor only. Note also that subdivisions in these two steps can create k -star grids. This does not happen for simple topological graphs for $k > 2$.

The next step represents the new idea in this paper. For many vertices we find a large number of edges emanating from that vertex with the property that they go “parallel” (with respect to H) for a long time and then one by one they “depart” from the rest of the edges. All these “departures” take place in separate cells of H . We call these sets of edges *bundles*.

Using that C is large enough we find a *cross-track configuration* in G_2 , i. e., k edges of a bundle, another k edges of a (perhaps different) bundle such that these $2k$ edges go parallel through $l - 1$ cells of H but still, eventually the first k edges cross the second k edges. For simple topological graphs we can choose $l = k$ and the proof ends here. Indeed, the $2k$ edges in the cross-track configuration plus k edges of H form a k -star grid. In the general case however, some of the edges of H crossed by the edges in the cross-track configuration may coincide or may be parts of the same edge of G_0 separated only by our subdivision process. We take l to be an exponential function of k and use the following result of Schaefer and Stefankovič [SS04] to take care of the technical difficulties mentioned above.

Theorem A. (Schaefer and Stefankovič, [SS04]) *Let T be a topological graph. Redraw T so that the resulting topological graph T' satisfies the following two conditions:*

- (i) *If two edges of T' cross each other, then the corresponding edges also cross in T ;*
- (ii) *T' has the minimum number of crossings among all drawings with property (i).*

Now for any $i > 0$ and any edge e , any 2^i consecutive crossings on e arise from at least i different edges.

By an application of Theorem A, we show that if l is large enough, then out of the l edges of H crossed by the parallel track of the edges of the cross-track configuration at least k must come from distinct edges of G_0 .

We continue with the detailed execution of the above plan.

Let $k \geq 1$ fixed, and let F be a topological graph with n' vertices and Cn' edges. Our goal is to prove that F contains a k -star grid if C is large enough. The bound on C depends on k but not on n' . This will establish the validity of Theorem 1.

Let F_0 be the densest non-empty connected subgraph of F that is, $F_0 \subseteq F$ connected and $|E(F_0)|/|V(F_0)|$ is maximal. Clearly, the requirement that F_0 has to be connected does not change the value of the maximum, so we have $|E(F_0)|/|V(F_0)| \geq |E(F)|/|V(F)| = C$. Removing a vertex of F_0 of degree d increases the ratio if $d < C$, therefore each vertex in F_0 has degree at least C . Let n denote the number of vertices of F_0 . Clearly, $n > C$ so we may assume $n \geq 5$.

Redraw F_0 so that the resulting topological graph G_0 satisfies the following two conditions:

- (i) *If two edges of G_0 cross each other, then the corresponding edges also cross in F_0 ;*
- (ii) *G_0 has the minimum number of crossings among all drawings with property (i).*

It is enough to find a k -star grid in G_0 as property (i) shows that the corresponding edges form a k -star grid in F_0 and thus in F too.

We will apply a *subdivision* to G_0 , i. e., we declare a certain intersection point of two edges as a *new vertex* and replace each of the two edges by their two segments up to and from that new vertex. Notice that this way we may create two edges connecting the same pair of vertices, thus we have to extend our definition of topological graph to allow for this. No pair of vertices will ever be connected by more than two edges. The graph obtained from G_0 by several subdivisions is called a *subdivision of G_0* . To distinguish from the new vertices of the subdivision, vertices of G_0 are called *old vertices*.

Notice that for $k > 2$, subdivision does not introduce a k -star grid in a simple topological graph, so if G_0 is simple it is enough to find a k -star grid in a subdivision of G_0 . The situation is somewhat more complex if G_0 is not simple. If G_0 contains two k -edge stars A and B such that each edge of A is crossed by each edge of B and another edge e_0 crosses every edge in $A \cup B$ k times, then the repeated subdivision of e_0 may result in a k -star grid.

