

Recognizing string graphs is decidable

János Pach*

Courant Institute, NYU
and Hungarian Academy of Sciences

Géza Tóth†

Massachusetts Institute of Technology
and Hungarian Academy of Sciences

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

*Supported by NSF grant CR-00-86013, PSC-CUNY Research Award 63382-0032 and OTKA-T-032452.

†Supported by NSF grant DMS-99-70071, OTKA-T-032452.

[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 $\text{ST}(G)$ denote the minimum number of crossings in a string representation of G , and let

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

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 $\text{ST}(n) \leq (2n)^{24n^2+48}$. On the other hand, it was shown by J. Kratochvíl and J. Matoušek [KM91] that $\text{ST}(n) \geq 2^{cn}$ 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 $\text{CR}(G, R)$. Note that the usual *crossing number* of G is equal to $\text{CR}(G, \binom{E(G)}{2})$. Let

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

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

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

Kratochvíl [K98] called it an “astonishing and challenging fact that so far there is no recursive” upper bound known on $\text{ST}(n)$. 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].

Shortly after, an alternative proof of the above results with slightly better bounds was found independently by Schaefer and Stefankovič [SS01].

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 *lens*.

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

Proof: By symmetry, it is sufficient to show that every lens contains a vertex. Suppose that there is a lens that does not contain any vertex of G . Then, there is such a lens which is minimal by containment. 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 lens must also cross the other one. Therefore, replacing the portion of f between A and B by an arc running outside the lens and very close to e , we would reduce the number of crossings in the drawing, contradicting its minimality. \square

Deleting from the plane any two arcs, e and f , the plane falls into a number of connected components, called *cells with respect to the pair (e, f)* or, simply, *cells*, whenever it will be clear from the context what e and f are. At most 4 of these cells contain an 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 $K \geq 16n^3$ and $n \geq 10$. Then e and f determine $M \geq K/(8n^3)$ empty four-cells, C_1, C_2, \dots, 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 lens 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. Since H' is a subtree of H , it has at most n leaves. This implies that H' contains a path with at least $M = (K' - 1)/(2n - 3) - 1 > K/(8n^3)$ vertices, each of degree 2. The sequence of four-cells corresponding to the vertices of this path meet the requirements of the lemma. \square




Figure 1. *An empty (e, f) -path of four-cells.*

A sequence C_1, C_2, \dots, 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. 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*, respectively, suggesting that on our pictures the boundary pieces belonging to e are “horizontal”; see Figure 1.) For every $1 \leq i < M$, 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 $D = f_M^r \cap e_M^b$.

In order to classify the empty paths, orient arbitrarily the edges of G . We distinguish two different types of the empty paths (see Fig. 2), and study them separately in the next two sections.

