Recognizing String Graphs Is Decidable

János Pach and Géza Tóth

Rényi Institute of Mathematics
Hungarian Academy of Sciences
`pach@cims.nyu.edu, geza@renyi.hu`

Abstract. A graph is called a *string graph* if its vertices can be represented by continuous curves ("strings") in the plane so that two of them cross each other if and only if the corresponding vertices are adjacent. It is shown that there exists a recursive function $f(n)$ with the property that every string graph of n vertices has a representation in which any two curves cross at most $f(n)$ times. We obtain as a corollary that there is an algorithm for deciding whether a given graph is a string graph. This solves an old problem of Benzer (1959), Sinden (1966), and Graham (1971).

1 Introduction

Given a simple graph G, is it possible to represent its vertices by simply connected regions in the plane so that two regions overlap if and only if the corresponding two vertices are adjacent? In other words, is G isomorphic to the *intersection graph* of a set of simply connected regions in the plane? This deceptively simple extension of propositional logic and its generalizations are often referred to in the literature as *topological inference problems* [CGP98a], [CGP98b],[CHK99]. They have proved to be relevant in the area of geographic information systems [E93], [EF91] and in graph drawing [DETT99]. In spite of many efforts [K91a], [K98] (and false claims [SP92], [ES93]), no algorithm was found for their solution. It is known that these problems are at least NP-hard [KM89], [K91b], [MP93].

Since each element of a finite system of regions in the plane can be replaced by a simple continuous arc ("string") lying in its interior so that the intersection pattern of these arcs is the same as that of the original regions, the above problem can be rephrased as follows. Does there exist an algorithm for recognizing *string graphs,* i.e., intersection graphs of planar curves? As far as we know, in this form the question was first asked in 1959 by S. Benzer [B59], who studied the topology of genetic structures. Somewhat later the same question was raised by F. W. Sinden [S66] in Bell Labs, who was interested in electrical networks realizable by printed circuits. Sinden collaborated with R. L. Graham, who communicated the question to the combinatorics community by posing it at the open problem session of a conference in Keszthely, in 1976 [G78]. Soon after G. Ehrlich, S. Even, and R. E. Tarjan [EET76] studied the *"string graph problem"* (see also [K83] and [EPL72] for a special case). The aim of this paper is to answer the

above question in the affirmative: there exists an algorithm for recognizing string graphs.

To formulate our main result precisely, we have to agree on the terminology. Let G be a graph with vertex set $V(G)$ and edge set $E(G)$. A *string representation* of G is an assignment of simple continuous arcs to the elements of $V(G)$ such that no three arcs pass through the same point and two arcs cross each other if and only if the corresponding vertices of G are adjacent. G is a *string graph* if it has a string representation. Every intersection point between two arcs is called a *crossing*. (That is, two arcs may determine many crossings.) For any string graph G, let $\mathrm{ST}(G)$ denote the minimum number of crossings in a string representation of G, and let

$$\mathrm{ST}(n) := \max_{|V(G)|=n} \mathrm{ST}(G),$$

where the maximum is taken over all string graphs G with n vertices.

Theorem 1. *Every string graph with n vertices has a string representation with at most $(2n)^{24n^2+48}$ crossings.*

Using the above notation, we have $\mathrm{ST}(n) \le (2n)^{24n^2+48}$. On the other hand, it was shown by J. Kratochvíl and J. Matoušek [KM91] that $\mathrm{ST}(n) \ge 2^{cn}$ for a suitable constant c.

Theorem 1 implies that string graphs can be recognized by a finite algorithm. Indeed, by brute force we can try all possible placements of the crossing points along the arcs representing the vertices of the graph and in each case test planarity (which can be done in linear time in the total number of crossings [HT74]).

As was pointed out in [KM91], the representation of string graphs is closely related to the following problem. Let $R \subseteq \binom{E(G)}{2}$ be a set of pairs of edges of G. We say that the pair (G, R) is *weakly realizable* if G can be drawn in the plane so that only pairs of edges belonging to R are *allowed* to cross (but they do not *have* to cross). Such a drawing is called a *weak realization* of (G, R). The minimum number of crossings in a weak realization of (G, R) is denoted by $\mathrm{CR}(G, R)$. Note that the usual *crossing number* of G is equal to $\mathrm{CR}(G, \binom{E(G)}{2})$.) Let

$$\mathrm{CR}(n) := \max_{|V(G)|=n, R} \mathrm{CR}(G, R),$$

where the maximum is taken over all weakly realizable pairs (G, R) on n vertices). It was proved in [KLN91] and [KM91] that the problems of recognizing string graphs and weakly realizing pairs are *polynomially equivalent*. In particular,

$$\mathrm{ST}(n) \le \mathrm{CR}(n^2) + \binom{n}{2}. \tag{1}$$

Kratochvíl [K98] called it an "astonishing and challenging fact that so far there is no recursive" upper bound known on $\mathrm{ST}(n)$. Our next theorem, which, combined with (1), immediately implies Theorem 1, fills this gap.

Theorem 2. *Let G be a simple graph with m edges, and let (G, R) be a weakly realizable pair. Then (G, R) has a weak realization with at most $(4m)^{12m+24}$ crossings.*

As before, it follows from Theorem 2 that there is a recursive algorithm for deciding whether a pair (G, R) is realizable.