Obviously, no edge of G_0 intersects itself, otherwise we could reduce the number of crossings by removing the loop. Suppose that G_0 has two distinct edges, e and f , that meet at least twice (including their common endpoints, in the case they have). A simply connected region whose boundary is composed of an arc of e and an arc of f is called a *lens*.

Claim 1. *Every lens in G_0 has a vertex in its interior.*

Proof. Suppose, for a contradiction, that there is a lens ℓ that contains no vertex of 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. \square

Clearly, the property of having no self-intersecting edge and the property stated in Claim 1 are both inherited from G_0 to its subdivisions.

Let G be a topological graph, H a subgraph of G . Let e be an edge of G not contained in H . We always consider e with an orientation. Each edge can be considered with either orientation. The edge e has a finite number of intersection points with edges of H , these points split the Jordan curve e into a finite number of shorter curves. We call these shorter curves the *segments* of the edge e determined by H and denote them by $s_1(e), s_2(e), \dots$ in the order they appear on e . The dependence on H is not explicit in the notation but H will always be clear from the context. If e does not cross the edges of H the entire edge is a single segment.

We consider a crossing-free subgraph H of a topological graph G that is connected and contains all vertices. Such a graph H subdivides the plane into *cells*. The boundary of a cell is closed walk in H that may visit vertices several times and may even pass through an edge twice. The *size* of a cell is the length of the corresponding walk, that is, the the number of edges in the walk, with multiplicity. A segments s of an edge e not in H inherits its orientation from e . It is contained in single cell α , the endpoints of s are on the boundary of α . We call the cell α and the vertex or edge of the boundary walk of α where s starts the *origin* of s . Similarly, α and the vertex or edge of this walk where s ends is *destination* of s . Notice that in case the boundary of α visits the relevant vertex or edge more than once the origin or destination of e contains more information than the vertex or edge itself, it tells us “which side” of the vertex or edge is involved. If two segments have the same origin and the same destination we call them *parallel* and say that their *type* is the same. If two segments s and s' have the same origin but different destinations, then they are contained in the same cell. We say that s *turns left from* s' if the common origin, the destination of s , and the destination of s' appear in this order in the clockwise tour of the boundary of the cell. Notice that the common origin must differ from either of the destinations. A segment with equal origin and destination would define an “empty lens” contradicting Claim 1. As a consequence, for segments s and s' with a common origin, either s and s' are parallel, or s turns left from s' , or s' turns left from s .

As in [PRT04] and [PPST05], first we construct a subdivision G_1 of G_0 that contains a crossing-free spanning tree T .

Since the abstract underlying graph of G_0 is connected, we can choose a sequence of edges

$e_1, e_2, \dots, e_{n-1} \in E(G_0)$ such that e_1, e_2, \dots, e_i form a tree T_i , for every $1 \leq i \leq n-1$. In particular, e_1, e_2, \dots, e_{n-1} form a spanning tree T_{n-1} of G .

Construct the crossing-free topological graphs $\tilde{T}_1, \tilde{T}_2, \dots, \tilde{T}_{n-1}$, as follows. Each is a subtree of a subdivision of G_0 . 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 $1 \leq i < n-1$, 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. If this is a vertex of \tilde{T}_i simply add e_{i+1} to \tilde{T}_i to get \tilde{T}_{i+1} . If v' is in the interior of an edge e then we apply subdivision: we introduce v' as a new vertex. We replace the edge e of \tilde{T}_i with the two resulting parts and add the segment of e_{i+1} between v and v' to obtain \tilde{T}_{i+1} .

We let $T = \tilde{T}_{n-1}$ and G_1 be the subdivision of G_0 obtained in the process. Note that G_1 has n old and at most $n-2$ new vertices.




Figure 3: Constructing \tilde{T}_5 from T_5 .

Next, just like in [PPST05], we further subdivide G_1 to obtain G_2 and a crossing-free subgraph H of G_2 .

