



Computing NodeTrix Representations of Clustered Graphs

*Giordano Da Lozzo Giuseppe Di Battista Fabrizio Frati
Maurizio Patrignani*

Department of Engineering, Roma Tre University, Italy

Abstract

NodeTrix representations are a popular way to visualize clustered graphs; they represent clusters as adjacency matrices and inter-cluster edges as curves connecting the matrix boundaries. We study the complexity of constructing NodeTrix representations focusing on planarity testing problems, and we show several NP-completeness results and some polynomial-time algorithms. Building on such algorithms we develop a JavaScript library for NodeTrix representations aimed at reducing the crossings between edges incident to the same matrix.

Submitted: November 2016	Reviewed: October 2017	Revised: November 2017	Accepted: December 2017	Final: December 2017
		Published: January 2018		
Article type: Regular paper		Communicated by: G. Liotta		

This research was supported in part by MIUR Project “MODE” under PRIN 20157EFM5C and by H2020-MSCA-RISE project 734922 “CONNECT”.

E-mail addresses: dalozzo@dia.uniroma3.it (Giordano Da Lozzo) gdb@dia.uniroma3.it (Giuseppe Di Battista) frati@dia.uniroma3.it (Fabrizio Frati) patrigna@dia.uniroma3.it (Maurizio Patrignani)

1 Introduction and Overview

NodeTrix representations have been introduced by Henry, Fekete, and McGuffin [19] in one of the most cited papers of the InfoVis conference [1]. A NodeTrix representation is a hybrid representation for the visualization of social networks where the node-link paradigm is used to visualize the overall structure of the network, within which adjacency matrices show communities.

Formally, a NodeTrix (NT for short) representation is defined as follows. A *flat clustered graph* (V, E, \mathcal{C}) is a graph (V, E) with a partition \mathcal{C} of V into sets V_1, \dots, V_k , called *clusters*, that can be defined according to the application needs. The word “flat” is used to underline that clusters are not arranged in a multi-level hierarchy (see, e.g., [11, 13] for two papers dealing with non-flat clustered graphs). An edge $(u, v) \in E$ with $u \in V_i$ and $v \in V_j$ is an *intra-cluster edge* if $i = j$ and is an *inter-cluster edge* if $i \neq j$.

In an *NT representation* the subgraphs induced by the clusters V_1, \dots, V_k are represented by non-overlapping symmetric adjacency matrices M_1, \dots, M_k , where M_i is drawn in the plane so that its boundary is a square Q_i with sides parallel to the coordinate axes. In the following, with a slight abuse of terminology, we call “matrix” both the abstract algebraic concept and the representation of an algebraic matrix in the plane. Thus, the matrices M_1, \dots, M_k convey the information about the intra-cluster edges of (V, E, \mathcal{C}) , while each inter-cluster edge (u, v) with $u \in V_i$ and $v \in V_j$ is represented by a curve connecting a point on Q_i with a point on Q_j , where the point on Q_i (on Q_j) belongs to the column or to the row of M_i (resp. of M_j) associated with u (resp. with v). Fig. 1 shows an NT representation of a flat clustered graph.




Figure 1: An NT representation of a flat clustered graph.

Several papers aimed at improving the readability of NT representations by reducing the number of crossings between inter-cluster edges. For this purpose, vertices can have duplicates in different matrices [18] or clusters can be computed so to have dense intra-cluster graphs and a planar inter-cluster graph [10].

In this paper we study the problem of automatically constructing an NT representation of a given flat clustered graph. This problem combines traditional

graph drawing issues, like the placement of a set of geometric objects in the plane (here the squares Q_1, \dots, Q_k) and the routing of the graph edges (here the inter-cluster edges), with a novel algorithmic challenge: To handle the degrees of freedom given by the choice of the *order* for the rows and the columns of the matrices and by the choice of the *sides* of the matrices to which the inter-cluster edges attach to. Indeed, the order of the rows and columns of a matrix M_i is arbitrary, as long as M_i is symmetric; further, an inter-cluster edge incident to M_i can arbitrarily exit M_i from four sides: left or right if it exits M_i from its associated row, or top or bottom if it exits M_i from its associated column.

When working on a new model for graph representations, the very first step is usually to study the complexity of testing if a graph admits a planar representation within that model. Hence, we deal with the problem of testing if a flat clustered graph admits a planar NT representation. An NT representation is *planar* if no inter-cluster edge e intersects any matrix M_i , except possibly at an end-point of e on Q_i , and no two inter-cluster edges e and e' cross each other, except possibly at a common end-point. The NODETRIX PLANARITY (NT PLANARITY for short) problem asks if a flat clustered graph admits a planar NT representation. Fig. 2 shows a planar NT representation of the flat clustered graph also represented in Fig. 1.

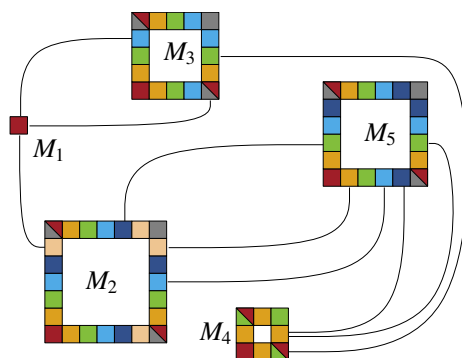


Figure 2: A planar NT representation of the flat clustered graph from Fig. 1.

Our findings show how tough the problem is (see Table 1). Namely, we show that NT PLANARITY is NP-complete and remains so even if the order of the rows and of the columns of the matrices is fixed (i.e., it is part of the input), or if the exit sides of the inter-cluster edges on the matrix boundaries are fixed. It is easy to show that NT PLANARITY becomes linear-time solvable if both the row-column order and the exit sides of the inter-cluster edges are fixed. But this is probably too restrictive for practical applications since all the degrees of freedom that are representation-specific are lost.

Motivated by such complexity results, we study a more constrained model that is still useful for practical applications. A *monotone NT representation* is an NT representation in which the matrices have prescribed positions and the inter-cluster edges are represented by xy -monotone curves inside the convex

		General Model		Monotone Model	
		Free Sides	Fixed Sides	Free Sides	Fixed Sides
R/C Order	Free	NPC [Th. 1]	NPC [Th. 2]	NPC [Th. 5]	NPC [Th. 6]
	Fixed	NPC [Th. 3]	P [Th. 4]	P [Th. 8]†	P [Th. 7]

Table 1: Complexity results for NT PLANARITY. The result marked † assumes that the number of clusters be constant.

hull of their incident matrices. We require that this convex hull, which might contain many edges, does not intersect any other matrix. We study this model for several reasons; we mention two of them here and some other later in the paper. First, in most of (although not in all) the available examples of NT representations the inter-cluster edges are represented by xy -monotone curves (see, e.g., the NodeTrix clips and prototype available online [2]). Second, we are interested in supporting a visualization system where the position of the matrices is decided by the user and the inter-cluster edges are automatically drawn with “few” crossings. Therefore, the crossings between inter-cluster edges not incident to a common matrix are somehow unavoidable, as they depend on the matrix positions selected by the users, and we are only interested in reducing the number of *local* crossings, that are the crossings between pairs of edges incident to the same matrix.

We say that an NT representation is *locally planar* if no two inter-cluster edges incident to the same matrix cross. Fig. 3 shows a locally-planar monotone NT representation of the flat clustered graph already depicted in Figs. 1 and 2

While testing if a flat clustered graph admits a monotone NT locally planar representation is NP-complete even if the sides are fixed (see Table 1), the problem becomes polynomial-time solvable in the reasonable scenario in which the number of matrices is constant, the order of the rows and columns is fixed, and the sides of the matrices to which the inter-cluster edges attach is variable.

Building on the insights for the last result, we developed a library for NT representations (a demo is available online [3]). The adopted techniques allow the user to move the matrices around while the layout of the inter-cluster edges is automatically recomputed; this happens without any slowdown of the interaction.

The paper is organized as follows. First, in Section 2 we provide some definitions and establish some notation. In Section 3 we deal with the problem of testing if a flat clustered graph admits a planar NT representation. Then, in Section 4 we study monotone NT representations. Further, in Section 5 we describe our library for NT representations. Conclusions and open problems are discussed in Section 6 where NT PLANARITY is related to graph drawing problems of theoretical interest.

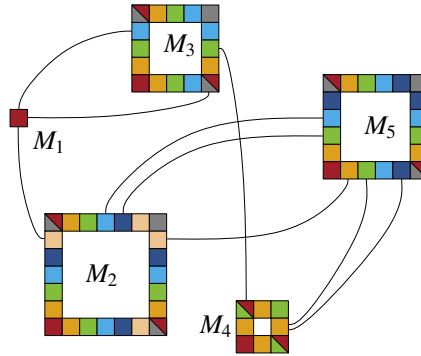


Figure 3: A locally-planar monotone NT representation of the flat clustered graph from Figs. 1 and 2. Note that the representation in Fig. 1 is monotone however not locally-planar, while the one in Fig. 2 is not monotone.

2 Definitions and Notation

In this section we establish formal definitions and notation. An *NT representation* consists of:

1. A *row-column order* σ_i for each cluster V_i , that is, a bijection $\sigma_i : V_i \leftrightarrow \{1, \dots, |V_i|\}$.
2. A *side assignment* s_i for each inter-cluster edge incident to V_i , that is, a mapping $s_i : \bigcup_{j \neq i} E_{i,j} \rightarrow \{T, B, L, R\}$, where $E_{i,j}$ is the set of inter-cluster edges between the clusters V_i and V_j (V_i and V_j are *adjacent* if $E_{i,j} \neq \emptyset$).
3. A *matrix*¹ M_i for each cluster V_i that is, a representation of V_i as a symmetric adjacency matrix such that:
 - (a) the boundary of M_i is a square Q_i with sides parallel to the coordinate axes; let $\min_x(Q_i)$ be the minimum x -coordinate of a point on Q_i ; $\min_y(Q_i)$, $\max_x(Q_i)$, and $\max_y(Q_i)$ are defined analogously;
 - (b) the left-to-right order of the columns and the top-to-bottom order of the rows in M_i is σ_i ; and
 - (c) any two distinct matrices are disjoint; if V_i has only one vertex, we often talk about the matrix representing that vertex, rather than the matrix representing V_i .
4. An *edge drawing* for each inter-cluster edge $e = (u, v)$ with $u \in V_i$ and $v \in V_j$, that is, a representation of e as a Jordan curve between two points p_u and p_v defined as follows. Let m_T^u be the mid-point of the line segment that is the intersection of the top side of Q_i with the column associated

¹Recall that we use the term “matrix” to indicate both the abstract algebraic concept and the representation of the algebraic matrix in the plane.

to u in M_i ; points m_B^u , m_L^u , and m_R^u are defined analogously. Then p_u coincides with m_T^u , m_B^u , m_L^u , or m_R^u if $s_i(e) = T$, $s_i(e) = B$, $s_i(e) = L$, or $s_i(e) = R$, respectively. Point p_v is defined analogously.

3 Testing NodeTrix Planarity

In this section we study the time complexity of testing NODETRIX PLANARITY for a flat clustered graph. We start with the following.

Theorem 1 NODETRIX PLANARITY is NP-complete even if at most three clusters contain more than one vertex.

In the following we prove Theorem 1. First, the membership in NP of NT PLANARITY will be proved in Lemma 4.

The Reduction

In order to establish the NP-hardness of NT PLANARITY, we give a reduction from the following NP-complete problem [8].

Name:	PARTITIONED 3-PAGE BOOK EMBEDDING
Input:	A graph (V, E) , together with a partition of E into three sets E_1, E_2 , and E_3 .
Question:	Does a total ordering \mathcal{O} of V exist such that the end-vertices of any two edges e and e' in the same set E_i do not alternate? Two edges e and e' <i>alternate</i> if an end-vertex of e' is between the two end-vertices of e in \mathcal{O} and vice versa.

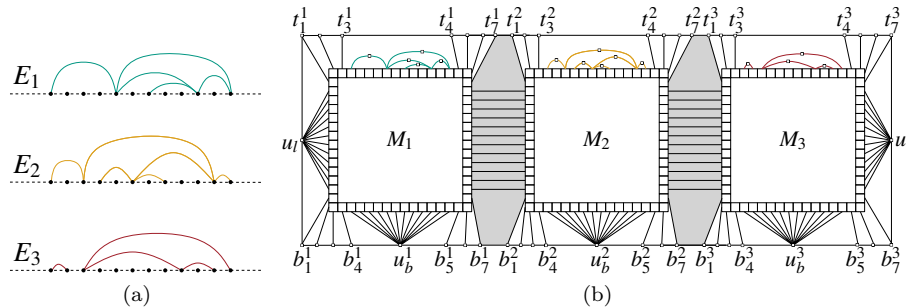


Figure 4: (a) An instance $(V, E = E_1 \cup E_2 \cup E_3)$ of PARTITIONED 3-PAGE BOOK EMBEDDING and (b) the corresponding instance (V', E', \mathcal{C}') of NT PLANARITY. The gray regions are $R^{1,2}$ and $R^{2,3}$.

We show how to construct in polynomial time an instance (V', E', C') of NT PLANARITY starting from an instance $(V, E = E_1 \cup E_2 \cup E_3)$ of PARTITIONED 3-PAGE BOOK EMBEDDING. Refer to Fig. 4. We define C' (and hence implicitly V') as follows.