N. Linial has pointed out that the above questions are closely related to estimating the Euclidean distortion of certain metrics induced by weighted planar graphs [LLR95], [R99].

Using completely different (and more elegant) methods, Schaefer and Stefankovič [SS01] have recently established the slightly better bounds $\mathrm{CR}(m) \le m2^m$ and $\mathrm{ST}(n) \le n^2 2^{n^2} + \binom{n}{2}$.

2   Two Simple Properties of Minimal Realizations

In the sequel, let G be a simple graph with n vertices and m edges, let R be a set of pairs of edges, and assume that (G, R) is weakly realizable. Fix a weak realization (drawing) with the *minimum* number of crossings, and assume that this number is at least $(2n)^{12m+24}$. With no loss of generality, all drawings in this paper are assumed to be in *general position*. That is, no edge passes through a vertex different from its endpoints, no three edges have a point in common, and no two edges "touch" each other (i.e., if two edges have a point in common, then they properly cross at this point).

Let A and B be intersection points of two edges $e, f \in E(G)$ and suppose that the portions of e and f between A and B do not have any other point in common. Then the region enclosed by these two arcs is called a *lense*.

Lemma 2.1. *Every lense and its complement contains a vertex of G.*

Proof. By symmetry, it is sufficient to show that every lense contains a vertex. Consider a lense which is minimal by containment, and assume that it is bounded by portions of $e, f \in E(G)$ between A and B. By the minimality, any other edge of G which intersects one side of the lense must also cross the other one. Therefore, replacing the portion of f between A and B by an arc running outside the lense and very close to e, we would reduce the number of crossings in the drawing, contradicting its minimality. □

Deleting from the plane any two arcs, e and f, the plane falls into a number of connected components, called *cells*. At most 4 of them contain and endpoint of e or f. A cell containing no endpoint of e or f is called a *k-cell,* if its boundary consists of k *sides* (subarcs of e and f). Obviously, k must be even, and the sides of a k-cell belong to e and f, alternately. We say that a cell is *empty* if it contains no vertex of G.

Lemma 2.2. *Let e and f be two portions of edges of G that cross each other K times, where $K \ge 16n^3$ and $n \ge 10$. Then e and f determine $M \ge K/(8n^3)$ empty four-cells, $C_1, C_2, \ldots, C_M$, such that C_i and C_{i+1} share a side belonging to f, for every i.*

Proof. The arcs e and f divide the plane into K cells. All but at most 4 of them contain no endpoint of e or f and have an even number of sides.

Define a graph H, whose vertices represent the K cells, and two vertices are joined by an edge if and only if the corresponding cells share a side which belongs to f. Since H is connected and has at most $K-1$ edges, it is a tree. Every leaf of H corresponds to a lense or possibly a cell containing an endpoint of e or f. The number of lenses is at most n, because, by Lemma 2.1, each of them contains a vertex of G. Thus, H has at most n leaves. Consequently, the degree of every vertex of H is at most n.

Delete every vertex of H which corresponds to a cell that (i) either contains an endpoint of e or f, (ii) or contains a vertex of G and has more than 2 sides. The number of deleted vertices is at most $n + 4$, each of them has degree at most n, so the resulting forest consists of at most $n(n + 4)$ trees. Hence, one of these trees, H', has at least $K' = (K - n - 4)/[n(n + 4)]$ vertices. Any leaf of H' is either a leaf of H or is connected to a deleted vertex of H. Therefore, H' has at most $2n$ leaves. This implies that H' contains a path with at least $M = (K' - 1)/(4n - 3) - 1 > K/(8n^3)$ vertices, each of degree 2. The sequence of four-cells corresponding to the vertices of this path meet the requirements of the lemma.   □
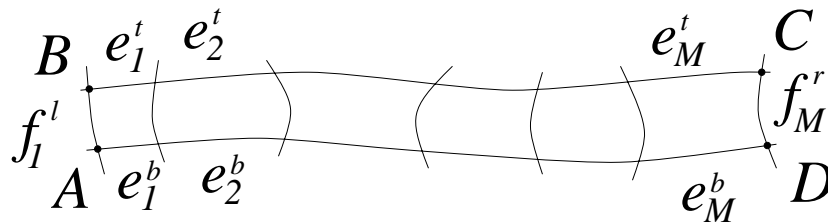

Fig. 1.

A sequence $C_1, C_2, \ldots, C_M$ of four-cells whose existence is guaranteed in Lemma 2.2 is called an *empty (e, f)-path* of four-cells. In what follows, we analyze the finer structure of such a path. We assume that e and f are oriented, and we denote the four sides of C_i by e_i^t, f_i^r, e_i^b and f_i^l, in clockwise order. (Here the superscripts stand for *top, right, bottom,* and *left*, repectively, suggesting that on our pictures the boundary pieces belonging to e are "horizontal"; see Figure 1.) For every $1 \le i < M$, we have $f_i^r = f_{i+1}^l$.

Let A denote the common endpoint of f_1^l and e_1^b, that is, $A = f_1^l \cap e_1^b$. Let $B = f_1^l \cap e_1^t$, $C = f_M^r \cap e_M^t$, and $D = f_M^r \cap e_M^b$.

We distinguish two different types of the empty paths, (See Fig. 2) and study them separately in the next two sections.