Start with $H_0 = T$ and $\tilde{G}_0 = G_1$. Define H_1, \dots, H_u and $\tilde{G}_1, \dots, \tilde{G}_u$ recursively, maintaining that H_i is a crossing-free connected subgraph of a subdivision \tilde{G}_i of G_0 . Furthermore H_i is connected, it contains all vertices of \tilde{G}_i and all the cells of H_i are of size at least 8. This clearly holds for H_0 and \tilde{G}_0 if $n \geq 5$.

Having defined H_i and \tilde{G}_i consider the segments of the edges of \tilde{G}_i as determined by H_i . Let s be such a segment. By *adding s to H_i* we mean constructing a subdivision of \tilde{G}_i by inserting new vertices for the endpoints of s if necessary, and defining a subgraph H_i^s of it by adding s to H_i . More precisely, we also have to replace any edge of H_i that contains in its interior an endpoint of s by the two new edges resulting from the subdivision. Notice that s itself is an edge after the subdivision. The resulting graph H_i^s is a crossing-free connected spanning subgraph of the resulting subdivision of \tilde{G}_i . The cell of H_i containing s is now subdivided into two cells, the other cells remain intact (but their size may increase). We call s a *proper cut of H_i* if both new cells of H_i^s are of size at least 8.

If there exists a proper cut of H_i , then we choose one such segment s and set $H_{i+1} = H_i^s$ and let \tilde{G}_{i+1} be the resulting subdivision of \tilde{G}_i . If there is no proper cut of H_i we set $u = i$, $H = H_u$ and $G_2 = \tilde{G}_u$.

The number of cells starts at 1 cell, at $H_0 = T$, and increases by 1 in every step, so H_i contains




Figure 4: A proper cut.

$i + 1$ cells. Each of these cells are of size at least 8, so we have at least $4i + 4$ edges in H_i . From the Euler formula, the number of vertices v_i of H_i is at least $3i + 5$. As $H_0 = T$ contains at most $2n - 2$ vertices and we introduce at most 2 new vertices in every step, so we also have $v_i \leq 2i + 2n - 2$. The upper and lower bounds on v_i imply $i \leq 2n - 7$. So the above process terminates in $u \leq 2n - 7$ steps. This proves the following

Claim 2. G_2 is a subdivision of G_0 with at most $6n - 16$ vertices. H is a connected, spanning, crossing-free subgraph of G_2 with no proper cut. H has at most $8n - 24$ edges.

We call an old vertex of G_2 *important* if its degree in H is less than 32. By Claim 2, H has less than $n/2$ vertices of degree 32 or more. Out of the n old vertices we must have more than $n/2$ important vertices.

Let $l = 2^{k+1}k^2 + 1$. Consider an edge e of G_2 not in H . Call any l consecutive of the segments $s_1(e), s_2(e), \dots$ a *track* of e . The *type* of a track is simply the sequence of the types of the l segments, $s_i(e), \dots, s_{i+l-1}(e)$. Tracks (of possibly different edges) of the same type are called *parallel*. Consider two edges e and f of G_2 that are not in H . Let $d(e, f)$ be the largest index $i \geq 1$ such that for all $1 \leq j < i$ the segments $s_j(e)$ and $s_j(f)$ exist and are parallel. For example, if e and f start at different vertices or in different cells we have $d(e, f) = 1$.

Notice that for any origin of a segment at most 24 destinations are possible. For large cells of H more choices would be possible but they yield proper cuts of H which do not exist by Claim 2. By the same claim there are less than $32n$ possible origins and therefore less than $768n$ types of segments. The destination of a segment determines the origin of the next segment, therefore there are less than $32 \cdot 24^l n$ different types of tracks.

Let $m = 300k \cdot 24^l$. We call the sequence e_1, \dots, e_{2m} of $2m$ edges of G_2 but not in H a *bundle* if $l \leq d(e_1, e_{2m}) < d(e_2, e_{2m}) < \dots < d(e_{2m-1}, e_{2m})$. Notice that the edges of a bundle start at a common vertex. We say that the bundle *emanates* from this common starting vertex.