- First, let V'_1 be a set whose elements are in bijection with V ; in particular, we denote by v'_1 the vertex in V'_1 that is in bijection with a vertex $v \in V$. Analogously, we define a set V'_2 (and a set V'_3) whose elements are in bijection with V and we denote by v'_2 (resp. by v'_3) the vertex in V'_2 (resp. in V'_3) that is in bijection with a vertex $v \in V$. Note that the vertices v'_1 , v'_2 , and v'_3 are in bijection with each other (due to their bijection with v). We now let the set C' contain three clusters V''_1 , V''_2 , and V''_3 , each with $|V| + 4$ vertices, defined as follows: $V''_1 = V'_1 \cup \{x_1, y_1, w_1, z_1\}$, $V''_2 = V'_2 \cup \{x_2, y_2, w_2, z_2\}$, and $V''_3 = V'_3 \cup \{x_3, y_3, w_3, z_3\}$;
- for every edge $e \in E$, the set C' contains a cluster $\{u'_e\}$; and
- the set C' contains clusters $\{u_b^1\}$, $\{u_b^2\}$, $\{u_b^3\}$, $\{u_l\}$, $\{u_r\}$ and, for $i = 1, \dots, 7$, clusters $\{t_i^1\}$, $\{t_i^2\}$, $\{t_i^3\}$, $\{b_i^1\}$, $\{b_i^2\}$, and $\{b_i^3\}$.

The set E' contains an arbitrary set of intra-cluster edges and the following inter-cluster edges.

- *equivalence edges*: edges (u'_e, r'_i) and (u'_e, s'_i) , for every edge $e = (r, s) \in E_i$;
- *bounding edges*: the edges of cycle $D = (u_l, t_1^1, \dots, t_7^1, t_1^2, \dots, t_7^2, t_1^3, \dots, t_7^3, u_r, b_7^3, b_6^3, b_5^3, u_b^3, b_4^3, \dots, b_1^3, b_7^2, b_6^2, b_5^2, u_b^2, b_4^2, \dots, b_1^2, b_7^1, b_6^1, b_5^1, u_b^1, b_4^1, \dots, b_1^1, u_l)$;
- *order-preserving edges*: edges (v'_1, v'_2) and (v'_2, v'_3) , for every vertex $v \in V$;
- *side-filling edges*: edges between u_l and every vertex in V'_1 , edges between u_r and every vertex in V'_3 , and edges between u_b^i and every vertex in V'_i for $i = 1, 2, 3$; and
- *corner edges*: for $i = 1, 2, 3$, edges (t_1^i, y_i) , (t_2^i, x_i) , (t_3^i, y_i) , (t_4^i, w_i) , (t_5^i, z_i) , (t_6^i, x_i) , (t_7^i, y_i) , (b_1^i, w_i) , (b_2^i, z_i) , (b_3^i, x_i) , (b_4^i, y_i) , (b_5^i, w_i) , (b_6^i, z_i) , and (b_7^i, w_i) .

The described construction can easily be performed in polynomial time. We now prove that an instance $(V, E = E_1 \cup E_2 \cup E_3)$ of PARTITIONED 3-PAGE BOOK EMBEDDING is positive if and only if the corresponding instance (V', E', C') of NT PLANARITY admits an NT planar representation.

The Necessity

The necessity is easy to prove. Suppose that $(V, E = E_1 \cup E_2 \cup E_3)$ is a positive instance of PARTITIONED 3-PAGE BOOK EMBEDDING and let \mathcal{O} be a total ordering of V such that the end-vertices of any two edges e and e' in the same set E_i do not alternate in \mathcal{O} . For $i = 1, 2, 3$, let σ_i be the total order of the

vertices in V_i'' such that x_i, y_i, w_i , and z_i are the first, second, last but one, and last vertex in σ_i , respectively, and the vertices in V_i' (which all follow x_i and y_i and precede w_i and z_i in σ_i) are ordered so that, for every $r'_i, s'_i \in V_i'$, vertex r'_i precedes vertex s'_i in σ_i if and only if r precedes s in \mathcal{O} . For $i = 1, 2, 3$, we represent V_i'' as a symmetric adjacency matrix M_i whose left-to-right order of the columns is σ_i ; every other cluster in \mathcal{C}' consists of a single vertex and we arbitrarily define a matrix for it. We embed these matrices in the plane (except for the matrices representing vertices u'_e with $e \in E$, which will be embedded later) so that no two of them overlap. The side assignment for the inter-cluster edges is as follows.

First, we assign:

- edges $(t_j^i, t_{j+1}^i), (t_7^i, t_1^{i+1}), (b_j^i, b_{j+1}^i), (b_4^i, u_b^i), (u_b^i, b_5^i)$, and (b_7^i, b_1^{i+1}) to the right (left) side of the matrix representing the first (resp. second) vertex in the pair;
- edges $(b_1^i, u_l), (u_l, t_1^i), (b_7^i, u_r)$, and (u_r, t_7^i) to the top (bottom) side of the matrix representing the first (resp. second) vertex in the pair;
- every other edge incident to t_j^i to the bottom side of the matrix representing t_j^i ;
- every other edge incident to b_j^i to the top side of the matrix representing b_j^i ;
- every other edge incident to u_b^i to the top side of the matrix representing u_b^i ;
- every other edge incident to u_l to the right side of the matrix representing u_l ; and
- every other edge incident to u_r to the left side of the matrix representing u_r .

By suitably routing the bounding edges, we get that M_1, M_2 , and M_3 we ensure that the cycle D bounds the outer face of the representation, with M_1, M_2 , and M_3 inside it.

Second, we assign the side-filling and order-preserving edges so that none of them is assigned to the top side of a matrix M_i . To achieve this, we assign: the side-filling edges incident to u_l and to u_r to the left side of M_1 and to the right side of M_3 , respectively; the side-filling edges incident to u_b^i to the bottom side of M_i ; and the order-preserving edges between vertices in V_i' and vertices in V_{i+1}' to the right side of M_i and to the left side of M_{i+1} . Route all these edges inside D . In particular, the order-preserving edges can be routed without crossings since the top-to-bottom order of the vertices in V_i' along the right side of M_i is $\sigma_i - \{x_i, y_i, w_i, z_i\}$, since the top-to-bottom order of the vertices in V_{i+1}' along the left side of M_{i+1} is $\sigma_{i+1} - \{x_{i+1}, y_{i+1}, w_{i+1}, z_{i+1}\}$, and since $\sigma_i - \{x_i, y_i, w_i, z_i\}$ and $\sigma_{i+1} - \{x_{i+1}, y_{i+1}, w_{i+1}, z_{i+1}\}$ are the same ordering according to the bijection between V_i' and V_{i+1}' .

Third, we assign the corner edges (t_2^i, x_i) , (t_3^i, y_i) , (t_4^i, w_i) , and (t_5^i, z_i) to the top side of M_i , (t_6^i, x_i) , (t_7^i, y_i) , and (b_7^i, w_i) to the right side of M_i , (b_3^i, x_i) , (b_4^i, y_i) , (b_5^i, w_i) , and (b_6^i, z_i) to the bottom side of M_i , and (t_1^i, y_i) , (b_1^i, w_i) , and (b_2^i, z_i) to the left side of M_i . The corner edges can be routed inside D without crossing the side-filling and order-preserving edges since by construction their end-vertices in V_i'' are the first, second, last but one, and last vertex in σ_i .

Fourth, embed the matrices representing u'_e with $e \in E_i$ in the region R_i delimited by edges (t_i^3, t_i^4) , (t_i^3, y_i) , and (t_i^4, w_i) , and by the top side of M_i . Assign the equivalence edges (u'_e, r'_i) and (u'_e, s'_i) to the top side of M_i and to any sides of the matrix representing u'_e ; route these edges in R_i . This can be done without introducing crossings, since the left-to-right order $\sigma_i - \{x_i, y_i, w_i, z_i\}$ of the vertices in V_i' along the top side of M_i is the same as the order \mathcal{O} of the vertices in V (according to the bijection between V_i' and V), hence no two pairs of edges $((p'_i, u'_e), (u'_e, q'_i))$ and $((r'_i, u'_f), (u'_f, s'_i))$ have alternating end-points along the top side of M_i given that edges (p, q) and (r, s) do not have alternating end-vertices in \mathcal{O} .

The Sufficiency

The proof of the sufficiency is more involved. Given a solution Γ for the instance (V', E', \mathcal{C}') of NT PLANARITY, we need to define an ordering \mathcal{O} for the vertex set V in the instance $(V, E = E_1 \cup E_2 \cup E_3)$ of PARTITIONED 3-PAGE BOOK EMBEDDING. In order to do that, we prove some claims and lemmata about the structure of Γ . For $i = 1, 2, 3$, let M_i be the matrix representing V_i'' in Γ , let Q_i be its boundary, and let σ_i be the left-to-right order of the vertices in V_i'' along the top side of Q_i . We have the following.

Claim 1 *The matrices M_1, M_2 , and M_3 , and the matrices representing vertices u'_e for all $e \in E$ are on the same side of D in Γ .*

Proof: The statement follows from the fact that the clusters V_i'' for $i = 1, 2, 3$ and $\{u'_e\}$ for $e \in E$ are connected by inter-cluster edges that are not bounding edges. Indeed, the clusters V_1'' , V_2'' , and V_3'' are connected by order-preserving edges and each cluster $\{u'_e\}$ is connected to a cluster V_i'' by two equivalence edges. Thus, if the matrices representing two of these clusters were on opposite sides of D in Γ , there would exist: (i) an order-preserving or an equivalence edge crossing a matrix representing a vertex in D or crossing a bounding edge, or (ii) a matrix among M_1, M_2 , and M_3 or a matrix representing a vertex u'_e with $e \in E$ overlapping a matrix representing a vertex in D or crossing a bounding edge. However, this would contradict the planarity of Γ . \square

We henceforth assume that M_1, M_2 , and M_3 , as well as the matrices representing the vertices u'_e for all $e \in E$, are *inside* D in Γ . Indeed, by Claim 1, these matrices are on the same side of D in Γ . If they are outside D , then the matrices representing vertices in D are all incident to an internal face f of Γ . Hence, changing the outer face of Γ to f ensures that M_1, M_2 , and M_3 and the

matrices representing the vertices u'_e for all $e \in E$ are inside D in Γ . This change of the outer face can be accomplished by rerouting the inter-cluster edges that are dual to a simple path in the dual of Γ from f to the outer face.

In the following we also assume that t_1^1 , t_2^1 , and t_3^1 are encountered in this order when traversing D in clockwise direction. This is not a loss of generality, up to a reflection of Γ .

Let Γ'' be the restriction of Γ to M_1 , M_2 , and M_3 , to the matrices representing the vertices t_j^i , b_j^i , u_b^i , u_l , and u_r , and to the bounding and corner edges of (V', E', C') . For $i = 1, 2$, let $R^{i,i+1}$ be the region of the plane delimited by the edges (t_7^i, y_i) , (b_7^i, w_i) , (t_7^i, t_1^{i+1}) , (b_7^i, b_1^{i+1}) , (t_1^{i+1}, y_{i+1}) , and (b_1^{i+1}, w_{i+1}) , by the boundaries of M_i and M_{i+1} , and by the boundaries of the matrices representing vertices t_7^i , b_7^i , t_1^{i+1} , and b_1^{i+1} . We have the following.

Claim 2 *Every order-preserving edge connecting a vertex in V'_i with a vertex in V'_{i+1} lies inside $R^{i,i+1}$.*

Proof: Assume that $i = 1$; the discussion with $i = 2$ is analogous.

We define three regions that partition the interior of Γ'' . Region R^1 is the minimal simple (i.e., without holes) region of the plane containing M_1 and containing the matrices and edges representing path $D^1 = (w_1, b_7^1, b_6^1, b_5^1, u_b^1, b_4^1, b_3^1, b_2^1, b_1^1, u_l, t_1^1, \dots, t_7^1, y_1)$. Region R^2 is the minimal simple region of the plane containing M_2 , M_3 , and containing the matrices and edges representing path $D^2 = (y_2, t_1^2, \dots, t_7^2, t_1^3, \dots, t_7^3, u_r, b_7^3, b_6^3, b_5^3, u_b^3, b_4^3, b_3^3, b_2^3, b_1^3, b_7^2, b_6^2, b_5^2, u_b^2, b_4^2, b_3^2, b_2^2, b_1^2, w_2)$. The third region is $R^{i,i+1}$.

Since D^1 and M_1 do not share vertices with D^2 , M_2 , and M_3 , since D^1 and D^2 are both incident to the outer face of Γ'' , and since R^1 and R^2 are simple, we have that R^1 and R^2 are disjoint. Further, the only faces of Γ'' regions R^1 and R^2 are both incident to (and which the order-preserving edges lie because of the planarity of Γ) are the outer face of Γ'' and $R^{i,i+1}$; however, neither M_i nor M_{i+1} is incident to the outer face of Γ'' , since M_i and M_{i+1} lie inside D . It follows that every order-preserving edge connecting a vertex in V'_i with a vertex in V'_{i+1} lies inside $R^{i,i+1}$. \square

We now present the following claim, which argues about the incidences between the corner edges and the square Q_i .

Claim 3 *For each $i \in \{1, 2, 3\}$, one of the following statements holds true (see respectively Figs. 5a and 5b).*

- (1) *Vertices x_i , y_i , w_i , and z_i are the first, second, last but one, and last vertex in σ_i , respectively. Further, (t_3^i, y_i) and (t_4^i, w_i) are assigned to the top side of Q_i , (t_7^i, y_i) and (b_7^i, w_i) are assigned to the right side of Q_i , (b_4^i, y_i) and (b_5^i, w_i) are assigned to the bottom side of Q_i , and (t_1^i, y_i) and (b_1^i, w_i) are assigned to the left side of Q_i .*
- (2) *Vertices z_i , w_i , y_i , and x_i are the first, second, last but one, and last vertex in σ_i , respectively. Further, (b_5^i, w_i) and (b_4^i, y_i) are assigned to the*




Figure 5: (a) Statement (1) of Claim 3. (b) Statement (2) of Claim 3.

top side of Q_i , (b_1^i, w_i) and (t_1^i, y_i) are assigned to the right side of Q_i , (t_4^i, w_i) and (t_3^i, y_i) are assigned to the bottom side of Q_i , and (b_7^i, w_i) and (t_7^i, y_i) are assigned to the left side of Q_i .