Type 1: For every i, the two sides of the cell C_i which belong to e have the *same* orientation (i.e., both of them are oriented towards f_i^r or both are oriented towards f_i^l). See Fig. 2 (a), (b).

Type 2: For every i, the two sides of the cell C_i which belong to e have *opposite* orientations. See Fig. 2 (c).
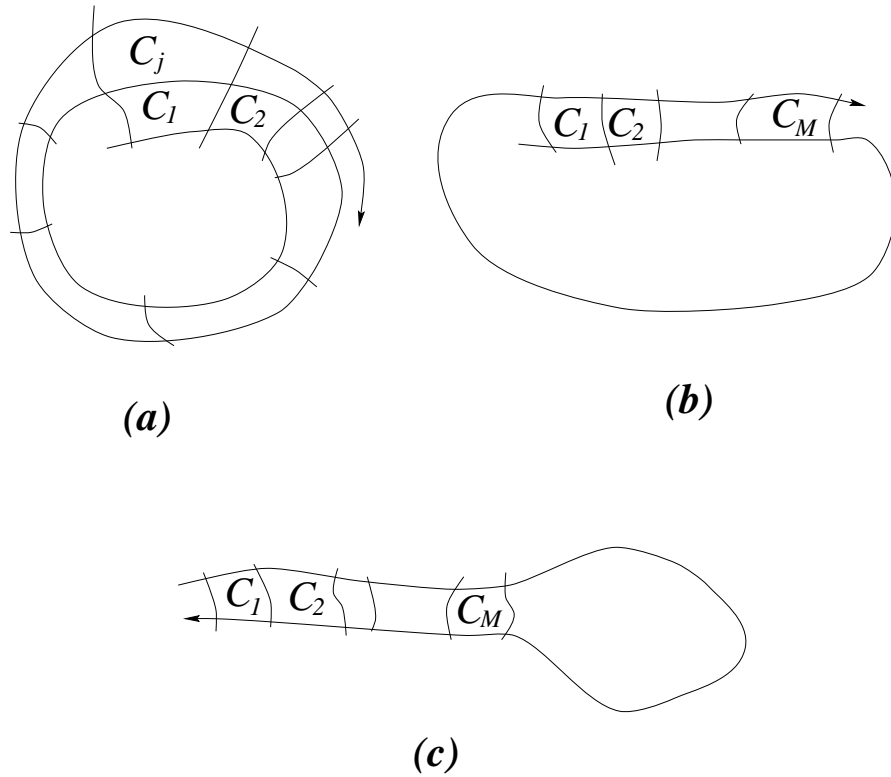

(a)

(b)

(c)

Fig. 2.

3   Empty Paths of Type 1

Lemma 3.1 *Let $C_1, C_2, \ldots, C_M$ be an empty (e, f)-path of four-cells of type 1. Then e_1^t and e_1^b cannot coincide with any side of C_i, $1 < i \leq M/(5m)$.*

Proof. Suppose, in order to obtain a contradiction, that e.g. $e_1^t = e_{j+1}^b$, for some $1 < j < M/(5m)$. This easily implies $e_i^t = e_{j+i}^b$, for every $i \leq M - j$. See Fig 2 (a).

Assume without loss of generality that M can be written in the form $M = M'j + 1$, for a suitable $M' \geq 5m$. Let $\hat{f} = f_1^l \cup f_{j+1}^l \cup \cdots \cup f_M^l$, which is a segment of f, and orient it from f_1^l towards f_M^l. Let $F_1 = f_1^l \cap e_1^b$ and $F_2 = f_M^l \cap e_M^t$ be the starting point and the endpoint of $\hat{f}$, resp., and let $I = C_1 \cup C_2 \cup \cdots \cup C_M$. Furthermore, let J_1 denote the region bounded by $j_1 = e_1^b \cup e_2^b \cup \ldots \cup e_j^b \cup f_1^l$

which does not contain I, and let J_2 be the region bounded by $j_2 = e_M^t \cup e_{M-1}^t \cup \ldots \cup e_{M-j+1}^t \cup f_M^r$ which does not contain I.

Let S denote the set of arcs obtained by intersecting the edges of G with I. Since every lens contains a vertex of G, and I consists of *empty* four-cells, there is no lens in I. Thus, one of the endpoints of every $x \in S$ is on j_1, and the other on j_2.

Definition 3.2. For any two oriented edges, $\overline{e}$ and $\overline{f}$, crossing at some point X, we say that $\overline{e}$ *crosses* $\overline{f}$ *from left to right* if the direction of $\overline{e}$ at X can be obtained from the direction of $\overline{f}$ at X by a clockwise turn of less than π.
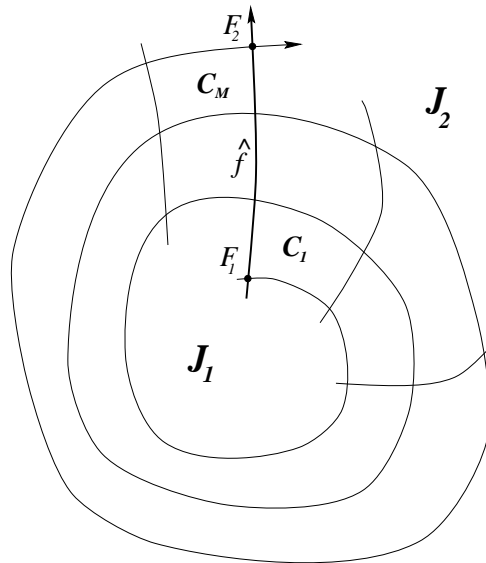

Fig. 3.