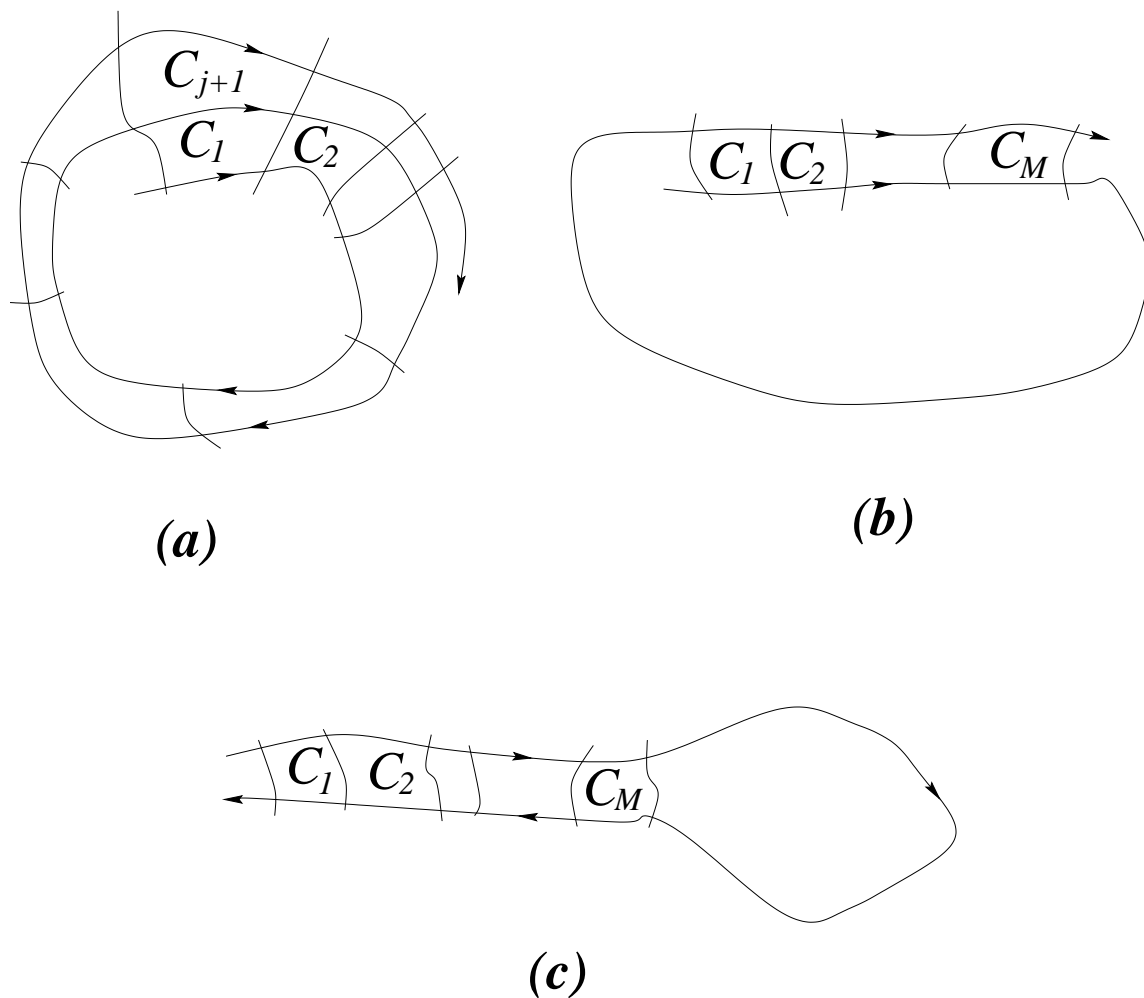


Figure 2. The empty paths (a) and (b) are of type 1, while (c) is of type 2.

Type 1: For every i , the two sides of the cell C_i belonging 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 belonging to e have *opposite* orientations. See Fig. 2 (c).

Observe that the type of a path is independent of the orientation of the edges. In the sequel,

we use some other auxiliary orientations of the arcs, which have nothing to do with the original orientation.

Suppose that C_1, C_2, \dots, C_M is an empty (e, f) -path of four-cells of type 1, such that e_1^t and e_1^b do not coincide with any side of C_i , $1 < i \leq M$. Then we say that C_1, C_2, \dots, C_M is an empty (e, f) -path of four-cells of *type 1b* (see Fig 2 (b)). A path of type 1 but not of type 1b is said to be of *type 1a*.

3 Empty paths of type 1

The aim of this section is to show that if a drawing of G has a sufficiently long empty path of type 1, then it also contains a long empty path of type 2. Therefore, it will be sufficient to study empty paths of the latter type.

Most of the section is devoted to the proof of the following

Lemma 3.1 *Let C_1, C_2, \dots, C_M be an empty (e, f) -path of four-cells of type 1. Then for $M' = \lfloor M/(5m) \rfloor$, $C_1, C_2, \dots, C_{M'}$ is an empty (e, f) -path of four-cells of type 1b.*

Proof: It is sufficient to show that e_1^t and e_1^b cannot coincide with any side of C_i , $1 < i \leq M/(5m)$. 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 $\hat{f} = f_1^l \cup f_{j+1}^l \cup \dots \cup f_M^l$, which is a segment of f , and orient it from f_1^l towards f_M^l . (This orientation has nothing to do with the original orientation.) 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 $I = C_1 \cup C_2 \cup \dots \cup C_M$. Furthermore, let J_1 denote the region bounded by $j_1 = e_1^b \cup e_2^b \cup \dots \cup e_j^b \cup f_1^l$ which does not contain I , and let J_2 be the region bounded by $j_2 = e_M^t \cup e_{M-1}^t \cup \dots \cup e_{M-j+1}^t \cup f_M^l$ 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, \bar{e} and \bar{f} , crossing at some point X , we say that \bar{e} *crosses \bar{f} from left to right* if the direction of \bar{e} at X can be obtained from the direction of \bar{f} at X by a clockwise turn of less than π .

Let \hat{e} be the portion of e starting with e_1^b and ending with e_M^t . Assume without loss of generality that \hat{e} (and e) is oriented from e_1^b towards e_M^t . Then \hat{e} crosses \hat{f} from left to right at every crossing. See Fig. 3.

Let $S_1 \subset S$ denote the set of all elements of S that do not cross \hat{f} . Orient each element of S_1 such that they start at a point of j_2 and end at a point of j_1 .

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 y from *right to left* at every crossing. In particular, we can pick an orientation of each $x \in S \setminus S_1$ such that at every crossing it crosses \hat{f} from left to right.