Proof: We prove that, if y_i precedes w_i in σ_i , then statement (1) holds true. A similar proof shows that, if w_i precedes y_i in σ_i , then statement (2) holds true.

First, each of the four corner edges (t_1^i, y_i) , (t_3^i, y_i) , (t_7^i, y_i) , and (b_4^i, y_i) incident to y_i is assigned to a distinct side of Q_i in Γ . Indeed assume, for a contradiction, that two of these corner edges, say (t_1^i, y_i) and (t_3^i, y_i) , are assigned to the same side of Q_i , as in Fig. 6a; this implies that the end-points of (t_1^i, y_i) and (t_3^i, y_i) on Q_i coincide. Let p^* be the end-point of these edges on Q_i ; thus, p^* is on the boundary of a row or column of M_i associated to y_i . Let R^* be the region delimited by the corner edges (t_1^i, y_i) and (t_3^i, y_i) , by the bounding edges (t_1^i, t_2^i) and (t_2^i, t_3^i) , and by the boundaries of the matrices representing t_1^i , t_2^i , and t_3^i . Suppose that Q_i is not contained in R^* . If edge (t_2^i, x_i) leaves the matrix representing t_2^i outside D , then it crosses D since M_i is inside D , a contradiction. Otherwise, (t_2^i, x_i) leaves the matrix representing t_2^i inside R^* . Since p^* is the only point on Q_i incident to R^* and since p^* is not on the boundary of a row or column of M_i associated to x_i , we have that (t_2^i, x_i) crosses the boundary of R^* , a contradiction. If Q_i is contained in R^* , a contradiction can be derived analogously by considering the routing of edge (t_6^i, x_i) rather than (t_2^i, x_i) .

A similar argument proves that each of the four corner edges (t_4^i, w_i) , (b_7^i, w_i) , (b_5^i, w_i) , and (b_1^i, w_i) incident to w_i is assigned to a distinct side of Q_i in Γ .

We now prove that (t_1^i, y_i) is assigned to the left side of Q_i . Assume the contrary, for a contradiction; refer to Fig. 6b, where (t_1^i, y_i) is assigned to the top side of Q_i . Consider the line segment Q^+ traversed when walking along Q_i in clockwise direction from the end-point of (t_1^i, y_i) to the end-point of (t_3^i, y_i) . If (t_1^i, y_i) is assigned to the top, right, or bottom side of Q_i , we have that Q^+ entirely contains the top side of the column of M_i associated to w_i , the right side of the row of M_i associated to w_i , or the left side of the row of M_i associated to w_i , respectively. Since one of the four corner edges incident to w_i has its end-point on this segment and since none of t_1^i , t_2^i , and t_3^i is adjacent to w_i , it




Figure 6: (a) Contradiction to the planarity of Γ if (t_1^i, y_i) and (t_3^i, y_i) are assigned to the same side of Q_i . The gray region is R^* . (b) Contradiction to the planarity of Γ if (t_1^i, y_i) is assigned to the top side of Q_i . (c) All the vertices in V_i' come after y_i and before w_i in σ_i . The gray region is R_b .

follows that: (i) one of the corner edges (t_1^i, y_i) and (t_3^i, y_i) , or one of the corner edges incident to w_i crosses a matrix representing a vertex in D or crosses one of the bounding edges; or (ii) two among the corner edges (t_1^i, y_i) and (t_3^i, y_i) , and among the corner edges incident to w_i cross each other. In both cases we get a contradiction to the planarity of Γ .

Since (t_1^i, y_i) is assigned to the left side of Q_i , since each of the four corner edges incident to y_1 is assigned to a distinct side of Q_i , and since t_1^i, t_3^i, t_7^i , and b_4^i appear in this clockwise order along D , we have that (t_3^i, y_i) , (t_7^i, y_i) , and (b_4^i, y_i) are assigned to the top, right, and bottom side of Q_i , respectively.

An analogous proof shows that (t_4^i, w_i) , (b_7^i, w_i) , (b_5^i, w_i) , and (b_1^i, w_i) are assigned to the top, right, bottom, and left side of Q_i , respectively.

Further, (t_2^i, x_i) is assigned to the left or top side of Q_i , since (t_1^i, y_i) and (t_3^i, y_i) are assigned to the left and top side of Q_i , respectively, and since t_1^i, t_2^i , and t_3^i appear in this clockwise order along D . In both cases, x_i precedes y_i in σ_i . An analogous argument proves that z_i follows w_i in σ_i . Hence, x_i, y_i, w_i , and z_i appear in this order in σ_i .

It remains to argue that x_i, y_i, w_i , and z_i are the first, second, last but one, and last vertex in σ_i , respectively; refer to Fig. 6c. Let R_b be the region delimited by the corner edges (b_4^i, y_i) and (b_5^i, w_i) , by the bounding edges (b_4^i, u_b^i) and (u_b^i, b_5^i) , and by the boundaries of the matrices representing b_4^i, u_b^i, b_5^i , and V_i'' . If any side-filling edge incident to u_b^i leaves the matrix representing u_b^i outside D , then this edge crosses D , a contradiction to the planarity of Γ . Otherwise, every side-filling edge incident to u_b^i leaves the matrix representing u_b^i inside R_b . Since (b_4^i, y_i) and (b_5^i, w_i) are both assigned to the bottom side of Q_i , it follows that all the side-filling edges incident to u_b^i have their other end-point on the bottom side of Q_i , between the end-point of (b_4^i, y_i) and the end-point of (b_5^i, w_i) . Thus, all the vertices in V_i' come after y_i and before w_i in σ_i ; this concludes the proof of statement (1). \square

We are now ready to state and prove the following two lemmata.

Lemma 1 *For each $i \in \{1, 2, 3\}$, all the equivalence edges belonging to the set $E'_i = \{(u'_e, r'_i), (u'_e, s'_i) : e = (r, s) \in E_i\}$ are assigned to the same side of Q_i , except possibly for the edges (u'_e, r'_i) and (u'_e, s'_i) such that r'_i and s'_i are consecutive in σ_i .*

Proof: We prove that every edge in $E'_i = \{(u'_e, r'_i), (u'_e, s'_i) : e = (r, s) \in E_i\}$ is assigned to the same side of Q_i edges (t^i_3, y_1) and (t^i_4, w_1) are assigned to, except possibly for the edges (u'_e, r'_i) and (u'_e, s'_i) such that r'_i and s'_i are consecutive in σ_i . The statement clearly implies Lemma 1. Note that (t^i_3, y_i) and (t^i_4, w_i) are assigned either both to the top or both to the bottom side of Q_i , by Claim 3.

We first prove that, for any edge $e = (r, s) \in E_i$, the edges (u'_e, r'_i) and (u'_e, s'_i) are assigned to the same side of Q_i . Indeed suppose, for a contradiction, that the edge (u'_e, r'_i) is assigned to, say, the top side of Q_i and the edge (u'_e, s'_i) is assigned to a different side of Q_i . By Claim 3, there exist two corner edges that are assigned to the top side of Q_i and that are incident to the second and to the last but one vertex in σ_i , respectively; these corner edges are (t^i_3, y_i) and (t^i_4, w_i) , or (b^i_5, w_i) and (b^i_4, y_i) . Assume that they are (t^i_3, y_i) and (t^i_4, w_i) , as the other case is analogous. Consider the region R_t delimited by (t^i_3, y_i) and (t^i_4, w_i) , by the bounding edge (t^i_3, t^i_4) , by the boundaries of the matrices representing t^i_3 and t^i_4 , and by the top side of Q_i . By Claim 3, the incidence point of edge (u'_e, r'_i) on the top side of Q_i is between the incidence points of (t^i_3, y_i) and (t^i_4, w_i) on the top side of Q_i , given that $r'_i \in V'_i$. Since edge (u'_e, r'_i) does not cross M_i , it leaves M_i inside R_t . Since (u'_e, s'_i) is assigned to a side of Q_i different from the top side, then edge (u'_e, r'_i) , or edge (u'_e, s'_i) , or the matrix representing u'_e crosses the boundary of R_t , a contradiction to the planarity of Γ . It follows that the edges (u'_e, r'_i) and (u'_e, s'_i) are assigned to the same side of Q_i . Now suppose that, for some $e = (r, s) \in E_i$, the edges (u'_e, r'_i) and (u'_e, s'_i) are both assigned to a side different from the one the edges (t^i_3, y_i) and (t^i_4, w_i) are assigned to, and suppose that r'_i and s'_i are not consecutive in σ_i . By Claim 3, the edges (u'_e, r'_i) and (u'_e, s'_i) are both assigned to the side (b^i_4, y_i) and (b^i_5, w_i) are assigned to, or to the side (t^i_1, y_i) and (b^i_1, w_i) are assigned to, or to the side (t^i_7, y_i) and (b^i_7, w_i) are assigned to.

Assume first that (u'_e, r'_i) and (u'_e, s'_i) are both assigned to the side (b^i_4, y_i) and (b^i_5, w_i) are assigned to. As in the proof of Claim 3, we can define R_b as the region delimited by the edges (b^i_4, y_i) , (b^i_5, w_i) , (b^i_4, u^i_b) , and (u^i_b, b^i_5) , and by the boundaries of the matrices representing b^i_4 , u^i_b , b^i_5 , and V^i_b ; then every side-filling edge incident to u^i_b leaves the matrix representing u^i_b inside R_b , as otherwise it would cross D . It follows that all the side-filling edges incident to u^i_b have their other end-point on the same side of Q_i edges (b^i_4, y_i) and (b^i_5, w_i) are assigned to. Let $p'_i \in V'_i$ be any vertex between r'_i and s'_i in σ_i . This vertex exists by hypothesis. Since p'_i is between r'_i and s'_i in σ_i , the incidence point of the side-filling edge (u^i_b, p'_i) on Q_i is between the incidence points of (u'_e, r'_i) and (u'_e, s'_i) on Q_i . It follows that (u^i_b, p'_i) crosses M_i , or the matrix representing u'_e , or (u'_e, r'_i) , or (u'_e, s'_i) .

The cases in which (u'_e, r'_i) and (u'_e, s'_i) are both assigned to the side (t^i_1, y_i) and (b^i_1, w_i) are assigned to (with $i = 1$), or to the side (t^i_7, y_i) and (b^i_7, w_i) are

assigned to (with $i = 3$) are analogous to the previous one, with u_l or u_r playing the role of u_b^i , respectively.

Assume next that (u'_e, r'_i) and (u'_e, s'_i) are both assigned to the side (t_7^i, y_i) and (b_7^i, w_i) are assigned to and $i \leq 2$. By Claim 3 and by the planarity of Γ , the matrix representing u'_e and the edges (u'_e, r'_i) and (u'_e, s'_i) lie inside region $R^{i,i+1}$. By Claim 2, every order-preserving edge between a vertex in V'_i and a vertex in V'_{i+1} also lies inside $R^{i,i+1}$, hence it is assigned to the same side of Q_i as (t_7^i, y_i) and (b_7^i, w_i) . Again by Claim 3, every order-preserving edge between a vertex in V'_i and a vertex in V'_{i+1} has its end-point between the end-points of (t_7^i, y_i) and (b_7^i, w_i) along the side of Q_i these edges are all assigned to. Let $p'_i \in V'_i$ be any vertex between r'_i and s'_i in σ_i . This vertex exists by hypothesis. Since p'_i is between r'_i and s'_i in σ_i , then the incidence point of the order-preserving edge (p'_i, p'_{i+1}) on Q_i is between the incidence points of (u'_e, r'_i) and (u'_e, s'_i) on Q_i . It follows that (p'_i, p'_{i+1}) crosses M_i , or the matrix representing u'_e , or (u'_e, r'_i) , or (u'_e, s'_i) .

The case in which (u'_e, r'_i) and (u'_e, s'_i) are both assigned to the side (t_1^i, y_i) and (b_1^i, w_i) are assigned to and $i \geq 2$ is analogous to the previous one.

This concludes the proof of Lemma 1. \square

Lemma 2 *For any $i, j \in \{1, 2, 3\}$, the orderings $\sigma_i - \{x_i, y_i, w_i, z_i\}$ of V'_i and $\sigma_j - \{x_j, y_j, w_j, z_j\}$ of V'_j either are the same ordering or are the reverse of each other (according to the bijection between V'_i and V'_j).*

Proof: By Claim 3, (t_7^i, y_i) and (b_7^i, w_i) are both assigned to the right or to the left side of Q_i , and (t_1^{i+1}, y_{i+1}) and (b_1^{i+1}, w_{i+1}) are both assigned to the right or to the left side of Q_{i+1} . We distinguish four cases, based on these two assignments.

If (t_7^i, y_i) and (b_7^i, w_i) are both assigned to the right side of Q_i , and (t_1^{i+1}, y_{i+1}) and (b_1^{i+1}, w_{i+1}) are both assigned to the left side of Q_{i+1} , then the orderings $\sigma_i - \{x_i, y_i, w_i, z_i\}$ of V'_i and $\sigma_{i+1} - \{x_{i+1}, y_{i+1}, w_{i+1}, z_{i+1}\}$ of V'_{i+1} are the same ordering (according to the bijection between V'_i and V'_{i+1}). In fact, traversing the right side of Q_i from the incidence point with (t_7^i, y_i) to the incidence point with (b_7^i, w_i) , the segments delimiting the rows associated to the vertices in V'_i are encountered in the order $\sigma_i - \{x_i, y_i, w_i, z_i\}$; this is because by Claim 3 the vertex y_i precedes the vertex w_i in σ_i . Analogously, traversing the left side of Q_{i+1} from the incidence point with (t_1^{i+1}, y_{i+1}) to the incidence point with (b_1^{i+1}, w_{i+1}) , the segments delimiting the rows associated to the vertices in V'_{i+1} are encountered in the order $\sigma_{i+1} - \{x_{i+1}, y_{i+1}, w_{i+1}, z_{i+1}\}$. Now if there were a pair of vertices p'_i and q'_i such that p'_i and q'_i are in this order in σ_i and such that q'_{i+1} and p'_{i+1} are in this order in σ_{i+1} , then the order-preserving edges (p'_i, p'_{i+1}) and (q'_i, q'_{i+1}) would cross each other, contradicting the planarity of Γ .