Claim 3. If $C \geq C_k := 31 \cdot 24^{2m+l} + 31$, then there exists a bundle emanating from every important vertex.

Proof. Consider an important vertex x . Let S_0 be the set of edges of G_2 not in H that start at x . The vertex x has degree at least C in G_0 and it has the same degree in its subdivision G_2 . Its degree in H is at most 31, so $|S_0| \geq C - 31$. For $i \geq 1$ we define S_i to be a subset of maximal size of S_{i-1} with $s_i(e)$ existing and having equal type for each $e \in S_i$. The number of possible origins for the type of segment $s_1(e)$ of an edge $e \in S_0$ is the degree of x in H . Since x is important, at most 31 origins and at most 744 types of $s_1(e)$ may exist for $e \in S_0$. Thus, $|S_1| \geq |S_0|/744$. Notice that the type of $s_i(e)$ determines if e ends with the segment $s_i(e)$ and if so, then it determines the ending vertex. So if one of the edges $e \in S_i$ ends with its i 'th segment, then all does, and they all connect the same pair of vertices. Thus, as long as $|S_i| > 2$, $s_{i+1}(e)$ exists for all $e \in S_i$. Furthermore, the type of $s_i(e)$ determines the origin of $s_{i+1}(e)$. So if $|S_i| > 2$ then $|S_{i+1}| \geq |S_i|/24$.

The finiteness of the entire topological graph G_2 implies that $S_i = \emptyset$ for large enough i . Let $l \leq d_1 < d_2 < \dots < d_v$ be all the indices $d \geq l$ such that $|S_{d+1}| < |S_d|$. The above calculations yield that $|S_{d_1}| \geq 24^{2m}$ and $S_{d_{i+1}} = S_{d_i+1} \geq 24^{2m-i}$ for $i \leq 2m$. We choose e_i to be an arbitrary element of $S_{d_i} \setminus S_{d_{i+1}}$. We have $d(e_i, e_{2m}) = d_i + 1$ for $i < 2m$. This establishes that (e_1, \dots, e_{2m}) form a bundle. \square

Fix a bundle $B^x = \{e_1^x, \dots, e_{2m}^x\}$ from every important vertex x . The existence is given by Claim 3. These will be all the bundles, and in fact all the edges of $G_2 \setminus H$ we consider from now on.

The segments $s_1(e_{2m}^x), s_2(e_{2m}^x), \dots, s_{d^x}(e_{2m}^x)$ for $d^x = d(e_m^x, e_{2m}^x)$ form the *backbone* of the bundle B^x . The tracks of e_{2m}^x contained in the backbone are called the *vertebras*. We denote the vertebra starting with the segment $s_i(e_{2m}^x)$ by t_i^x . Notice that the vertebrae interleave: the last $l - 1$ segments of a vertebra are the first $l - 1$ segments of the next vertebra. With any vertebra t_i^x we find $m - 1$ parallel tracks: the tracks starting with the segments $s_i(e_{m+1}^x), \dots, s_i(e_{2m-1}^x)$.

Let $e = t_i^x$ and $f = t_j^y$ be two distinct parallel vertebrae. Notice that $i > 1$ and $j > 1$ must hold, since we only consider a single bundle from any (important) vertex. Let e' and f' be the inverse orientation of the ‘‘previous’’ segments $s_{i-1}(e_{2m}^x)$ and $s_{j-1}(e_{2m}^y)$, respectively. Notice that e' and f' have the same origin. We say that $e < f$ if e' turns left from f' . We also say that $e < f$ if t_{i-1}^x and t_{j-1}^y are parallel, and $t_{i-1}^x < t_{j-1}^y$. Notice that the recursive definition is well founded and it defines a linear order among parallel vertebrae. We call a vertebra *extremal* if it is smallest or largest among the vertebrae of its type. If e is a non-extremal vertebra we let e^+ stand for the next larger vertebra of the same type, while e^- stands for the next smaller vertebra. We say that a vertebra e is *special* if it is either extremal or one of e^+ or e^- is the last vertebra in a backbone.