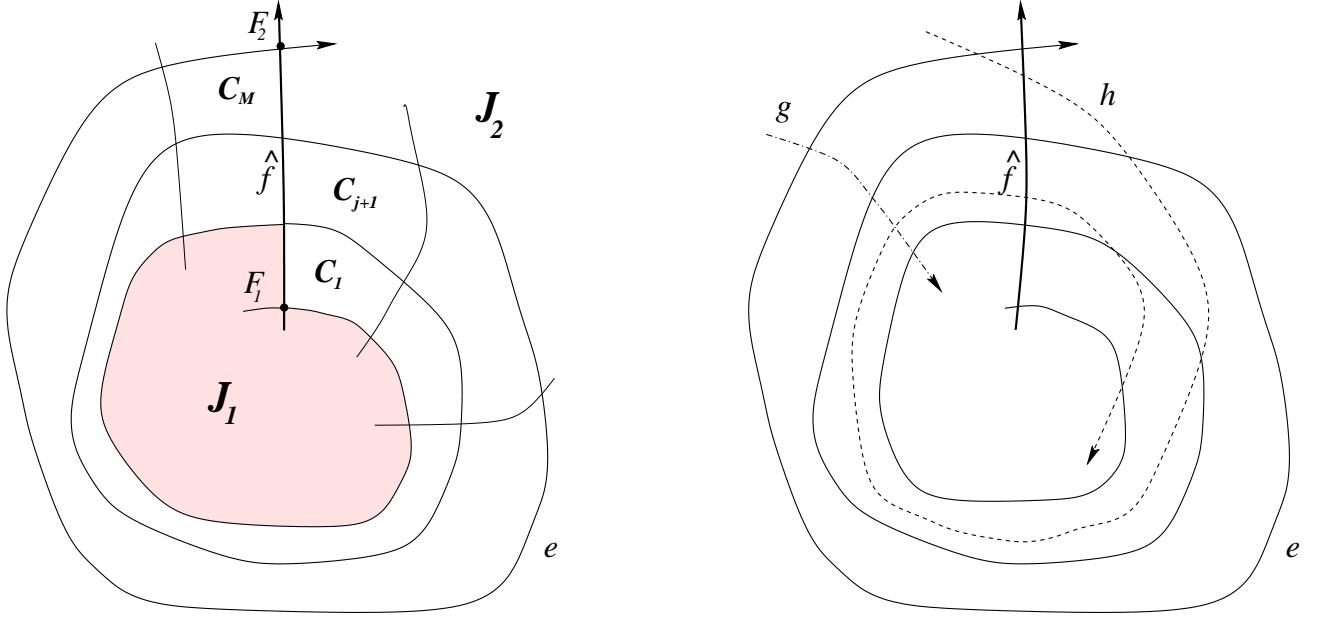


Figure 3. An empty path of type 1, where $e_1^t = e_{j+1}^b$, $g \in S_1$, $e \in S_2$, $h \in S_3$.

Let $S_2 \subset S \setminus S_1$ (and $S_3 \subset S \setminus S_1$) 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).

Perhaps not all elements of S_1 and S_3 cross the elements of S_2 , but if they do, they always cross from left to right.

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

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

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

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 $|\text{rank}(x) - \text{rank}(y)| \geq 2$ 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 $|\text{rank}(x) - \text{rank}(y)| \leq 1$. Therefore, for any edge g of G , the set $\text{rank}(g)$ is either empty, or it consists of one integer or two consecutive integers.

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

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

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

Let w be a *minimal* element of S^h with respect to the partial ordering \prec . Let W_1, W_2, \dots, 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} .

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} , 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 between y and \hat{f} that 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 elementary flip.

Similarly, every element of S_1 or S_3 intersects the elements of S_2 from left to right, therefore, any element of $S_1 \cup S_3$ intersects x the same number of times before and after the elementary flip. See Fig. 4.




Figure 4. The partial ordering \prec , $y \prec x$, and an 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 .

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), where Y_1^r is closest to F_1 . We can assume that the segments $Y_i^l Y_i^r$ are horizontal. Delete $Y_1^l Y_2^r$ from y and replace it by the $Y_1^r Y_2^l$ straight-line segment inside T (see Fig. 5). 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. \square

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




Figure 5. *Reducing the number of crossings.*

Proof: Let A, B, C , and D denote the same as in the definition of an (e, f) -path after the proof of Lemma 2.2 (see Fig. 1). These points follow each other along e in the order B, C, A, D .