Let $\hat{e}$ be the portion of e starting with e_1^b and ending with e_M^t. Orient it from e_1^b towards e_M^t. Then $\hat{e}$ crosses $\hat{f}$ from left to right at every crossing. See Fig. 3.

Since I contains no lens, for any orientation of two elements $x, y \in S$, every crossing of these two curves are of the same type, i.e., either x crosses y from left to right at every crossing, or x crosses from right to left at every crossing. In particular, we can pick an orientation of $x \in S$ such that at every crossing it crosses $\hat{f}$ from left to right.

Let $S_1 \subset S$ denote the set of all elements of S that do not cross $\hat{f}$. Let $S_2 \subset S$ (and $S_3 \subset S$) denote the set of all elements starting at a point of j_1 and ending at a point of j_2 (starting at a point of j_2 and ending at a point of j_1, respectively).

We will modify the drawing by re-routing the elements of S_2 so as to reduce the number of crossings. Notice that, since there is no lens in I, the intersection points of x and $\hat{f}$ follow each other in the same order on both x and $\hat{f}$.

Define a binary relation on S_2 as follows. For any $x, y \in S_2$, we say that y *precedes* x (and write $y \prec x$), if x and $\hat{f}$ have two consecutive crossings, X and X', such that y does not intersect the portion of $\hat{f}$ between X and X'.

Claim. The relation $\prec$ is a partial ordering on S_2.

Proof. Suppose that $y \prec x$. The union of the portions of x and $\hat{f}$ between X and X' divides I into two pieces, separating j_1 from j_2. Since y does not cross the portion of $\hat{f}$ between X and X', it must cross the portion of x between X and X' from right to left (see Fig. 4). But then at every other crossing y has to cross x from right to left. This shows that $\prec$ is antisymmetric, because assuming that $x \prec y$, the same argument would show that at each of their crossings y must cross x from left to right, a contradiction.

To show that $\prec$ is transitive, suppose that $z \prec y$ and $y \prec x$. Let Y and Y' be two consecutive crossings between y and $\hat{f}$ such that z does not intersect the portion of $\hat{f}$ between Y and Y'. If the YY' portion of $\hat{f}$ contains the XX' portion, then z does not cross the XX' portion, hence we are done: $z \prec x$. Thus, we can assume that the XX' and YY' portions are disjoint. Suppose without loss of generality that along $\hat{f}$ the order of these four points is X, X', Y, Y'. Let $Y_1, Y_2, \ldots, Y_k$ be all crossings of $\hat{f}$ and y, between the points X' and Y', where $Y_{k-1} = Y, Y_k = Y'$. By our assumptions, every portion $Y_{i-1}Y_i$ of $\hat{f}$ contains at least one crossing with x and at most one crossing with z. Therefore, the XY' portion of $\hat{f}$ contains at least $k+1$ crossings with x and at most $k-1$ crossings with z. Thus, we have $z \prec x$, concluding the proof of the Claim. $\square$

For any $x \in S_2$, define $\mathrm{rank}(x) = |x \cap \hat{f}|$. For any *edge* g of G, let

$$\mathrm{rank}(g) = \{\mathrm{rank}(x) \mid x \in S_2, x \subset g\},$$

and let $\mathrm{rank}(G) = \bigcup_{g \in E(G)} \mathrm{rank}(g)$.

In the proof of the Claim, we showed that if x has two consecutive crossings with $\hat{f}$ with the property that y does not intersect the XX' portion of $\hat{f}$, then y must cross the XX' portion of x. Therefore, if $|\mathrm{rank}(x) - \mathrm{rank}(y)| \geq 2$ for some $x, y \in S_2$, then x and y cross each other. If x and y belong to the same edge of G, then they cannot cross, so in this case we have $|\mathrm{rank}(x) - \mathrm{rank}(y)| \leq 1$. Therefore, for any edge g of G, the set $\mathrm{rank}(g)$ is either empty, or it consists of one integer or two consecutive integers.

Since we have $\mathrm{rank}(\hat{e}) = M' + 1 > 5m$, but $|\mathrm{rank}(G)| \leq 2m$, there is an integer L, $5m > L > 3$ such that $L, L-1, L-2 \notin \mathrm{rank}(G)$.

Now let $S_2 = S^l \cup S^h$ (where l and h stand for 'low' and 'high,' resp.) such that

$$S^l = \{x \in S_2 \mid \mathrm{rank}(x) < L - 2\}, \quad S^h = \{x \in S_2 \mid \mathrm{rank}(x) > L\}.$$

Let w be a *minimal* element of S^h with respect to the partial ordering $\prec$. Let $W_1, W_2, \ldots, W_k$ denote the crossings of w and $\hat{f}$, in this order. By the minimality of w, every element of S^h intersects $\hat{f}$ between any W_i and W_{i+1}.
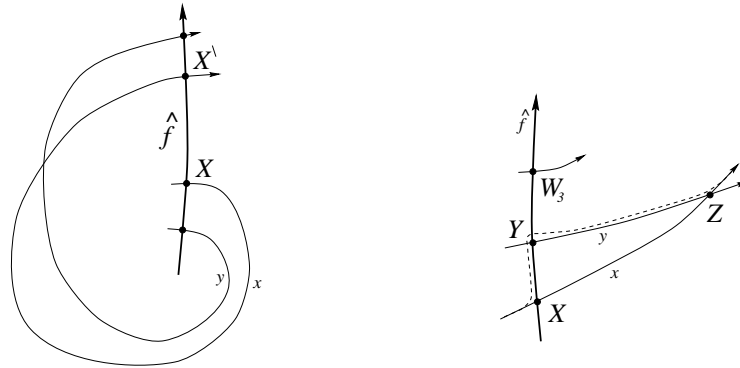
Fig. 4.