Claim 4. *The number of special vertebrae is at most $65 \cdot 24^l n$.*

Proof. We have at most two extremal vertebrae for every type, that is at most $64 \cdot 24^l n$ extremal vertebrae. We have one last vertebra in every backbone, that is at most n last vertebrae. Each last vertebra makes its at most two neighbors special, so the claimed bound holds. \square

We define a *cross-track configuration* as two sets of k edges such that every edge from the first set crosses every edge from the second set, and all $2k$ edges go parallel for a long time. More precisely, let A and B both be sets of k edges. We say that $A \cup B$ is a *cross-track configuration* if the following conditions are satisfied.

- (i) Every $a \in A$ crosses every $b \in B$.
- (ii) Every $a \in A$ is incident to an old vertex x and every $b \in B$ is incident to an old vertex y .
- (iii) There is $\alpha, \beta > 0$ such that for every $a \in A$, $b \in B$, and $0 \leq i < l - 1$, $s_{\alpha+i}(a)$ and $s_{\beta+i}(b)$ exist and are parallel.

Notice that for simple topological graphs a cross-track configuration $A \cup B$ can be appended with the set $X \subseteq E(H)$ consisting of k of the origins of the the segments in the parallel tracks of the edges in $A \cup B$. These edges cross every edge in $A \cup B$, therefore $A \cup B \cup X$ form a k -star grid. Unfortunately, if G_2 is not simple, then X may contain fewer than k edges, in extreme situations X might consist of a single edge (the edges in $A \cup B$ go round and round crossing this single edge many times). Also, finding k -star grid in G_2 is not enough in this case.

Our immediate goal is to find a cross-track configuration in G_2 , see Claim 6. As explained above this leads immediately to a k -star grid in G_2 , and also in G_0 if G_0 is simple. For non-simple topological graphs we will also use the cross-track configuration to find k -star grids in G_0 , but the argument is more involved.

The following claim is based on a similar observation in [AAPPS97].

Claim 5. *Let e and f be two consecutive vertebras of the bundle B^x , neither special. Then e^+ and f^+ are also consecutive vertebras of a backbone or there exists a cross-track configuration in G_2 . The same holds for e^- and f^- .*

Proof. Assume f follows e in B^x and let $e^+ = t_i^y$. We have to show that $f^+ = t_{i+1}^y$. Suppose that $f^+ = t_j^z$.

Since e is not special, $e^* = s_{i+l}(e_{2m}^y)$ is still in the backbone of B^y . Let f^* be the last segment of f . These two segments have a common origin. We distinguish three cases. See Fig. 5.

Case 1: e^* and f^* are parallel. Then, by the definition of the order of vertebras t_{i+1}^y must be f^+ .

Case 2: f^* turns left from e^* . In this case all edges e_a^x intersect all edges e_b^y for $m < a, b \leq 2m$. This provides a cross-track configuration. See Fig 5 (a).

Case 3: e^* turns left from f^* . Now the edges e_a^y and e_b^z must cross for $m < a, b \leq 2m$, and this also provides a cross-track configuration. See Fig 5 (b).