If (t_7^i, y_i) and (b_7^i, w_i) are both assigned to the left side of Q_i , and (t_1^{i+1}, y_{i+1}) and (b_1^{i+1}, w_{i+1}) are both assigned to the right side of Q_{i+1} , the proof is analogous.

If (t_7^i, y_i) and (b_7^i, w_i) are both assigned to the right side of Q_i , and (t_1^{i+1}, y_{i+1}) and (b_1^{i+1}, w_{i+1}) are both assigned to the right side of Q_{i+1} , then the orderings $\sigma_i - \{x_i, y_i, w_i, z_i\}$ of V_i' and $\sigma_{i+1} - \{x_{i+1}, y_{i+1}, w_{i+1}, z_{i+1}\}$ of V_{i+1}' , respectively, are the reverse of each other (according to the bijection between V_i' and V_{i+1}'). In fact, traversing the right side of Q_i from the incidence point with (t_7^i, y_i) to the incidence point with (b_7^i, w_i) , the segments delimiting the rows associated to the vertices in V_i' are encountered in the order $\sigma_i - \{x_i, y_i, w_i, z_i\}$. However, traversing the right side of Q_{i+1} from the incidence point with (t_1^{i+1}, y_{i+1}) to the incidence point with (b_1^{i+1}, w_{i+1}) , the segments delimiting the rows associated to the vertices in V_{i+1}' are encountered in the reverse order of $\sigma_{i+1} - \{x_{i+1}, y_{i+1}, w_{i+1}, z_{i+1}\}$; this is because by Claim 3 the vertex y_i follows the vertex w_i in σ_i . Now if there is a pair of vertices p'_i and q'_i such that p'_i and q'_i are in this order in σ_i and such that p'_{i+1} and q'_{i+1} are in this order in σ_{i+1} , then the order-preserving edges (p'_i, p'_{i+1}) and (q'_i, q'_{i+1}) cross each other, contradicting the planarity of Γ .

If (t_7^i, y_i) and (b_7^i, w_i) are assigned to the left side of Q_i , and (t_1^{i+1}, y_{i+1}) and (b_1^{i+1}, w_{i+1}) are assigned to the left side of Q_{i+1} the proof is analogous. \square

We are now ready to present a lemma that concludes the proof of sufficiency.

Lemma 3 *The graph (V, E) , together with the partition $E = E_1 \cup E_2 \cup E_3$, is a positive instance of PARTITIONED 3-PAGE BOOK EMBEDDING.*

Proof: The proof employs Lemmata 1 and 2.

Consider the ordering \mathcal{O} of V corresponding to $\sigma_1 - \{x_1, y_1, w_1, z_1\}$ according to the bijection between V and V_1' . By Lemma 2, for $i = 2, 3$, either $\sigma_i - \{x_i, y_i, w_i, z_i\}$ is the same or is the reverse of $\sigma_1 - \{x_1, y_1, w_1, z_1\}$ (according to the bijection between V_i' and V_1'). Then, for $i = 1, 2, 3$, no two edges $e = (p, q)$ and $f = (r, s)$ in E_i have alternating end-vertices in \mathcal{O} . Indeed, e and f cannot have alternating end-vertices if p'_i and q'_i are consecutive in σ_i , or if r'_i and s'_i are consecutive in σ_i , as in this case p and q , or r and s would be consecutive in \mathcal{O} , respectively. If p'_i and q'_i are not consecutive in σ_i and r'_i and s'_i are not consecutive in σ_i , then by Lemma 1 the edges (u'_e, p'_i) , (u'_e, q'_i) , (u'_f, r'_i) , and (u'_f, s'_i) are all assigned to the same side of Q_i , hence if e and f had alternating end-vertices then: (i) (u'_e, p'_i) or (u'_e, q'_i) would cross (u'_f, r'_i) or (u'_f, s'_i) , or (ii) (u'_e, p'_i) or (u'_e, q'_i) would cross the matrix representing u'_f or M_i , or (iii) (u'_f, r'_i) or (u'_f, s'_i) would cross the matrix representing u'_e or M_i , or (iv) two among M_i and the matrices representing u'_e and u'_f would cross each other. In each case we would get a contradiction to the planarity of Γ . \square

Lemma 3 completes the proof that $(V, E = E_1 \cup E_2 \cup E_3)$ is a positive instance of PARTITIONED 3-PAGE BOOK EMBEDDING if and only if (V', E', \mathcal{C}') is a positive instance of NT PLANARITY. Thus, the NP-completeness of PARTITIONED 3-PAGE BOOK EMBEDDING [8] concludes the proof of Theorem 1.

We now turn our attention to flat clustered graphs with given side assignments. Let $G = (V, E, \mathcal{C})$ be a flat clustered graph with a given side assignment




Figure 7: Illustration for the proof of Theorem 2.

s_i , for each $V_i \in \mathcal{C}$. We say that G is *NT planar with fixed sides* if G admits an NT planar representation Γ such that, for every edge $e = (u, v) \in E$ with $u \in V_i$ and $v \in V_j$, the incidence points of e with the matrices M_i and M_j representing V_i and V_j in Γ lie on the straight-line segments corresponding to the $s_i(e)$ side of M_i and to the $s_j(e)$ side of M_j , respectively.

Theorem 2 NODETRIX PLANARITY WITH FIXED SIDES is NP -complete even for instances with two clusters.

Proof: Lemma 4 will prove that NT PLANARITY WITH FIXED SIDES is in NP .

We give a reduction from the NP -complete problem BETWEENNESS [20], where an instance is a collection of ordered triplets of items and the target is to find a total order of the items in which, for each of the given triplets, the middle item in the triplet appears somewhere between the other two items.

Consider an instance \mathcal{T} of BETWEENNESS, i.e., a set of h items $\{a_1, a_2, \dots, a_h\}$ and a collection of t ordered triplets of $\tau_j = \langle a_{b_j}, a_{c_j}, a_{d_j} \rangle$, with $j = 1, \dots, t$. We construct the corresponding instance of NT PLANARITY WITH FIXED SIDES by defining a flat clustered graph (V, E, \mathcal{C}) , where $\mathcal{C} = \{V_1, V_2\}$, and a side assignment s_i , with $i \in \{1, 2\}$, as follows:

- cluster V_1 contains $2 + h \times t$ vertices: that is, v_α, v_β , and one vertex $v_{[i,j]}$ for each $i = 1, \dots, h$ and each $j = 1, \dots, t$;
- cluster V_2 contains $2 + h \times (t - 1) + 2t$ vertices: that is, u_α, u_β , plus one vertex $u_{[i,j]}$ for each $i = 1, \dots, h$ and each $j = 1, \dots, t - 1$, plus two vertices u'_j and u''_j for each $j = 1, \dots, t$.

The set E contains an arbitrary set of intra-cluster edges and the following inter-cluster edges.

- inter-cluster edge $e_\alpha = (v_\alpha, u_\alpha)$, with $s_1(e_\alpha) = R$ and $s_2(e_\alpha) = L$;
- inter-cluster edge $e_\beta = (v_\beta, u_\beta)$, with $s_1(e_\beta) = L$ and $s_2(e_\beta) = R$;
- for each $i = 1, \dots, h$ an inter-cluster edge $e_{[i,1]}^B = (u_\beta, v_{[i,1]})$, with $s_1(e_{[i,1]}^B) = B$ and $s_2(e_{[i,1]}^B) = R$;
- for each $i = 1, \dots, h$ and $j = 1, \dots, t - 1$ an inter-cluster edge $e_{[i,j]}^B = (u_{[i,j]}, v_{[i,j+1]})$, with $s_1(e_{[i,j]}^B) = s_2(e_{[i,j]}^B) = B$;
- for each $i = 1, \dots, h$ and $j = 1, \dots, t - 1$ an inter-cluster edge $e_{[i,j]}^T = (v_{[i,j]}, u_{[i,j]})$, with $s_1(e_{[i,j]}^T) = s_2(e_{[i,j]}^T) = T$;
- for each triplet $t_j = \langle a_{b_j}, a_{c_j}, a_{d_j} \rangle$, with $j = 1, \dots, t$, a path of four inter-cluster edges joining the five vertices $v_{[a_{b_j},j]}$, u'_j , $v_{[a_{c_j},j]}$, u''_j , and $v_{[a_{d_j},j]}$ in this order. Each edge e of such a path has $s_1(e) = R$ and $s_2(e) = L$.

(\implies) Suppose that the items of \mathcal{T} admit a total order $a_{\pi_1}, a_{\pi_2}, \dots, a_{\pi_h}$ in which for each of the given triplets, the middle item in the triplet appears somewhere between the other two items. We show how to construct a NodeTrix planar representation of (V, E, \mathcal{C}) .

Use for the matrix M_1 representing V_1 a row-column order σ_1 such that $\sigma_1(v_{[\pi_1,1]}) < \sigma_1(v_{[\pi_2,1]}) < \dots < \sigma_1(v_{[\pi_h,1]}) < \sigma_1(v_{[\pi_1,2]}) < \dots < \sigma_1(v_{[\pi_h,2]}) < \dots < \sigma_1(v_{[\pi_1,t]}) < \dots < \sigma_1(v_{[\pi_h,t]}) < \sigma_1(e_\alpha) < \sigma_1(e_\beta)$. Use for the matrix M_2 representing V_2 a row-column order σ_2 such that $\sigma_2(u'_t) < \sigma_2(u''_t) < \sigma_2(u_{[\pi_h,t-1]}) < \sigma_2(u_{[\pi_{h-1},t-1]}) < \dots < \sigma_2(u_{[\pi_1,t-1]}) < \sigma_2(u'_{t-1}) < \sigma_2(u''_{t-1}) < \sigma_2(u_{[\pi_h,t-2]}) < \sigma_2(u_{[\pi_{h-1},t-2]}) < \dots < \sigma_2(u_{[\pi_1,t-2]}) < \sigma_2(u'_{t-2}) < \sigma_2(u''_{t-2}) < \dots < \sigma_2(u_{[\pi_h,1]}) < \sigma_2(u_{[\pi_{h-1},1]}) < \dots < \sigma_2(u_{[\pi_1,1]}) < \sigma_2(u'_1) < \sigma_2(u''_1) < \sigma_2(e_\alpha) < \sigma_2(e_\beta)$. It can be easily seen that the inter-cluster edges can be drawn attached to the sides imposed by s_1 and s_2 without crossings, as in Fig. 7.

(\impliedby) Suppose that (V, E, \mathcal{C}) admits a NodeTrix planar representation where, for $i \in \{1, 2\}$, each inter-cluster edge attaches according to the edge assignment s_i to the matrix M_i representing the cluster V_i . We show that \mathcal{T} admits a total order in which for each triplet, the middle item in the triplet appears somewhere between the other two items.

First observe that, whatever the row-column orders σ_1 and σ_2 chosen for matrices M_1 and M_2 are, respectively, the matrices M_1 and M_2 form, together with the edges e_α and e_β , a cycle that separates the top sides of the two matrices from their bottom sides. It follows that all inter-cluster edges that attach to the top side (to the bottom side) of M_1 or M_2 are drawn inside the same region delimited by the boundaries of M_1 and M_2 , by e_α , and by e_β ; we denote by R_T (by R_B) the one of these regions comprising the top side (resp. the bottom side) of M_1 and M_2 . Refer again to Fig. 7 for an illustration.

Consider the inter-cluster edges $e_{[i,1]}^B$, for $i = 1, \dots, h$. Since they all attach to u_β and since $s_1(e_{[i,1]}^B) = B$, they are all drawn inside R_B . Denote by $\pi = \pi_1, \dots, \pi_h$ the permutation of the indices $1, \dots, h$ such that $\sigma_1(v_{[\pi_1,1]}) <$

$\sigma_1(v_{[\pi_2,1]}) < \dots < \sigma_2(v_{[\pi_h,1]})$ (recall that these are the end-vertices of the edges $e_{[i,1]}^B$).

We claim that $\sigma_1(v_{[\pi_1,1]}) < \dots < \sigma_1(v_{[\pi_h,1]}) < \sigma_1(v_{[\pi_1,2]}) < \dots < \sigma_1(v_{[\pi_h,2]}) < \dots < \sigma_1(v_{[\pi_1,t]}) < \dots < \sigma_1(v_{[\pi_h,t]})$ holds true. First, we prove that, for each $j \in \{1, \dots, t\}$, the vertices $v_{[\pi_1,j]}, v_{[\pi_2,j]}, \dots, v_{[\pi_h,j]}$ appear in this order in the row-column order of M_1 ; indeed, by the definition of π , this is the case for $j = 1$. Observe that the inter-cluster edges $e_{[i,1]}^T$, for $i = 1, \dots, h$, are drawn inside R_T and force σ_2 to be such that $\sigma_2(u_{[\pi_h,1]}) < \sigma_2(u_{[\pi_{h-1},1]}) < \dots < \sigma_2(u_{[\pi_1,1]})$ (recall that the end-vertices of the edge $e_{[i,1]}^T$ are $u_{[i,1]}$ and $v_{[i,1]}$). Analogously, the inter-cluster edges $e_{[i,1]}^B$, for $i = 1, \dots, h$, are drawn inside R_B and force σ_1 to be such that $\sigma_1(v_{[\pi_1,2]}) < \sigma_1(v_{[\pi_2,2]}) < \dots < \sigma_1(v_{[\pi_h,2]})$ (recall that the end-vertices of the edge $e_{[i,1]}^B$ are $u_{[i,1]}$ and $v_{[i,2]}$). For $j = 2, \dots, t-1$, the same argument can be repeated alternately for all the inter-cluster edges $e_{[i,j]}^T$ and then for all the inter-cluster edges $e_{[i,j]}^B$; then for any $j \in \{1, \dots, t\}$, it holds true that $\sigma(v_{[\pi_1,j]}) < \sigma(v_{[\pi_2,j]}) < \dots < \sigma_1(v_{[\pi_h,j]})$. Also, it is easy to see that, in order for the drawing to be crossing-free, $\sigma_1(v_{[i',j']}) < \sigma_1(v_{[i'',j'']})$ whenever $j' < j''$. This concludes the proof of the claim.