By 'shifting' the 'bad' crossings further, we will modify the elements of S^l so that none of them will cross the portion of $\hat{f}$ between F_1 and W_3. Suppose that for some $x \in S^l$, $y \in S^h$, X (resp. Y) is a crossing of x (resp. y) and $\hat{f}$, and no element of S_2 crosses the XY portion of $\hat{f}$, and F_1, X, Y, W_3 follow each other in this order on $\hat{f}$. There are at least $L-3$ crossings of y and $\hat{f}$, which come after Y. Since x and $\hat{f}$ cross at most $L-4$ times, there exists a crossing Z of x and y, which comes after X along x and after Y along y. Let X' be a point on x slightly before X, let Y' be a point of $\hat{f}$ slightly after Y, and let Z' be a point of x slightly after Z. Replace the $X'Z'$ portion of x by a curve running from X' to Y' very close to the XY portion of $\hat{f}$, and from Y' to Z' running very close to the YZ portion of y. This is called an *elementary flip*. See Fig. 4. Any element of S_2 intersects the XZ portion of x and the YZ portion of y the same number of times, therefore, x intersects every other element of S_2 precisely the same number of times as before the the elementary flip.

Do as many elementary flips as possible. When we get stuck, no element of S^l crosses the portion of $\hat{f}$ between F_1 and W_3.

Let P be a point of $\hat{f}$. For any $b \in S_2$, let $B(P)$ and $B'(P)$ be those two crossing points of $\hat{f}$ and b, which follow immediately after P along $\hat{f}$ (if two such point exist). The portions of b and $\hat{f}$ between $B(P)$ and $B'(P)$ divide I into two components. Let $I_b(P)$ denote the one containing j_2. Let $I_2(P) = \bigcup_{b \in S_2} I_b(P)$ and let $I(P) = I \setminus I_2(P)$.

Suppose without loss of generality that $\hat{f}$ is a straight-line segment. Let T be a (very thin) rectangle, whose left side, T^l, coincides with the segment F_1W_3, and whose right side T^r is very close to T^l. For any $y \in S^h$, let $Y_1^l, Y_2^l, \dots, Y_k^l$ (resp. $Y_1^r, Y_2^r, \dots, Y_k^r$) be the intersections of y and T^l (resp. T^r). We can assume that the segments $Y_i^l Y_i^r$ are horizontal. Delete $Y_1^r Y_k^l$ from y.

Connect Y_i^r to Y_{i+1}^l by a straight-line segment (inside T), and connect Y_{i+1}^l with Y_{i+1}^r by a curve in $I(Y_{i+1}^l)$ ($1 \le i \le k-1$), running very close to its boundary. In this way, we obtain a weak realization of (G, R). Any $x \in S^h$ intersects $\hat{f}$ one time less than previously. It is not hard to see that the number of crossings between x and $y \in S$ remained the same or decreased. Therefore,

the number of crossings in the above realization is smaller than in the original drawing, a contradiction.   □

Lemma 3.1 combined with the next statement shows that it is sufficient to study empty paths of type 2.

Lemma 3.3. *Let $C_1, C_2, \ldots, C_M$ be an empty (e, f)-path of four-cells. Let $\bar{e} = e_1^t \cup e_2^t \cup \cdots \cup e_M^t$, and suppose that $\bar{e}$ and f form an empty $(\bar{e}, f)$-path of type 1, whose length is $L \leq M$.*
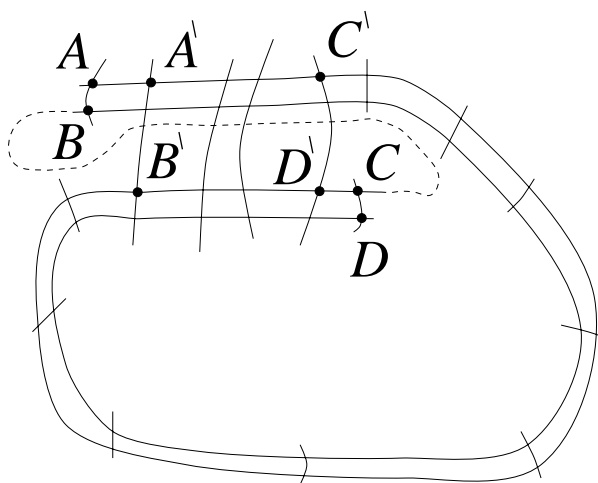*Then there is an empty (e, f)-path of type 2, whose length is L.*

Fig. 5.

Proof. We can assume that $C_1, C_2, \ldots, C_M$ is a path of type 1, otherwise there is nothing to prove. Let A, B, C, and D denote the same as in Figure 1. Then these points follow each other along e in the order A, C, B, D.