The proof for e^- and f^- is similar. □

We considered at least $n/2$ bundles. By Claim 4 we have at most $65 \cdot 24^l n$ special vertebras, so the pigeonhole principle gives the existence of a bundle B^x with at most $130 \cdot 24^l$ special vertebras. We fix such a bundle B^x and let e_i stand for the i th segment in the backbone of B^x : $e_i = s_i(e_{2m}^x)$ for $1 \leq i \leq d(e_m^x, e_{2m}^x)$. We call e_i a *departure point* if $i = d(e_j^x, e_{2m}^x)$ for some $1 \leq j \leq m$. We look for an interval of the backbone of B^x without special vertebras but with the largest number of departure points. There are m departure points, so at least $\lfloor m/(130 \cdot 24^l + 1) \rfloor$ of them are in an interval that has no special vertebra. Formally, we have $1 \leq i < j \leq d(e_m^x, e_{2m}^x) - l + 1$, such that none of the vertebras $t_i^x, t_{i+1}^x, \dots, t_j^x$ are special, but for some indices $1 \leq i' < j' \leq m$ we have $i + l \leq d(e_{i'}^x, e_{2m}^x) < d(e_{j'}^x, e_{2m}^x) \leq j + l - 1$ and $j' - i' + 1 \geq \lfloor m/(130 \cdot 24^l + 1) \rfloor$.

By Claim 5 we either have a cross-track configuration or the vertebras $(t_i^x)^+, (t_{i+1}^x)^+, \dots, (t_j^x)^+$ are consecutive tracks of some bundle B^y , while $(t_i^x)^-, (t_{i+1}^x)^-, \dots, (t_j^x)^-$ are also consecutive tracks of some bundle B^z . In the latter case for any $i' \leq v \leq j'$ the edge e_v^x crosses all edges e_w^y with $m < w \leq 2m$ or it crosses all edges e_w^z with $m < w \leq 2m$. One of the options must occur with at least $\lfloor m/(260 \cdot 24^l + 2) \rfloor \geq k$ edges. This provides us a set A of k edges of the bundle B^x , another set B of k edges of a bundle such that the properties of cross-track configuration are satisfied. Thus, a cross-tack configuration must exist. See Figure 6. This proves the following

Claim 6. *For $C \geq C_k$ there exists a cross-track configuration in G_2 .*




Figure 5: e^+ and f^+ are consecutive vertexes.

Let $A \cup B$ be a cross-track configuration in G_2 . We use it to find a k -star grid in G_0 .

There are $\alpha, \beta > 0$ such that for every $a \in A$, $b \in B$ and $0 \leq i < l - 1$, the segments $s_{\alpha+i}(a)$ and $s_{\beta+i}(b)$ are parallel. Let $s_i^*(e) = s_{\alpha+i}(e)$ for $e \in A$ and $s_i^*(e) = s_{\beta+i}(e)$ for $e \in B$. We say that $0 \leq i < l - 1$ is *bad* if two distinct segments from the set $\{s_i^*(e) \mid e \in A \cup B\}$ intersect.

Observe that we counted at most one crossing for each pair of edges in $A \cup B$, otherwise we would get an “empty lens”. Therefore, there are at most $\binom{2k}{2}$ bad values of i . So there are $0 \leq i_0 < i_1 \leq l - 1$, with $i_1 - i_0 + 2 > l / (\binom{2k}{2} + 1) > 2^k + 1$ such that there is no bad i with $i_0 \leq i \leq i_1$. For $i_0 \leq i \leq i_1$, let h_i be the edge of H that is the common origin of the segments $s_i^*(e)$ for $e \in A \cup B$. Order the edges $e \in A \cup B$ according to the order the starting points of $s_i^*(e)$ appear on h_i . Notice that we get the same order for each i . Let a and b be the first and last edges in this order. Let p_i and q_i be the starting points of $s_i^*(a)$ and $s_i^*(b)$, respectively. Let a^* be “relevant” part of a , that is, a^* is the interval of a between p_{i_0} and p_{i_1} .

At this point we shift attention from G_2 and H to the original graph G_0 and modify its drawing in the plane. Let S be the set of edges of G_0 containing the edges $A \cup B$ of G_2 . Note that the edges in A are incident to the same old vertex, therefore they cannot be different segments of an edge of G_0 . The same holds for the edges in B . Moreover, any edge of A and any edge of B intersect, so they are not different segments of the same edge. Consequently, S contains $2k$ distinct edges. We do not redraw the edge containing a but redraw some segments of other edges making sure that conditions (i) and (ii) of the definition of G_0 are maintained and furthermore every edge that intersects a^* intersects also all edges in S .