Let $\bar{C}_1, \bar{C}_2, \dots, \bar{C}_L$ be the four-cells of the empty (\bar{e}, f) -path and let $\bar{e}_i^t, \bar{f}_i^r, \bar{e}_i^b$ and \bar{f}_i^l denote the sides of C_i^t . Further, let B' be the common endpoint of \bar{f}_1^r and \bar{e}_1^t , let $A' = \bar{f}_1^r \cap \bar{e}_1^b$, $C' = \bar{f}_L^l \cap \bar{e}_L^t$, and $D' = \bar{f}_L^l \cap \bar{e}_L^b$. The points B', C', A' , and D' follow each other in this order along e , and they all lie between B and C (see Fig. 6).

Let I denote the region bounded by \bar{f}_1^r , and by the portion of e between $\bar{f}_1^r \cap \bar{e}_1^t$ and A' (cf. Figure 6). Clearly, C and A are in the exterior and in the interior of I , respectively. Since e cannot cross itself, the portion e'' of e between C and A must intersect \bar{f}_1^r . It follows that e'' intersects $\bar{f}_L^l, \bar{f}_{L-1}^l, \dots, \bar{f}_1^l$ in this order. (If e'' enters some \bar{C}_j through the side \bar{f}_{j+1}^l , then it must leave \bar{C}_j through the opposite side, \bar{f}_j^l , otherwise we would obtain an empty lens contradicting Lemma 2.1.

Thus, f and the portion of e between A' and A form an empty (e, f) -path of type 2, whose length is L . \square




Figure 6. Proof of the existence of an (e, f) -path of type 2.

4 Empty paths of type 2 and the proof of Theorem 2

The idea of the proof of Theorem 2, to be presented in this section, is roughly the following. Starting with a pair of edges, e and f , that cross a large number of times, we will set up a procedure to find a long “nested” sequence of empty paths of type 2. These paths will give rise to distinct vertices. If the number of crossings between e and f is too large, then we obtain more than n vertices, and this contradiction will conclude the proof.

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, \dots, C_M . We use the notation on Figures 1 and 7. Now the arcs $e_1^t, e_2^t, \dots, e_M^t, e_M^b, \dots, e_1^b$ follow each other on e in this order. Let I denote the union of the cells C_i ($1 \leq i \leq M$), which is a curvilinear quadrilateral, whose vertices are A, B, C , and D , in clockwise order. Furthermore, let J denote the region, disjoint from I , 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 $f_1^l, f_2^l, \dots, f_M^l$ 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 , $1 \leq i \leq M$, then we say that g (with vertex v) *cuts all the way through* the path P . (See Fig 7.) 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 .




Figure 7. An edge g with a vertex v that cuts all the way through P .

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 $I \cup J$. (The definitions of I' and J' are analogous to the definition of I and J , respectively.) Obviously, every edge that cuts all the way through P' also cuts all the way through P .

If f_1^l is not crossed by any edge, we can replace the AB portion of e by a curve running very close to f_1^l . We would get a weak realization of G with fewer crossings, a contradiction. Thus, f_1^l is crossed by at least one edge.

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 $f_1^l, f_2^l, \dots, f_M^l, f_M^r$ 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 $f_1^l, f_2^l, \dots, f_M^l, f_M^r$, 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.

Definition 4.2. For any portion e' of an edge e , an (e', f) -path of empty four-cells is called a *weak (e, f) -path*.

Proof of Theorem 2: Consider a weak realization of (G, R) , in which the number of crossings is

minimum, and assume 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 $i = 1, 2, \dots, 2m + 1$, 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 type 2 and 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 no pair is charged. In STEP i , we define g_i with an endpoint v_i such that (g_i, v_i) cuts 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 type 2 and 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.




Figure 8. A charged pair (g_i, v_i) .

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 lemmas 3.1 and 3.3, there is an (e_{i-1}^-, f) -path P^- of type 2, whose length is $(8n^3)^{2m-2i+3}$. 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 $P^+ \subset P^-$ holds. 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} with 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 e crosses f at F from left to right.

If 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 (see Fig. 8). 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 1b inside P_i , whose length is $(8n^3)^{2m-2i+3}$. 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 $i + 1$.

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 $(2n)^{12m+24} \leq (4m)^{12m+24}$ crossings. This concludes the proof of Theorem 2. \square

Acknowledgement. An extended abstract of this paper was presented at Graph Drawing 2001, Vienna (see [PT01]). We are very grateful to Sean McGuiness for several very important remarks.

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) Sūrikaiseikikenkyūsho 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*, **96** (2001), 316-325.
- [PT01] J. Pach and G. Tóth, Recognizing string graphs is decidable, in: *Graph Drawing 2001 (P. Mutzel, ed.)*, *Lecture Notes in Computer Science* **2265**, Springer Verlag, Berlin 2002, 247-260.
- [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, in: *Proceedings of the 33rd Annual Symposium on the Theory of Computing (STOC 2001)*, 2001, 241–246.
- [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.