Let $\overline{C}_1, \overline{C}_2, \ldots, \overline{C}_L$ be the four-cells of the empty $(\bar{e}, f)$-path and let $\bar{e}_i^t, \overline{f}_i^r, \bar{e}_i^b$ and $\overline{f}_i^l$ denote the sides of C_i'. Further, let A' be the common endpoint of $\overline{f}_1^l$ and $\bar{e}_1^t$, let $B' = \overline{f}_1^l \cap \bar{e}_1^b$, $C' = \overline{f}_L^r \cap \bar{e}_L^t$, and $D' = \overline{f}_L^r \cap \bar{e}_L^b$. Then A', C', B', and D' follow each other in this order along e, and they all lie between A and C.

Let I denote the region bounded by $\bar{e}_1^b, \overline{f}_1^r$, and by the portion of e between $\overline{f}_1^r \cap \bar{e}_1^t$ and B' (cf. Figure 5). Clearly, C and D are in the exterior and in the interior of I, respectively. Since e cannot cross itself, the portion e'' of e between C and B must intersect $\overline{f}_1^r$. It follows that e'' intersects $\overline{f}_L^l, \overline{f}_{L-1}^l, \ldots, \overline{f}_1^l$ in this order. (If e'' enters some $\overline{C}_j$ through the side $\overline{f}_{j+1}^l$, then it must leave $\overline{C}_j$ through the opposite side, $\overline{f}_j^l$, otherwise we would obtain an empty lense contradicting Lemma 2.1. Thus, f and the portion of e between B' and C form an empty (e, f)-path of type 2, whose length is L.   □

4   Empty Paths of Type 2 and the Proof of Theorem 2

Let e, f be two portions of edges of G. Suppose that they form an empty (e, f)-path P of type 2, consisting of M four-cells, $C_1, C_2, \ldots, C_M$. We use the notation on Figure 1. Now the arcs $e_1^t, e_2^t, \ldots, e_M^t, e_M^b, \ldots, e_1^b$ follow each other on e in this order. Let I denote the union of the cells C_i $(1 \le i \le M)$, which is a curvilinear quadrilateral, whose vertices are A, B, C, and D, in clockwise order. Furthermore, let J denote the region bounded by f_M^r and the portion of e between C and D.

Suppose that there is a portion of an edge g which intersects $f_1^l, f_2^l, \ldots f_M^l$ in this order and ends at a vertex $v \in V(G)$ lying in J. If g does not intersect any of the sides e_i^t, e_i^b, $1 \le i \le M$, then we say that g (with vertex v) *cuts all the way through* the path P. (See Fig 6.) The CD portion of e is said to be the *top* of the path P, the BC and AD portions are the *sides* of P. Note that there is no piece of e inside I or J.
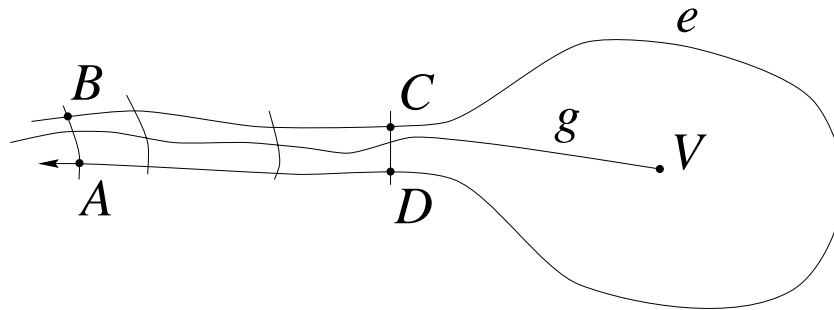

Fig. 6.

Lemma 4.1. *Let e, f be two portions of edges of G that form an empty (e, f)-path P of type 2.*
    Then there exists a portion of an edge g which cuts all the way through P.

Proof. We can and will assume without loss of generality that P is *minimal* in the sense that there is no other (e', f)-path P' of type 2, consisting of M four-cells, such that $I' \cup J'$ is strictly contained in $I \cup J$. Obviously, every edge that cuts all the way through P' also cuts all the way through P.

Fix an edge g of G. The boundary of $I \cup J$ may cut g into several portions. Any portion of g which enters $I \cup J$ through the side f_1^l must leave $I \cup J$ at some point, otherwise it meets the requirements in the lemma. Indeed, such a portion can only end at a vertex in J, because I is empty, and it must cross the arcs $f_1^l, f_2^l, \ldots, \ldots, f_M^l, f_M^r$ in this order, because otherwise it would create an empty lens with f.

Suppose there is a portion g' of g entering and leaving $I \cup J$ through the side f_1^l. Using again that all four-cells in P are empty, we obtain that g' must first cross the arcs $f_1^l, f_2^l, \ldots, f_M^l, f_M^r$, in this order, then turn back and cross them another time in the opposite order. Hence, there is a (g', f)-path P' of length M inside $I \cup J$, contradicting the minimality of P.

Therefore, every portion of an edge which enters $I \cup J$ through f_1^l must leave it through the portion of e between B and A. Replacing this portion of e by a curve running very close to f_1^l, but not entering I, we obtain another weak realization of (G, R) with a smaller number of crossings. This contradicts our assumption that the initial realization was minimal. $\square$

Before we turn to the proof of Theorem 2, we have to introduce another notion. For any portion e' of an edge e, an (e', f)-path of empty four-cells is called a *weak (e, f)-path*.