Now, for each $j = 1, \dots, t$, consider the inter-cluster edges of the path $v_{[a_{b_j},j]}, u'_j, v_{[a_{c_j},j]}, u''_j$, and $v_{[a_{d_j},j]}$. In any NT planar representation of (V, E, \mathcal{C}) such a path forces $v_{[a_{c_j},j]}$ to be in the middle of $v_{[a_{b_j},j]}$ and $v_{[a_{d_j},j]}$, that is, it forces a_{c_j} to be in the middle of a_{b_j} and a_{d_j} in π . It follows that π is an ordering of the items $a_{\pi_1}, a_{\pi_2}, \dots, a_{\pi_h}$ in which, for each triplet, the middle item in the triplet appears between the other two items. \square

Let $G = (V, E, \mathcal{C})$ be a flat clustered graph with a given row-column order σ_i , for each $V_i \in \mathcal{C}$. We say that G is *NT planar with fixed order* if it admits an NT planar representation Γ where, for each cluster $V_i \in \mathcal{C}$, each vertex $v \in V_i$ is associated with the $\sigma_i(v)$ -th row and column of the matrix M_i representing V_i in Γ .

Theorem 3 NODETRIX PLANARITY WITH FIXED ORDER is NP -complete even if at most one cluster contains more than one vertex.

Proof: The membership in NP of NT PLANARITY WITH FIXED ORDER will be proved in Lemma 4.

For the NP -hardness, we give a reduction from the NP -complete problem that asks to determine whether a proper 4-coloring exists for a *circle graph* [23], which is an intersection graph of chords of a circle. Let $G = (N, A)$ be a circle graph. First, by means of the algorithm in [22], we construct in polynomial time an intersection representation $\langle \mathcal{P}, \mathcal{O} \rangle$ of G , where \mathcal{P} is a linear sequence of distinct points on a circle and \mathcal{O} is a set of chords between pairs of points in \mathcal{P} such that: (i) each chord $c \in \mathcal{O}$ corresponds to a vertex $n \in N$ and (ii) two chords $c', c'' \in \mathcal{O}$ intersect if and only if $(n', n'') \in A$, where n' and n'' are the vertices in N corresponding to c' and c'' , respectively. Then G admits a proper 4-coloring if and only if the chords in \mathcal{O} can be 4-colored so that no

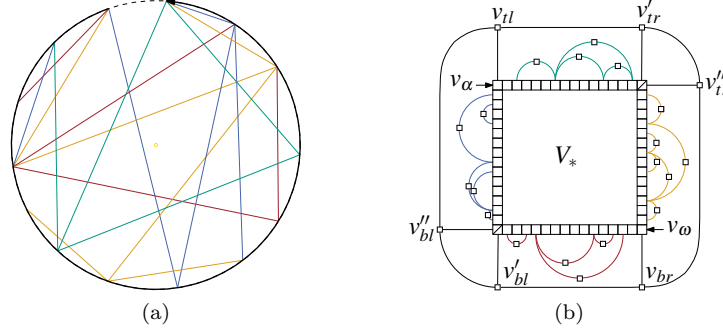


Figure 8: (a) An intersection representation $\langle \mathcal{P}, \mathcal{O} \rangle$ of a circle graph $G = (N, A)$. (b) Instance (V, E, \mathcal{C}) of NODETRIX PLANARITY corresponding to $\langle \mathcal{P}, \mathcal{O} \rangle$.

two chords of the same color intersect. Starting from $\langle \mathcal{P}, \mathcal{O} \rangle$ we construct an instance (V, E, \mathcal{C}) of NODETRIX PLANARITY WITH FIXED ORDER as follows (refer to Fig. 8). Set \mathcal{C} contains:

- a cluster V_* containing one vertex v_i for each point $p_i \in \mathcal{P}$, plus two additional vertices v_α and v_ω ;
- six clusters $\{v_{tl}\}$, $\{v'_{tr}\}$, $\{v''_{tr}\}$, $\{v_{br}\}$, $\{v'_{bl}\}$, and $\{v''_{bl}\}$, respectively; and
- a cluster $\{v_c\}$, for each chord $c \in \mathcal{O}$.

The set E contains an arbitrary set of intra-cluster edges and the following inter-cluster edges.

- *bounding edges* (v_{tl}, v'_{tr}) , (v'_{tr}, v''_{tr}) , (v''_{tr}, v_{br}) , (v_{br}, v'_{bl}) , (v'_{bl}, v''_{bl}) , and (v''_{bl}, v_{tl}) ;
- *corner edges* (v_{tl}, v_α) , (v'_{tr}, v_ω) , (v''_{tr}, v_α) , (v_{br}, v_ω) , (v'_{bl}, v_α) , and (v''_{bl}, v_ω) ; and
- *chord edges*: for each chord $c = (p_i, p_j) \in \mathcal{O}$, edges (v_i, v_c) and (v_c, v_j) .

Finally, we fix the row-column order σ_* of the only non-unitary cluster V_* to be $v_\alpha \circ \mathcal{P} \circ v_\omega$ (where with a slight abuse of notation we denote by \mathcal{P} not only the order of the points on the circle in the given intersection representation of G , but also the corresponding order of the vertices in $V_* - \{v_\alpha, v_\omega\}$). We now prove the equivalence between the problem of properly 4-coloring $\langle \mathcal{P}, \mathcal{O} \rangle$ and the constructed instance of NODETRIX PLANARITY WITH FIXED ORDER.

(\implies) Suppose that the chords of $\langle \mathcal{P}, \mathcal{O} \rangle$ can be assigned colors 1, 2, 3, 4 so that no two chords with the same color intersect. We show how to construct a NodeTrix planar representation with fixed order of (V, E, \mathcal{C}) . Represent clusters V_* , $\{v_{tl}\}$, $\{v'_{tr}\}$, $\{v''_{tr}\}$, $\{v_{br}\}$, $\{v'_{bl}\}$, and $\{v''_{bl}\}$ by matrices M_* , M_{tl} , M'_{tr} , M''_{tr} , M_{br} , M'_{bl} , and M''_{bl} , respectively, where the row-column order of M_* is σ_* . Draw

the bounding edges so that M_* is inside the cycle D they compose together with M_{tl} , M'_{tr} , M''_{tr} , M_{br} , M'_{bl} , and M''_{bl} . Draw the corner edges also inside D . The corner edges, together with the boundary of M_* , subdivide the region of the plane inside D into five regions, namely one region internal to the boundary of M_* and four regions incident to the top, right, bottom and left side of M_* . We refer to these regions as to the *top*, *right*, *bottom*, and *left region*, respectively. Depending on whether a chord $c = (p_i, p_j)$ has color 1, 2, 3, or 4, we draw the chord edges (v_i, v_c) and (v_c, v_j) , as well as the matrix M_c representing cluster $\{v_c\}$, inside the top, right, bottom, and left region, respectively. We claim that the obtained NT representation of (V, E, \mathcal{C}) is planar. Suppose, for a contradiction, there is a crossing between two paths $(v_i, v_{c'}, v_j)$ and $(v_k, v_{c''}, v_h)$ corresponding to two chords $c' = (p_i, p_j)$ and $c'' = (p_k, p_h)$. Then these two paths are in the interior of the same (top, right, bottom, or left) region, hence they attach to the same side of M_* ; it follows that the two chords c' and c'' have the same color. By the definition of σ_* , since the end-vertices of the two paths alternate along the side of M_* , the end-points of c' and c'' alternate in \mathcal{P} . Hence, c' and c'' cross, thus contradicting the fact that c' and c'' have the same color.

(\Leftarrow) Suppose that (V, E, \mathcal{C}) admits a NodeTrix-planar representation Γ with a row-column order σ_* for the unique non-unitary cluster V_* . Denote by M_* the matrix representing V_* in Γ . We show that the chords of $\langle \mathcal{P}, \mathcal{O} \rangle$ are 4-colorable so that no two chords of the same color intersect.

Similarly to the proof of Theorem 1, it can be proved that the corner edges subdivide the side of D that contains M_* into five regions, defined as in the previous direction, so that all the vertices on the top, right, bottom, and left side of M_* are incident to the top, right, bottom, and left region, respectively; while in the proof of Theorem 1 this was ensured by Claim 3, it is here a trivial consequence of the fact that v_α and v_ω are the first and the last vertex in σ_* . By the planarity of Γ , both the incidence points of a path (v_i, v_c, v_j) with the boundary of M_* are on the same side of M_* . Then color all the chords $c = (p_i, p_j)$ in \mathcal{O} such that path (v_i, v_c, v_j) is in the top, right, bottom, or left region with color 1, 2, 3, or 4, respectively.

We claim that the obtained 4-coloring of the chords of $\langle \mathcal{P}, \mathcal{O} \rangle$ is proper. Suppose, for a contradiction, there is a crossing in $\langle \mathcal{P}, \mathcal{O} \rangle$ between two chords $c' = (p_i, p_j)$ and $c'' = (p_k, p_h)$ both with color 1 – the discussion for the other colors is analogous. Then c' and c'' have alternating end-points in $\langle \mathcal{P}, \mathcal{O} \rangle$. Since the order of the points in \mathcal{P} coincides with the order of the corresponding vertices in σ_* , it follows that paths $(v_i, v_{c'}, v_j)$ and $(v_k, v_{c''}, v_h)$ have alternating end-points on the top side of M_* , hence they cross, a contradiction to the planarity of Γ . This concludes the proof. \square

Let $G = (V, E, \mathcal{C})$ be a flat clustered graph with a given row-column order σ_i and side assignment s_i , for each $V_i \in \mathcal{C}$. Then G is *NT planar with fixed order and fixed side* if it is simultaneously planar with fixed order and with fixed sides.

Theorem 4 NODETRIX PLANARITY WITH FIXED ORDER AND FIXED SIDES can be solved in linear time.

Proof: Consider the graph G' obtained from an instance $G = (V, E, \mathcal{C})$ of NODETRIX PLANARITY WITH FIXED ORDER AND FIXED SIDES by collapsing each cluster $V_i \in \mathcal{C}$ into a vertex v_i . Intuitively, instance G is NT planar with fixed order and fixed sides if and only if G' is planar with the additional constraint that the clockwise order of the edges incident to each vertex v_i is “compatible” with the row-column order σ_i and the side assignment s_i for the cluster V_i .

More formally, denote by \mathcal{E}_i the set of the inter-cluster edges incident to V_i and denote by $v_i(e)$ the vertex of V_i incident to an edge $e \in \mathcal{E}_i$. The edges in \mathcal{E}_i can be decomposed into a circular sequence of sets $\mathcal{S} = \mathcal{E}_{T,1}, \mathcal{E}_{T,2}, \dots, \mathcal{E}_{T,|V_i|}, \mathcal{E}_{R,1}, \mathcal{E}_{R,2}, \dots, \mathcal{E}_{R,|V_i|}, \mathcal{E}_{B,|V_i|}, \mathcal{E}_{B,|V_i|-1}, \dots, \mathcal{E}_{B,1}, \mathcal{E}_{L,|V_i|}, \mathcal{E}_{L,|V_i|-1}, \dots, \mathcal{E}_{L,1}$, where each $\mathcal{E}_{X,j}$, with $X \in \{T, B, L, R\}$ and $j \in \{1, \dots, |V_i|\}$, contains the edges $e \in \mathcal{E}_i$ such that $s_i(e) = X$ and $\sigma_i(v_i(e)) = j$. Let Γ' be a planar embedding of G' and let λ_i denote the clockwise order of the edges incident to vertex v_i of G' in Γ' . The embedding Γ' of G' is *compatible* with functions σ_i and s_i if: (i) all the edges belonging to the same set $\mathcal{E}_{X,j}$ appear consecutively in λ_i , and (ii) for any three edges $e' \in \mathcal{E}_{X',j'}$, $e'' \in \mathcal{E}_{X'',j''}$, and $e''' \in \mathcal{E}_{X''',j'''}$, where $\mathcal{E}_{X',j'}$, $\mathcal{E}_{X'',j''}$, and $\mathcal{E}_{X''',j'''}$ are all distinct, appear in this clockwise order in λ_i if and only if $\mathcal{E}_{X',j'}$, $\mathcal{E}_{X'',j''}$, and $\mathcal{E}_{X''',j'''}$ appear in this circular order in \mathcal{S} .

It can be easily seen that an instance of NODETRIX PLANARITY WITH FIXED ORDER AND FIXED SIDES has a solution if and only if G' admits an embedding Γ' that is compatible with σ_i and s_i , for all vertices v_i of G' . We obtain an instance of constrained planarity for G' that can be tested in linear time with known techniques [17]. \square

We conclude the section with the following lemma.

Lemma 4 NODETRIX PLANARITY, NODETRIX PLANARITY WITH FIXED SIDES, and NODETRIX PLANARITY WITH FIXED ORDER are in NP.