Let $i_0 \leq i < i_1$ and consider the following four intervals: (1) the interval of h_i between p_i and




Figure 6: H_y and H_z envelope a vertebra of H_x .




Figure 7: Procedure REDRAW.

q_i , (2) $s_i^*(a)$, (3) the interval of h_{i+1} between p_{i+1} and q_{i+1} , and (4) $s_i^*(b)$. These segments bound a quadrilateral shaped region R_i , with “vertices” p_i , q_i , p_{i+1} , and q_{i+1} . See Figure 7. We cannot rule out that some of the regions R_i are not disjoint and, in fact, we cannot even rule out that $h_i = h_{i+1}$ (see Fig. 8.) but it does not effect the argument to be presented.

The region R_i does not contain vertices, therefore no edge of G_0 entering R_i through $s_i^*(a)$ may leave R_i through $s_i^*(a)$ again, as that would contradict Claim 1. We distinguish three types of edges of G_0 entering R_i through $s_i^*(a)$. Note that an edge can cross $s_i^*(a)$ several times, in this case we consider all the segments of e inside R_i separately.

Type 1: The edge e enters R_i through $s_i^*(a)$ and leaves R_i through $s_i^*(b)$. In this case, e crosses each edge in S .

Type 2: The edge e enters R_i through $s_i^*(a)$ and leaves R_i through h_i .




Figure 8: We cannot even rule out that $h_i = h_{i+1}$.

Type 3: The edge e enters R_i through $s_i^*(a)$ and leaves R_i through h_{i+1} .

We describe procedure REDRAW. If there exists $i_0 \leq i < i_1$ with an edge of type 2 crossing $s_i^*(a)$, then we choose an arbitrary such i and the edge e of type 2 crossing $s_i^*(a)$ closest to p_i . Let e_a be the point of e where it enters R_i and e_h be the point where it leaves R_i . Let e'_a and e'_h be points on e outside R_i but close to e_a and e_h , respectively. Replace the interval $e'_a e'_h$ of e by a curve outside R_i , which follows very closely the interval of a between e_a and p_i , and then the interval of h_i between p_i and e_h . In case $h_i = h_{i+1}$ the the new curve is drawn similarly, but it does not go outside the region R_i . It is easy to verify that if the new segment of e follows the boundary of R_i close enough, then no new crossings are created and therefore the modified topological graph satisfies properties (i) and (ii). See Figure 7.

If there exists $i_0 \leq i < i_1$ and an edge of type 3 crossing $s_i^*(a)$, then we proceed analogously. We choose such an i arbitrarily, we choose a type 3 edge that crosses $s_i^*(a)$ closest to p_{i+1} and redraw the segment of the edge in R_i taking a detour around p_{i+1} .

As long as there is an i , $i_0 \leq i < i_1$ with a type 2 or type 3 edge, execute REDRAW.

If a^* enters the region R_i (we cannot rule out this possibility), then REDRAW choosing this i effects other regions R_j . In the extreme case when p_{i+1} is on h_i between p_i and q_i , by redrawing edges of type 2 we create another crossing with $s_i^*(a)$ itself, possibly another type 2 crossing. Nevertheless, it can be shown that the procedure terminates after finitely many steps. To see this, consider an edge e . The set $\cup_{i=i_0}^{i_1-1} R_i$ divides e into several intervals. Let e^* be one of them. For each crossing p of e^* and a^* let $r(p) = i$ if and only if p is on $s_i^*(a)$. Let $r(e^*, a^*)$ be the sum of all $r(p)$ over all crossings. This sum will either always decrease or always increase when we execute REDRAW involving e^* , therefore e^* is involved in finitely many steps only. To see this “monotonicity condition” notice that each segment of a^* entering R_i has the “same orientation”, that is, it enters R_i through h_i and leaves through h_{i+1} .