Proof. Consider a weak realization of (G, R), in which the number of crossings is minimum, and suppose that this number is at least $(2n)^{12m+24} = (8n^3)^{2m+4}$. Find two edges, e and f, that cross each other $K \geq (8n^3)^{2m+3}$ times.

For every $i = 1, 2, \ldots, 2m + 1$, we construct the following objects:

(1) A subsegment $e_i \subseteq e_{i-1}$, where $e_0 = e$;

(2) an (e_i, f)-path P_i of length at least $(8n^3)^{2m-2i+3}$;

(3) a set of *charged pairs*, (g_i, v_i), where each pair consists of an edge g_i of G and one of its endpoints, v_i. Initially, there is no charged pair. In STEP i, we define g_i with endpoint v_i such that (g_i, v_i) will cut all the way through an (e, f)-path $P(g_i, v_i)$ of length $(8n^3)^{2m-2i+3}$. $P(g_i, v_i) \subset P_i$.

(4) $P'(g_i, v_i)$ will be a weak (e, f)-path of length $(8n^3)^{2m-2i+3}$ such that $P(g_i, v_i) \subset P'(g_i, v_i) \subset P_i$, and such that $P'(g_i, v_i)$ is *maximal* with this property.

STEP i: For $i = 1$, let $e_0^+ = e_0^- = e_0 = e$. For $i > 1$, P_{i-1} has two sides, call them e_{i-1}^+ and e_{i-1}^-.

The arc e_{i-1}^- intersects f at least $(8n^3)^{2m-2i+5}$ times, therefore, by the previous lemmas, there is a (e_{i-1}^-, f)-path P^- of length $(8n^3)^{2m-2i+3}$. Since e_{i-1}^+ and e_{i-1}^- are two sides of an (e_{i-1}, f)-path, there is also also a corresponding (e_{i-1}^+, f)-path P^- of length $(8n^3)^{2m-2i+3}$, and either $P^- \subset P^+$ or $P^+ \subset P^-$. Suppose that $P^- \subset P^+$, the other case can be treated analogously. Let $e_i = e_{i-1}^+$, $P_i = P^+$. Let $\hat{f} \subset f$ be a side of one of the four-cells of P_i. Color each point of $\hat{f}$ red, if it is inside $P'(g_j, v_j)$ for some $j < i$, and color the other points blue. The red points form some disjoint arcs of $\hat{f}$, whose endpoints are intersections of $\hat{f}$ and e.

Let F be one of the intersections of e_i^- and $\hat{f}$, and let F' be the intersection of e and $\hat{f}$, closest to F. By our construction, the portion of $\hat{f}$ between F and F' is blue. Suppose without loss of generality that at F e crosses f from left to right.

If the edge e crosses f from right to left at F', then we have an (e, f)-path P of type 2 inside P_i, whose length is $(8n^3)^{2m-2i+3}$. Therefore, there is a pair (g_i, v_i) cutting all the way through P. Since FF' was a blue arc, (g_i, v_i) has not been charged before. Now charge (g_i, v_i) and let $P(g_i, v_i) = P$. Let $P'(g_i, v_i)$ be a weak (e, f)-path of length $(8n^3)^{2m-2i+3}$ such that $P(g_i, v_i) \subset P'(g_i, v_i) \subset P_i$, and such that $P'(g_i, v_i)$ is *maximal* with this property. Go to STEP $i + 1$.

On the other hand, if e crosses f from left to right at F', then we have an (e, f)-path P of type 1 inside P_i, whose length $(8n^3)^{2m-2i+3}$. But then the FF'

arcs of e and f together separate the two endpoints of e. (See also the proof of Lemma 3.3.) Let F'' be a point of f very close to F, such that F is between F' and F''. The FF' portions of e and f together separate one of the endpoints, v, from F''. Since e cannot cross itself, there is also an (e, f)-path P of type 2 inside P^- such that (e, v) cuts all the way through it. Clearly, (e, v) was not charged before since the FF' arc of $\hat{f}$ is blue. Now charge $(g_i, v_i) = (e, v)$, and let $P(g_i, v_i) = P$. Let $P'(g_i, v_i)$ be a weak (e, f)-path of length $(8n^3)^{2m-2i+3}$ such that $P(g_i, v_i) \subset P'(g_i, v_i) \subset P_i$, and such that $P'(g_i, v_i)$ is *maximal* with this property. Go to STEP $i + 1$.
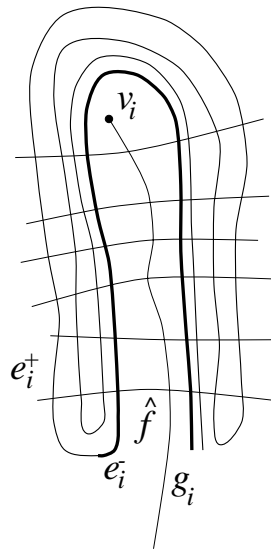

Fig. 7.