Proof: We prove the statement for NT PLANARITY; the other proofs are analogous. Consider an instance (V, E, \mathcal{C}) of NT PLANARITY. For each $V_i \in \mathcal{C}$, guess a row-column order σ_i and a side assignment s_i ; then use the algorithm described in the proof of Theorem 4 to test in linear time whether (V, E, \mathcal{C}) is NT planar with fixed order and fixed sides. Since the number of distinct row-column orders and side assignments is a function of $|V| + |E|$, we get the NP membership. \square

4 Monotone NodeTrix Representations

Let $G = (V, E, \mathcal{C})$ be a flat clustered graph and γ be a square assignment for G that maps each cluster in \mathcal{C} to an axis-aligned square in the plane. A curve is *x-monotone* (resp. *y-monotone*) if no two of its points have the same

projection on the x -axis (resp. on the y -axis) and is xy -monotone if it is either a horizontal or a vertical segment or it is both x - and y -monotone. A *monotone* NT representation Γ of $\langle G, \gamma \rangle$ is an NT representation such that:

1. all the inter-cluster edges are represented by xy -monotone curves;
2. for each cluster $V_i \in \mathcal{C}$, the boundary of the matrix M_i representing V_i is $Q_i = \gamma(V_i)$;
3. for each pair of adjacent clusters V_i and V_j , with $i \neq j$, the convex hull of Q_i and Q_j does not intersect any other square Q_k , with $k \neq i, j$ – we call this convex hull the *pipe* of Q_i and Q_j ; and
4. all the inter-cluster edges between vertices in V_i and vertices in V_j lie inside the pipe of Q_i and Q_j .

In a monotone NT representation Γ of G let $\chi_i(\Gamma)$ denote the number of edge crossings between pairs of inter-cluster edges incident to V_i . Let $\chi(\Gamma) = \sum_i \chi_i(\Gamma)$, where the sum is over all the clusters $V_i \in \mathcal{C}$; we say that Γ is *locally planar* if $\chi(\Gamma) = 0$ and no inter-cluster edge intersects any matrix except at its incidence points. The notions of *fixed order* and *fixed sides* easily extend to monotone NT representations.

In this section we study the complexity of testing if a flat clustered graph with fixed square assignment admits a monotone locally-planar NT representation, a problem which we call MONOTONE NT LOCAL PLANARITY (MNTLP).

There are several reasons to study the MNTLP problem. First, graph representations with monotone edges are natural and have been well investigated in the graph drawing literature (see, e.g., [4, 16, 21]). Second, straight-line NT representations, in which inter-cluster edges are straight-line segments, arguably provide a high readability of the adjacencies between vertices in different clusters. Note that straight-line NT representations are a particular type of monotone NT representations. Furthermore, we will prove in Lemma 5 that monotone NT representations are in fact equivalent to straight-line NT representations in terms of local planarity; that is, a flat clustered graph admits a monotone locally-planar NT representation if and only if it admits a straight-line locally-planar NT representation. Third, as shown in the previous section (see Theorems 1–3), the general NT PLANARITY problem is intractable in most cases, hence the need for a more constrained and manageable model; we will in fact show that, while the MNTLP problem is also intractable (see Theorems 5 and 6), it becomes polynomial-time solvable if the matrices have fixed orders and the number of clusters is bounded by a constant.

We start by proving the NP-hardness of the MNTLP problem and of its variant with fixed side assignment.

Theorem 5 *MNTLP is NP-complete.*

Proof: The proof that the problem is in NP is similar to the proof of Theorem 4: one can guess a row-column order and a side assignment for each cluster; then




Figure 9: Illustration for the proof of Theorem 5, with $m = 3$.

the monotone NT local planarity of the given clustered graph with the given square assignment and the guessed row-column order and side assignment can be tested in polynomial time by Theorem 7, to be presented later.

For the NP-hardness we give a reduction from the NP-complete problem BETWEENNESS [20], defined in the proof of Theorem 2. Consider an instance \mathcal{T} of BETWEENNESS consisting of a set of h items $\{a_1, \dots, a_h\}$ and m ordered triplets $\tau_i = \langle a_{b_i}, a_{c_i}, a_{d_i} \rangle$, with $i = 1, \dots, m$. We construct an instance $\langle G = (V, E, \mathcal{C}), \gamma \rangle$ of MNTLP as follows (refer to Fig. 9). Let \mathcal{C} consist of:

- for $i = 1, \dots, m$, a cluster V_i containing $h + 1$ vertices v_1^i, \dots, v_h^i, x^i ;
- two clusters $\{x^0\}$ and $\{u\}$; and
- for $i = 1, \dots, m$, a cluster $\{\alpha^i, \beta^i\}$.

The set E contains an arbitrary set of intra-cluster edges and the following inter-cluster edges.

- *order-preserving edges* connecting v_j^i with v_j^{i+1} , for $1 \leq i \leq m - 1$ and $1 \leq j \leq h$;
- *side-filling edges* connecting u with v_j^1 , for $1 \leq j \leq h$;
- *protecting edges* connecting x^i with x^{i+1} , for $0 \leq i \leq m - 1$;
- *corner edges* connecting x^i with α^i and β^i , for $1 \leq i \leq m$; and
- *betweenness edges* connecting, for each triplet $\tau_i = \langle a_{b_i}, a_{c_i}, a_{d_i} \rangle$ with $i \in \{1, \dots, m\}$, the vertex $v_{b_i}^i$ with α^i , the vertex $v_{d_i}^i$ with β^i , and the vertex $v_{c_i}^i$ with both α^i and β^i .

Finally, square assignment γ is defined as follows:

- for $i = 1, \dots, m$, the cluster V_i is assigned to a square Q_i so that, for every $1 \leq i < j \leq m$, we have $\min_y\{Q_i\} = \min_y\{Q_j\}$, $\max_y\{Q_i\} = \max_y\{Q_j\}$, and $\max_x\{Q_i\} < \min_x\{Q_j\}$;
- the cluster $\{x^0\}$ (the cluster $\{u\}$) is assigned to a square Q_0 (resp. Q_u), so that $\min_x\{Q_0\} = \min_x\{Q_u\} < \max_x\{Q_0\} = \max_x\{Q_u\} < \min_x\{Q_1\}$, and $\min_y\{Q_1\} < \min_y\{Q_u\} < \max_y\{Q_u\} < \min_y\{Q_0\} < \max_y\{Q_0\} < \max_y\{Q_1\}$; and

- for $i = 1, \dots, m$, the cluster $\{\alpha^i, \beta^i\}$ is assigned to a square Q_i^α so that $\max_x\{Q_{i-1}\} < \min_x\{Q_i^\alpha\} < \max_x\{Q_i^\alpha\} < \min_x\{Q_i\}$, and $\max_y\{Q_i\} < \min_y\{Q_i^\alpha\}$.

Notice that the above definition of γ is such that, for each pair of adjacent clusters V_i and V_j , with $i \neq j$, the convex hull of Q_i and Q_j does not intersect any other square Q_k , with $k \neq i, j$.

We now prove the equivalence between the given instance of BETWEENNESS and the constructed instance of MNTLP.

(\implies) If \mathcal{T} admits an order $\sigma = (a_{\pi_1}, \dots, a_{\pi_h})$ in which a_{c_i} appears between a_{b_i} and a_{d_i} in σ , for each triplet $\tau_i = (a_{b_i}, a_{c_i}, a_{d_i})$, we construct a monotone locally-planar NT representation Γ with fixed square assignment of G as follows. For $i = 1, \dots, m$, represent V_i as a matrix M_i with boundary Q_i and with row-column order $x^i, v_{\pi_1}^i, \dots, v_{\pi_h}^i$; represent $\{\alpha^i, \beta^i\}$ as a matrix M_i^α with boundary Q_i^α and with row-column order α^i, β^i or β^i, α^i depending on whether $v_{b_i}^i$ follows or precedes $v_{d_i}^i$ in σ , respectively; the representation of the unitary clusters $\{x^0\}$ and $\{u\}$ in Γ is trivially defined. Assign every inter-cluster edge incident to x^0 (or u) to the right side of Q_0 (or Q_u) and to the left side of Q_1 , every order-preserving or protecting edge between a vertex in V_i and a vertex in V_{i+1} to the right side of Q_i and to the left side of Q_{i+1} , every corner edge incident to x^i to the top side of Q_i and to the bottom side of Q_i^α , and every betweenness edge incident to a vertex in V_i to the top side of Q_i and to the right side of Q_i^α . Finally, draw all the inter-cluster edges as straight-line segments in Γ . Observe that such segments are xy -monotone curves inside the corresponding pipe.

Representation Γ has no crossing between any inter-cluster edge and any matrix, as a consequence of the square and side assignments and independently of the row-column order of the matrices. Further, the order-preserving, side-filling, and protecting edges do not cross the corner and betweenness edges since they are separated by the horizontal line $y = \max_y\{Q_1\}$, and do not cross each other since the row-column orders of any two matrices M_i and M_{i+1} both correspond to σ with x^i and x^{i+1} as the first element, respectively. The corner edges incident to x^i do not cross the betweenness edges incident to vertices in V_i , because of the side assignment of these edges to Q_i^α and since x^i precedes $v_{\pi_1}^i, \dots, v_{\pi_h}^i$ in the row-column order of M_i . Finally, the betweenness edges do not cross each other because of the row-column order defined for M_i^α and since the top side of the column representing $v_{c_i}^i$ in M_i is between the top sides of the columns representing $v_{b_i}^i$ and $v_{d_i}^i$ in M_i .

(\impliedby) Suppose that G admits a monotone locally-planar NT representation Γ with fixed square assignment γ . Let M_1, \dots, M_m be the matrices representing the clusters V_1, \dots, V_m in Γ , respectively. First, the monotonicity of Γ and the placement of squares Q_0, Q_1 , and Q_u imply that the side-filling edges, as well as the edge (x^0, x^1) , lie to the right of Q_0 and Q_u , to the left of Q_1 , above the line $y = \min_y\{Q_1\}$, and below the line $y = \max_y\{Q_1\}$. Since $\max_y\{Q_u\} < \min_y\{Q_0\}$, we have that all the side-filling edges lie below edge (x^0, x^1) , hence the planarity of Γ implies that x^1 is the first vertex in the row-column order of M_1 ; let $x^1, v_{\pi_1}^1, \dots, v_{\pi_h}^1$ be such an order, for some permutation π_1, \dots, π_h of

$\{1, \dots, h\}$. We claim that the total ordering $\sigma = (a_{\pi_1}, \dots, a_{\pi_h})$ is a solution to instance \mathcal{T} of BETWEENNESS.

We first prove that the order $x^1, v_{\pi_1}^1, \dots, v_{\pi_h}^1$ is “preserved” in M_2, \dots, M_m . The monotonicity of Γ and the placement of squares Q_i imply that, for $1 \leq i \leq m - 1$, the order-preserving and protecting edges between vertices in V_i and vertices in V_{i+1} are to the right of Q_i , to the left of Q_{i+1} , above the line $y = \min_y\{Q_i\}$, and below the line $y = \max_y\{Q_i\}$. Then the planarity of Γ implies that, if the row-column order of M_i is $x^i, v_{\pi_1}^i, \dots, v_{\pi_h}^i$, the row-column order of M_{i+1} is $x^{i+1}, v_{\pi_1}^{i+1}, \dots, v_{\pi_h}^{i+1}$.

Now consider any triplet $\tau_i = \langle a_{b_i}, a_{c_i}, a_{d_i} \rangle$ in \mathcal{T} . The monotonicity of Γ and the placement of the squares Q_i and Q_i^α imply that every corner or betweenness edge is assigned to the top or left side of Q_i and to the bottom or right side of Q_i^α . Further, since x^i is the first element in the row-column order of M_i and since $\max_y\{Q_i\} < \min_y\{Q_i^\alpha\}$, no betweenness edge is assigned to the left side of Q_i , as otherwise it would cross edge (x^{i-1}, x^i) . Hence, all the betweenness edges are assigned to the top side of Q_i . Since x^i is the first vertex in the row-column order of M_i and by the planarity of Γ , we have that, when traversing Q_i^α in clockwise direction starting from its top-right corner, the incidence points between Q_i^α and the betweenness edges are encountered before the incidence points between Q_i^α and the corner edges; in particular, no betweenness edge incident to the first vertex – say α^i as other case is analogous – in the row-column order of M_i^α is assigned to the bottom side of M_i^α , as otherwise this edge would cross the corner edge incident to β^i . The planarity of Γ also implies that, when traversing the top side of Q_i from left to right, the end-points of the betweenness edges incident to β^i are encountered all before the end-points of the betweenness edges incident to α^i . Since $v_{c_i}^i$ is the only vertex among $v_{b_i}^i, v_{c_i}^i$, and $v_{d_i}^i$ that is neighbor of both α^i and β^i , then its associated column is between the columns associated to $v_{b_i}^i$ and $v_{d_i}^i$, hence a_{c_i} is between a_{b_i} and a_{d_i} in σ . \square

Theorem 6 MNTLP WITH FIXED SIDES is NP-complete.

Proof: The reduction presented in the proof of Theorem 5, equipped with the side assignment for the inter-cluster edges described in the direction (\implies) implies the statement. \square

Since the instances of MNTLP used in the proof of Theorem 5 are planar whenever they are locally planar, testing the existence of a monotone planar NT representation with fixed square assignment is also NP-complete. Further, the instances of NT PLANARITY used in the proof of Theorem 1 can be drawn planar with straight-line (hence monotone) edges, whenever they are planar. Hence, testing whether a flat clustered graph admits a monotone planar NT representation – without square assignment – is also NP-complete.

Consider now a flat clustered graph $G = (V, E, \mathcal{C})$ and a monotone NT representation Γ of G with fixed square assignment γ . Consider two clusters $V_a, V_b \in \mathcal{C}$ and let $Q_a = \gamma(V_a)$ and $Q_b = \gamma(V_b)$. Since Q_a and Q_b are disjoint,




Figure 10: Possible arrangements for squares Q_a and Q_b . Thick red segments represent sides of Q_a and Q_b edge (u, v) cannot be assigned to. Red curves show further forbidden side assignment pairs for edge (u, v) .

there exists either a vertical or a horizontal line separating them. Suppose that the former holds, the other case being analogous. Also suppose that $\max_x(Q_a) < \min_x(Q_b)$ and $\max_y(Q_a) \geq \max_y(Q_b)$, the other cases being analogous up to reflections of the Cartesian axes (refer to Fig. 10). Also, consider an inter-cluster edge $e = (u, v) \in E_{a,b}$. Depending on the relative positions of Q_a and Q_b in Γ , not all the possible combinations of side assignments for e might be allowed, as described in the following property. Notice that, by the assumptions on the relative positions of Q_a and Q_b in Γ and by the monotonicity and the local planarity of Γ , we have that $s_a(e) \neq L, T$ and $s_a(e) \neq R$.