Let G'_0 be the topological graph obtained in the process. All edges of G'_0 crossing the curve a^* cross all edges in S . We did not create any additional crossing, so the graph G'_0 satisfies properties (i) and (ii) in the definition of G_0 . These properties and a result of Schaefer and Stefanković [SS04] imply the following.

Claim 7. For any edge e of G'_0 and for any $i > 0$, any 2^i consecutive crossings on e arise from at least i different edges.

The interval a^* of a crosses H at least 2^k times and we did not “redraw” these segments of edges of G_0 . We can therefore take 2^k consecutive crossings of a^* in G'_0 and by Claim 7 they are from at least k edges. Let X be a set of k edges of G'_0 crossing a^* . Clearly, $S \cup X$ is a k -star grid in G'_0 .

Clearly, the corresponding edges form a k -star grid in F too. This finishes our proof of Theorem 1. \square

References

- [A06] E. Ackerman, On the maximum number of edges in topological graphs with no four pairwise crossing edges, Proceedings of the 22th Annual ACM Symposium on Computational Geometry (SoCG 2006), 2006, 259–263.
- [AT07] E. Ackerman and G. Tardos, The maximum number of edges in quasi-planar graphs, *J. Combin. Theory, Ser. A* **114**, (2007), 563–571.
- [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.
- [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.
- [PPST05] J. Pach, R. Pinchasi, M. Sharir, G. Tóth, Topological graphs with no large grids, *Graphs and Combinatorics* **21** (2005), 355–364.
- [PRTT04] J. Pach, R. Radoičić, G. Tardos, and G. Tóth, Improving the Crossing Lemma by finding more crossings in sparse graphs, Proceedings of the 20th Annual ACM Symposium on Computational Geometry (SoCG 2004), 2004, 68–75.
- [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.
- [PRT03] J. Pach, R. Radoičić, and G. Tóth, Relaxing planarity for topological graphs, in: Discrete and Computational Geometry (J. Akiyama, M. Kano, eds.), Lecture Notes in Computer Science **2866**, Springer-Verlag, Berlin, 2003, 221–232.
- [PRT04] J. Pach, R. Radoičić, and G. Tóth, A generalization of quasi-planarity in: Towards a Theory of Geometric Graphs, (J. Pach, ed.), Contemporary Mathematics **342**, AMS, 2004, 177–183.
- [PSS96] J. Pach, F. Shahrokhi, and M. Szegedy, Applications of the crossing number, *Algorithmica* **16** (1996), 111–117. Also in: Proceedings of the 10th Annual ACM Symposium on Computational Geometry (SoCG 1994), 1994, 198–202.
- [PT97] J. Pach and G. Tóth, Graphs drawn with few crossings per edge, *Combinatorica* **17** (1997), 427–439.

- [PT05] J. Pach and G. Tóth, Disjoint edges in topological graphs, Combinatorial Geometry and Graph Theory: Indonesia-Japan Joint Conference, Bandung, Indonesia, 2003, Revised Selected Papers (J. Akiyama, E. T. Baskoro, M. Kano, eds.) Lecture Notes in Computer Science **3330**, Springer-Verlag, Berlin, 2005, 133–140.
- [SS04] M. Schaefer and D. Stefankovič, Decidability of string graphs, *Journal of Computer and System Sciences* **68** (2004), 319–334. Also in: Proceedings of the 33rd Annual Symposium on the Theory of Computing (STOC 2001), 2001, 241–246.
- [TT05] G. Tardos and G. Tóth, Crossing stars in topological graphs, in: Proceedings of the Japan Conference on Discrete and Computational Geometry 2004, in Honor of János Pach on His 50th Year, Lecture Notes in Computer Science **3742**, Springer-Verlag, Berlin, 2005, 184–197.