The algorithm will terminate after $2m + 1$ steps. At each step we charge a new pair (e, v). This is a contradiction, because the total number of pairs that can be charged is $2m$. Therefore, a minimal weak realization has at most $(2n)^{12m+24} \le (4m)^{12m+24}$ crossings. This concludes the proof of Theorem 2. □

References

[B59]   S. Benzer: On the topology of the genetic fine structure, *Proceedings of the National Academy of Sciences of the United States of America* **45** (1959), 1607–1620.

[CGP98a] Z.-Z. Chen, M. Grigni, and C. H. Papadimitriou, Planar map graphs, in: *STOC '98*, ACM, 1998, 514–523.

[CGP98b] Z.-Z. Chen, M. Grigni, and C. H. Papadimitriou, Planar topological inference, (Japanese) in: *Algorithms and Theory of Computing (Kyoto, 1998)* *Sūrikaisekikenkyūsho Kōkyūroku* **1041** (1998), 1–8.

[CHK99]   Z.-Z. Chen, X. He, and M.-Y. Kao, Nonplanar topological inference and political-map graphs, in: *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, MD, 1999),* ACM, New York, 1999, 195–204.

[DETT99] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis, *Graph Drawing,* Prentice Hall, Upper Saddle River, NJ, 1999.

[E93]   M. Egenhofer, A model for detailed binary topological relationships, *Geomatica* **47** (1993), 261-273.

[EF91]   M. Egenhofer and R. Franzosa, Point-set topological spatial relations, *International Journal of Geographical Information Systems* **5** (1991), 161-174.

[ES93]   M. Egenhofer and J. Sharma, Assessing the consistency of complete and incomplete topological information, *Geographical Systems* **1** (1993), 47–68.

[EET76]   G. Ehrlich, S. Even, and R. E. Tarjan, Intersection graphs of curves in the plane, *Journal of Combinatorial Theory, Series B* **21** (1976), 8–20.

[EHP00]   P. Erdős, A. Hajnal, and J. Pach, A Ramsey-type theorem for bipartite graphs, *Geombinatorics* **10** (2000), 64–68.

[EPL72]   S. Even, A. Pnueli, and A. Lempel, Permutation graphs and Transitive graphs, *Journal of Association for Computing Machinery* **19** (1972), 400–411.

[G80]   M. C. Golumbic, *Algorithmic Graph Theory and Perfect Graphs,* Academic Press, New York, 1980.

[G78]   R. L. Graham: Problem, in: *Combinatorics, Vol. II (A. Hajnal and V. T. Sós, eds.)*, North-Holland Publishing Company, Amsterdam, 1978, 1195.

[HT74]   J. Hopcroft and R. E. Tarjan, Efficient planarity testing, *J. ACM* **21** (1974), 549–568.

[K83]   J. Kratochvíl, String graphs, in: *Graphs and Other Combinatorial Topics (Prague, 1982),* Teubner-Texte Math. **59**, Teubner, Leipzig, 1983, 168–172.

[K91a]   J. Kratochvíl, String graphs I: The number of critical nonstring graphs is infinite, *Journal of Combinatorial Theory, Series B* **52** (1991), 53–66.

[K91b]   J. Kratochvíl, String graphs II: Recognizing string graphs is NP-hard, *Journal of Combinatorial Theory, Series B* **52** (1991), 67–78.

[K98]   J. Kratochvíl, Crossing number of abstract topological graphs, in: *Graph drawing (Montreal, QC, 1998), Lecture Notes in Comput. Sci.* **1547**, Springer, Berlin, 1998, 238–245.

[KLN91]   J. Kratochvíl, A. Lubiw, and J. Nešetřil, Noncrossing subgraphs in topological layouts, *SIAM J. Discrete Math.* **4** (1991), 223–244.

[KM89]   J. Kratochvíl and J. Matoušek, NP-hardness results for intersection graphs, *Comment. Math. Univ. Carolin.* **30** (1989), 761–773.

[KM91]   J. Kratochvíl and J. Matoušek, String graphs requiring exponential representations, *Journal of Combinatorial Theory, Series B* **53** (1991), 1–4.

[KM94]   J. Kratochvíl and J. Matoušek, Intersection graphs of segments, *Journal of Combinatorial Theory, Series B* **62** (1994), 289–315.

[LLR95]   N. Linial, E. London, and Y. Rabinovich, The geometry of graphs and some of its algorithmic applications, *Combinatorica* **15** (1995), 215–245.

[MP93]   M. Middendorf and F. Pfeiffer, Weakly transitive orientations, Hasse diagrams and string graphs, in: *Graph Theory and Combinatorics (Marseille-Luminy, 1990), Discrete Math.* **111** (1993), 393–400.

[PS01]   J. Pach and J. Solymosi, Crossing patterns of segments, *Journal of Combinatorial Theory, Ser. A,* to appear.

[R99]   S. Rao, Small distortion and volume preserving embeddings for planar and Euclidean metrics, in: *Proceedings of the Fifteenth Annual Symposium on Computational Geometry (Miami Beach, FL, 1999),* ACM, New York, 1999, 300–306.

[SS01]   M. Schaefer and D. Stefankovič, Decidability of string graphs, *STOC 01*, to appear.

[S66]   F. W. Sinden, Topology of thin film RC circuits, *Bell System Technological Journal* (1966), 1639-1662.

[SP92]   T. R. Smith and K. K. Park, Algebraic approach to spatial reasoning, *International Journal of Geographical Information Systems* **6** (1992), 177–192.