Property 1 *Let y_u and y_v be the y -coordinate of points m_R^u and m_L^v , respectively. The following three arrangements are possible for Q_a and Q_b in Γ .*

Arrangement 1: $\max_y(Q_b) < \min_y(Q_a)$. Then $s_b(e) \neq B$ and all other four side assignments $\langle s_a(e) = R, s_b(e) = T \rangle$, $\langle s_a(e) = R, s_b(e) = L \rangle$, $\langle s_a(e) = B, s_b(e) = T \rangle$, and $\langle s_a(e) = B, s_b(e) = L \rangle$ are allowed for e .

Arrangement 2: $\min_y(Q_b) < \min_y(Q_a) \leq \max_y(Q_b)$. Then $s_b(e) \neq B$; also, pair $\langle s_a(e) = B, s_b(e) = T \rangle$ is not allowed, while pair $\langle s_a(e) = R, s_b(e) = L \rangle$ is allowed. The remaining two possible pairs $\langle s_a(e) = R, s_b(e) = T \rangle$ and $\langle s_a(e) = B, s_b(e) = L \rangle$ are or are not allowed, depending on y_u and y_v . In particular, if $y_u \leq \max_y(Q_b)$, then $\langle s_a(e) = R, s_b(e) = T \rangle$ is not allowed, otherwise it is; also, if $y_v \geq \min_y(Q_a)$, then $\langle s_a(e) = B, s_b(e) = L \rangle$ is not allowed, otherwise it is.

Arrangement 3: $\min_y(Q_a) \leq \min_y(Q_b)$. Then $s_a(e) \neq B$; also, pair $\langle s_a(e) = R, s_b(e) = L \rangle$ is allowed. The remaining two possible pairs $\langle s_a(e) = R, s_b(e) = T \rangle$ and $\langle s_a(e) = R, s_b(e) = B \rangle$ are or are not allowed, depending on y_u . In particular, if $y_u \leq \max_y(Q_b)$, then $\langle s_a(e) = R, s_b(e) = T \rangle$ is not allowed, otherwise it is, and if $y_u \geq \min_y(Q_b)$, then $\langle s_a(e) = R, s_b(e) = B \rangle$ is not allowed, otherwise it is.

Note that if an edge e can be drawn as an xy -monotone curve not crossing any matrix, then it can also be drawn as a straight-line segment not crossing any matrix, since the pipe of Q_a and Q_b does not intersect any matrix other than M_a and M_b . The next lemma extends this observation by arguing that the xy -monotonicity constraint can be replaced by a straight-line requirement




Figure 11: Illustration for the proof of Lemma 5. (a) Case 1. (b) Case 2.

also for what concerns crossings between inter-cluster edges incident to the same matrix.

Lemma 5 *An instance $\langle G = (V, E, \mathcal{C}), \gamma \rangle$ of MNTLP WITH FIXED ORDER AND FIXED SIDES is locally planar if and only if it admits a monotone locally planar NT representation in which all the inter-cluster edges are drawn as straight-line segments.*

Proof: Since a straight-line segment is an xy -monotone curve, one direction of the proof is trivial. Consider an NT representation Γ of $\langle G = (V, E, \mathcal{C}), \gamma \rangle$ with a fixed row-column order, a fixed side assignment, and a fixed square assignment, in which all the inter-cluster edges are straight-line segments. Suppose that Γ is not locally planar and consider two crossing inter-cluster edges $e = (v_{a,1}, v_b)$ and $f = (v_{a,2}, v_c)$ such that $v_{a,1}$ and $v_{a,2}$ belong to the same cluster $V_a \in \mathcal{C}$. We show that e and f cross in any monotone NT representation Γ' with the same row-column order, side assignment, and square assignment as Γ . Two are the cases: either v_b and v_c belong to the same cluster V_b as in Fig. 11a (Case 1), or they belong to different clusters V_b and V_c , respectively, as in Fig. 11b (Case 2).

In Case 1, consider the region R_{ab} of the plane inside the pipe of Q_a and Q_b and outside each of Q_a and Q_b . Edge e splits R_{ab} into two regions R_1 and R_2 . Since e and f cross in Γ , the end-points of f are one incident to R_1 and one incident to R_2 in Γ' . Since the representation of f has to lie inside R_{ab} , it follows that e and f cross in Γ' .

In Case 2, consider the region R_{ab} defined as in Case 1 and consider the region R_{ac} of the plane inside the pipe of Q_a and Q_c and outside each of Q_a and Q_c . Since e lies in R_{ab} and f in R_{ac} , and since e and f cross in Γ , it follows that the intersection of R_{ab} and R_{ac} is a non-empty region R_\cap . The part of e inside R_\cap partitions R_\cap into two regions R_1 and R_2 . Since e and f cross in Γ , the end-point of f on the boundary of Q_a is incident to the one between R_1 and R_2 that does not share the boundary with region $R_{ac} - R_\cap$. Since the representation of f has to lie inside R_{ac} , it follows that e and f cross in Γ' . \square

The previous lemma, in contrast to the negative results of Theorems 5 and 6, allows us to show that MNTLP WITH FIXED ORDER AND FIXED SIDES is a

polynomial-time solvable problem.

Theorem 7 MNTLP WITH FIXED ORDER AND FIXED SIDES *can be solved in polynomial time.*

Proof: We check whether every edge can be represented as an xy -monotone curve by Property 1. Further, we check whether all the pairs of inter-cluster edges incident to the same cluster admit a non-crossing straight-line drawing; by Lemma 5 this is equivalent to testing the given instance for local planarity with fixed row-column order, fixed side assignment, and fixed square assignment. \square

The remaining piece of the complexity puzzle for MNTLP is the setting with fixed row-column order and free side assignment. Although we are not able to establish the complexity of the corresponding decision problem, we show that testing MNTLP with fixed order is a polynomial-time solvable problem if the number of clusters is constant. In order to do that, we show how to transform the instances of our problem into instances of 2-SAT.

Assuming the hypotheses stated before Property 1 about the relative positions of Q_a and Q_b , we say that an inter-cluster edge $e = (u \in V_a, v \in V_b)$ is S -drawn in Γ if:

- (i) Q_a and Q_b are arranged as in Arrangement 1 of Property 1 and either $\langle s_a(e) = R, s_b(e) = L \rangle$ or $\langle s_a(e) = B, s_b(e) = T \rangle$; or
- (ii) Q_a and Q_b are arranged as in Arrangement 2 of Property 1 and it holds that (a) $\langle s_a(e) = R, s_b(e) = L \rangle$, (b) $y_u > \max_y(Q_b)$, and (c) $y_v < \min_y(Q_a)$.

Note that if Q_a and Q_b are arranged as in Arrangement 3 of Property 1, then e is not S -drawn in Γ , by definition. The representation of an S -drawn edge is an S -drawing. We have the following.

Lemma 6 *Let $\langle G = (V, E, \mathcal{C} = \{V_a, V_b\}), \gamma, \sigma \rangle$ be an instance of MNTLP WITH FIXED ORDER. Consider the following two cases:*

- *Case 1: an inter-cluster edge $e^* \in E$ has a given S -drawing Γ_{e^*} , or*
- *Case 2: no inter-cluster edge in E has an S -drawing.*

Both in Case 1 and in Case 2, we can construct in $O(|E|^2)$ time a 2-SAT formula $\phi(a, b, \Gamma_e)$ and $\phi(a, b)$, respectively, with length $O(|E|^2)$ that is satisfiable if and only if $\langle G, \gamma, \sigma \rangle$ admits a monotone locally planar NT representation with fixed order satisfying the constraint of the corresponding case.

Proof: Consider the squares $Q_a = \gamma(V_a)$ and $Q_b = \gamma(V_b)$. If they are not disjoint, no NT representation of G exists, hence the statement is trivially true. Otherwise, there exists either a vertical line or a horizontal line separating them. Suppose that the former holds, the other case being analogous. Suppose that




Figure 12: Illustrations for the proof of Lemma 6, Case 1, Arrangement 1.

$\max_x(Q_a) < \min_x(Q_b)$ and $\max_y(Q_a) \geq \max_y(Q_b)$, the other cases being analogous up to reflections of the Cartesian axes.

Suppose that an inter-cluster edge e^* is required to have a drawing Γ_e as in Case 1. By the definition of an S-drawn edge, if Q_a and Q_b are arranged as in Arrangement 3 of Property 1, then the required NT representation does not exist, thus the statement trivially holds. Hence, we can assume that Q_a and Q_b are arranged as in Arrangement 1 or 2 of Property 1. Let $e \neq e^* \in E$ be any inter-cluster edge not adjacent to e . Denote by σ_a and σ_b the row-column orders associated to V_a and V_b in σ , respectively.

Consider Arrangement 1 and suppose $s_a(e^*) = R$ and $s_b(e^*) = L$. The end-vertices of e and e^* in V_a (in V_b) have two possible relative positions in σ_a (resp. in σ_b). This leads to four possible combinations for these relative positions.

If $\sigma_a(e^*) < \sigma_a(e)$ and $\sigma_b(e) < \sigma_b(e^*)$, then any xy -monotone curve representing e crosses e^* , independently of the side assignment for e , and the statement trivially holds. See Fig. 12a. For each of the three remaining combinations, *exactly two* side assignments for e create no crossing with e^* . Indeed:

- If $\sigma_a(e) < \sigma_a(e^*)$ and $\sigma_b(e) < \sigma_b(e^*)$, then it holds true that either $s_a(e) = R$ and $s_b(e) = T$, or that $s_a(e) = R$ and $s_b(e) = L$. See Fig. 12b.
- If $\sigma_a(e) < \sigma_a(e^*)$ and $\sigma_b(e^*) < \sigma_b(e)$, then it holds true that either $s_a(e) = R$ and $s_b(e) = T$, or that $s_a(e) = B$ and $s_b(e) = L$. See Fig. 12c.
- If $\sigma_a(e^*) < \sigma_a(e)$ and $\sigma_b(e^*) < \sigma_b(e)$, then it holds true that either $s_a(e) = R$ and $s_b(e) = L$, or that $s_a(e) = B$ and $s_b(e) = L$. See Fig. 12d.

The discussion for the case in which Q_a and Q_b are arranged as in Arrangement 1, $s_a(e^*) = B$, and $s_b(e^*) = T$ is analogous to the previous one.

Consider now Arrangement 2. According to the definition of S-drawing it holds true for $e^* = (u, v)$ that (a) $s_a(e^*) = R$ and $s_b(e^*) = L$, (b) the y -coordinate of p_u is greater than $\max_y(Q_b)$, and (c) the y -coordinate of p_v is smaller than $\min_y(Q_a)$.

Similarly to Arrangement 1, there are four possible combinations for the relative positions of the end-vertices of e and e^* in σ_a and σ_b . If $\sigma_a(e^*) < \sigma_a(e)$ and $\sigma_b(e) < \sigma_b(e^*)$, then any xy -monotone curve representing e crosses e^* , independently of the side assignment for e , and the statement trivially holds.




Figure 13: Illustrations for the proof of Lemma 6, Case 1, Arrangement 2.

See Fig. 13a. For each of the three remaining combinations, *exactly two* side assignments for e create no crossing with e^* .

- If $\sigma_a(e) < \sigma_a(e^*)$ and $\sigma_b(e) < \sigma_b(e^*)$, then it holds true that either $s_a(e) = R$ and $s_b(e) = T$, or that $s_a(e) = R$ and $s_b(e) = L$. See Fig. 13b.
- If $\sigma_a(e) < \sigma_a(e^*)$ and $\sigma_b(e^*) < \sigma_b(e)$, then it holds true that either $s_a(e) = R$ and $s_b(e) = T$, or that $s_a(e) = B$ and $s_b(e) = L$. See Fig. 13c.
- If $\sigma_a(e^*) < \sigma_a(e)$ and $\sigma_b(e^*) < \sigma_b(e)$, then it holds true that either $s_a(e) = R$ and $s_b(e) = L$, or that $s_a(e) = B$ and $s_b(e) = L$. See Fig. 13d.

Hence, for each inter-cluster edge $e \neq e^* \in E$ not adjacent to e^* , there exist two side assignments for e that allow it to be represented as an xy -monotone curve not intersecting e^* .

We are now ready to show, for Case 1 of the lemma, that a monotone locally planar NT representation of $\langle G = (V, E, \mathcal{C} = \{V_a, V_b\}), \gamma \rangle$ in which e^* is represented by Γ_e exists if and only if a suitable 2-SAT formula $\phi(a, b, \Gamma_e)$ is satisfiable.

For each inter-cluster edge $e \neq e^* \in E$ not adjacent to e^* , we define a Boolean variable x_e . The above discussion shows that, if we did not conclude that a trivially false formula exists, then there are exactly two distinct side assignments for e . We select one arbitrarily, which we call *canonical side assignment*, and associate $x_e = \text{TRUE}$ to it and $x_e = \text{FALSE}$ to the other.

For each pair of non-adjacent inter-cluster edges $e_1, e_2 \neq e^* \in E$, consider the four possible side assignments for them. We add to $\phi(a, b, \Gamma_e)$ at most four clauses defined as follows.

- If the canonical side assignment for e_1 and the canonical side assignment for e_2 generate a crossing between e_1 and e_2 , then we add clause $\{\overline{x_{e_1}} \vee \overline{x_{e_2}}\}$ to $\phi(a, b, \Gamma_e)$.
- If the canonical side assignment for e_1 and the non-canonical side assignment for e_2 generate a crossing between e_1 and e_2 , then we add clause $\{\overline{x_{e_1}} \vee x_{e_2}\}$ to $\phi(a, b, \Gamma_e)$.
- If the non-canonical side assignment for e_1 and the canonical side assignment for e_2 generate a crossing between e_1 and e_2 , then we add clause $\{x_{e_1} \vee \overline{x_{e_2}}\}$ to $\phi(a, b, \Gamma_e)$.




Figure 14: Illustrations for the proof of Lemma 6, Case 2, Arrangement 2.

- If the non-canonical side assignment for e_1 and the non-canonical side assignment for e_2 generate a crossing between e_1 and e_2 , then we add clause $\{x_{e_1} \vee x_{e_2}\}$ to $\phi(a, b, \Gamma_e)$.

As a consequence of the above discussion $\langle G = (V, E, \mathcal{C} = \{V_a, V_b\}), \gamma \rangle$ admits a monotone locally planar NT representation in which e^* is represented by Γ_e if and only if $\phi(a, b, \Gamma_e)$ is satisfiable. Further, since the number of clauses in $\phi(a, b, \Gamma_e)$ is upper-bounded by $O(|E|^2)$ and since it can be determined in constant time whether a side assignment for any two edges produces a crossing, then formula $\phi(a, b, \Gamma_e)$ can be constructed in $O(|E|^2)$ time and has $O(|E|^2)$ size. Since 2-SAT formulae can be tested for satisfiability in linear time [9], the statement of Case 1 follows.

Suppose now that Case 2 of the statement holds. According to Property 1, squares Q_a and Q_b can be arranged as in Arrangement 1, 2, or 3.

Consider Arrangement 1. By the hypothesis of the case, no edge is allowed to be S-drawn. Hence, for each inter-cluster edge e , we have either $s_a(e) = R$ and $s_b(e) = T$ or $s_a(e) = B$ and $s_b(e) = L$.

Consider Arrangement 2. Let $e = (u, v)$ be an inter-cluster edge. We distinguish four cases depending on the y -coordinate y_u of m_R^u with respect to $\max_y(Q_b)$ and on the y -coordinate y_v of m_L^v with respect to $\min_y(Q_a)$. In each of the four cases, *at most two* side assignments for e are possible so that e is not S-drawn.

- If $y_u > \max_y(Q_b)$ and $y_v \geq \min_y(Q_a)$, then it holds true that either $s_a(e) = R$ and $s_b(e) = T$, or that $s_a(e) = R$ and $s_b(e) = L$. See Fig. 14a.
- If $y_u > \max_y(Q_b)$ and $y_v < \min_y(Q_a)$, then it holds true that either $s_a(e) = R$ and $s_b(e) = T$, or that $s_a(e) = B$ and $s_b(e) = L$. See Fig. 14b; notice that the side assignment $s_a(e) = R$ and $s_b(e) = L$ would imply that e is S-drawn, which is not possible by hypothesis.
- If $y_u \leq \max_y(Q_b)$ and $y_v \geq \min_y(Q_a)$, then it holds true that $s_a(e) = R$ and $s_b(e) = L$. See Fig. 14c.
- If $y_u \leq \max_y(Q_b)$ and $y_v < \min_y(Q_a)$, then it holds true that either $s_a(e) = R$ and $s_b(e) = L$, or that $s_a(e) = B$ and $s_b(e) = L$. See Fig. 14d.

Consider Arrangement 3. Let $e = (u, v)$ be an inter-cluster edge. By definition e cannot be S-drawn. We distinguish three cases depending on the




Figure 15: Illustrations for the proof of Lemma 6, Case 2, Arrangement 3.

y -coordinate y_u of m_r^u with respect to $\min_y(Q_b)$ and $\max_y(Q_b)$. In each of the three cases, *at most two* side assignments for e are possible.

- If $y_u > \max_y(Q_b)$, then it holds true that either $s_a(e) = R$ and $s_b(e) = T$, or that $s_a(e) = R$ and $s_b(e) = L$. See Fig. 15a.
- If $\min_y(Q_b) \leq y_u \leq \max_y(Q_b)$, then it holds true that $s_a(e) = R$ and $s_b(e) = L$. See Fig. 15b.
- If $y_u < \min_y(Q_b)$, then it holds true that either $s_a(e) = R$ and $s_b(e) = L$, or that $s_a(e) = R$ and $s_b(e) = B$. See Fig. 15c.

Hence, regardless of whether Q_a and Q_b are arranged as in Arrangement 1, 2, or 3, and regardless of the y -coordinate of m_r^u and m_l^v , there exist at most two side assignments for e that allow it to be represented as an xy -monotone curve.

The construction of the 2-SAT formula and the bound on its size can be derived analogously to Case 1; the only difference is that, when only one side assignment is possible, a clause with a single literal is generated. This concludes the proof of the lemma. \square

We now turn to the study of flat clustered graphs with three clusters.

Lemma 7 *Let $\langle G = (V, E, \mathcal{C} = \{V_a, V_b, V_c\}), \gamma, \sigma \rangle$ be an instance of MNTLP WITH FIXED ORDER. Consider the four cases that are generated by assuming that an edge $e^* \in E_{a,b}$ has a prescribed S -drawing or not and that an edge $f^* \in E_{a,c}$ has a prescribed S -drawing or not. In each case, we can construct in $O(|E|^2)$ time a 2-SAT formula $\phi(a, b, c)$ with length $O(|E|^2)$ that is satisfiable if and only if $\langle G, \gamma, \sigma \rangle$ admits a monotone NT representation with fixed order that satisfies the constraints of the corresponding case, such that no inter-cluster edge intersects any matrix except at its incidence points, and such that there are no two edges, one in $E_{a,b}$ and one in $E_{a,c}$, that cross each other.*

Proof: In each of the four cases, the hypotheses lead us in either Case 1 or Case 2 of Lemma 6 for the edges in $E_{a,b}$ and the same holds for the edges in $E_{a,c}$. Hence, by Lemma 6, each of these edges admits at most two side assignments in each case. Moreover, each of these side assignments corresponds to a directed or negated literal. For each pair of edges $e \in E_{a,b}$ and $f \in E_{a,c}$ and for each of the at most four side assignments for them, we exploit Lemma 5 to test whether a side assignment for e and f leads to a crossing and in the case of a crossing we introduce suitable clauses to rule out that side assignment. \square

We finally get the following.

Theorem 8 MNTLP WITH FIXED ORDER *can be tested in $|E|^{O(|\mathcal{C}|^2)}$ time for an instance $\langle G = (V, E, \mathcal{C}), \gamma, \sigma \rangle$.*

Proof: For each pair V_a, V_b of adjacent clusters in \mathcal{C} , we guess whether V_a, V_b belongs to a set \mathcal{P}_S or to a set \mathcal{P}_N . The set \mathcal{P}_S contains all the pairs V_a, V_b of clusters that have an inter-cluster edge that is S-drawn. The set \mathcal{P}_N contains all the pairs V_a, V_b of clusters that do not have an inter-cluster edge that is S-drawn. For each pair V_a, V_b of clusters in \mathcal{P}_S we guess an inter-cluster edge $e \in E_{a,b}$ that can be S-drawn and one of its possible S-drawings Γ_e for e ; we remark that the guess of Γ_e consists of a guess of the side assignment for e , hence there are a constant number of possible guesses for each edge e .

By means of Lemma 6 we compute the following formula:

$$\phi_{\text{pairs}} = \bigwedge_{V_a, V_b \in \mathcal{P}_S} \phi(a, b, \Gamma_e) \bigwedge_{V_a, V_b \in \mathcal{P}_N} \phi(a, b).$$

Further, let $\mathcal{P}_{\text{triplet}}$ be the set of triplets V_a, V_b, V_c of clusters in \mathcal{C} such that V_b and V_c are adjacent to V_a . We write one of the formulae $\phi(a, b, c)$ of the four cases of Lemma 7 according to the presence in \mathcal{P}_S of an inter-cluster edge between V_a and V_b or of an inter-cluster edge between V_a and V_c . By means of Lemma 7, we compute the following:

$$\phi_{\text{triplets}} = \bigwedge_{V_a, V_b, V_c \in \mathcal{P}_{\text{triplet}}} \phi(a, b, c).$$

Finally, we define

$$\phi = \phi_{\text{pairs}} \wedge \phi_{\text{triplets}}.$$

We have that instance $\langle G = (V, E, \mathcal{C}), \gamma, \sigma \rangle$ is a positive instance if and only if there exists a guess such that the corresponding formula ϕ is satisfiable.

About the time complexity, for each guess $O(|E|^2)$ time is needed to compute the corresponding formula ϕ and to check it for satisfiability, due to Lemmata 6 and 7. The number of guesses can be bounded as follows. For each pair of adjacent clusters V_a, V_b we have to guess among $2|E_{a,b}| + 1$ possibilities, corresponding to the choice of $|E_{a,b}|$ edges to be S-drawn, each in two possible ways, plus the possibility of not having any S-drawn edge. This leads to $O((2|E| + 1)^{|\mathcal{C}|^2})$, which is in $|E|^{O(|\mathcal{C}|^2)}$, guesses. \square

Observe that the computational complexity of the algorithm described in the proof of Theorem 8 is polynomial if the number of clusters is constant.

5 A JavaScript Library for Monotone NodeTrix Representations

In this section we consider the following scenario. A user moves (e.g. via a drag-and-drop primitive) the matrices representing the clusters of a flat clustered

graph, choosing also her preferred row-column order. A system automatically selects sides for the inter-cluster edges so to produce a monotone NT representation Γ with a “small” $\chi(\Gamma)$.

The algorithm in the proof of Theorem 8 suggests the following strategy:

1. Compute the 2-SAT formula associated to each possible guess of S-drawings of edges between adjacent matrices. In each formula the value of a variable represents the two possible side assignments for an edge; further, each unsatisfied clause corresponds to a crossing in the monotone NT representation.
2. If one of such formulae admits a solution, draw the edges according to the values of the associated variables, obtaining an MNTLP representation.
3. Otherwise, for each formula, heuristically search for a solution of the corresponding MAX-2-SAT problem and keep the solution with the smallest number of false clauses, corresponding to a drawing with few local crossings.

Such a strategy requires polynomial time if the number of clusters is constant (Theorem 8) and the selected MAX-2-SAT heuristic is polynomial. However, solving a MAX-2-SAT instance for each of the guesses of the proof of Theorem 8 is unfeasible even in a static setting.

Therefore, we modify the above strategy as follows. We restrict to monotone NT representations without S-drawn edges. A locally-planar flat clustered graph may become non-planar with this restriction, hence this choice corresponds to trading accuracy for efficiency. However, the proof of Theorem 8 shows that in this setting there is a unique formula associated to an instance, hence we need to solve one MAX-2-SAT instance.

A JavaScript library implementing the above heuristic has been designed and used in a proof-of-concept editor available at [3] (see Fig. 16). The internal part of the matrices is not shown in the editor, and inter-cluster edges are polylines or splines; this is not intended to be the best choice and many alternatives for the actual geometry of the edges are possible whose visual appeal should be considered according to the specific application domain.

We performed the following experiments in order to check if the strategy is usable on medium-size instances, that is, on instances whose size is reasonable with respect to the need of guaranteeing the visibility of the matrix labels on the display of a laptop. We computed the running time of the layout algorithm on 1,200 instances, each composed of two 30-vertex matrices. The instances are grouped into 30 sets I_1, \dots, I_{30} , such that instances in set I_k contain k random inter-cluster edges between the two matrices. Our current implementation, which is not meant to achieve good time performance, computes the layout within 3 ms on average for all sets (we used the Google Chrome 62.03 browser on a 64 bit 2.2 GHz Intel Core i7 with 16 GB RAM). This implies that an instance with ten 30-vertex matrices of which 20 pairs are connected by 30 inter-cluster edges would allow 15 frames per second, which enables a fluid interaction.




Figure 16: A NodeTrix Representation created by the demo editor [3].

We did not compare the number of crossings produced by our heuristic with other approaches because, as far as we know, this is the first attempt to reduce local crossings in NT representations. Also, it would be pointless to compare our approach with the original NT software, since in that case the edges just attach to the nearest sides. The JavaScript software of our library is freely available and can be integrated in any NodeTrix-style interface. As an example, it can be coupled with an algorithm that automatically places matrices based on a force-directed approach or with one that computes row-column order for the matrices with the purpose of clarifying the internal structure of the clusters [14].

6 Conclusions and Open problems

We have shown that clustered graphs for Nodetrix planarity is NP-complete even if the order of the rows and columns is fixed or if the matrix sides to which the inter-cluster edges attach is fixed. We have also studied the setting where matrices have fixed positions and inter-cluster edges are xy -monotone curves. In this case we established negative and positive results; leveraging on the latter, we developed a library that computes a layout of the inter-cluster edges with few crossings. A demo [3] shows that the computation allows the user to move matrices without any slowdown of the interaction.

Several theoretical problems are related to the planarity of Nodetrix representations. First, the NP-completeness of Nodetrix planarity can be interpreted as a proof of the NP-completeness of clustered planarity (see, for exam-

ple, [5, 7, 12, 15]) when a specific type of representation is required. Observe, though, that a flat clustered graph may be Nodetrix planar even if its underlying graph is not planar. Second, planarity of hybrid representations have been recently studied [6] in the setting in which clusters are represented as the intersections of geometric objects. Our results can be viewed as a further progress in this area. Third, given a flat clustered graph with two clusters, computing a locally planar Nodetrix representation in which the clusters are represented as matrices aligned along their principal diagonal is equivalent to solve the 2-page bipartite book embedding with spine crossings problem [6]. Interestingly, if the two matrices are aligned along their secondary diagonal this equivalence is not evident anymore.

Among the future research directions, we mention the one of automatically embedding the matrices in order to minimize the number of crossings in monotone Nodetrix representations.

References

- [1] CiteVis: Visualizing citations among InfoVis conference papers. URL: <http://www.cc.gatech.edu/gvu/ii/citevis>.
- [2] NodeTriX. URL: <http://www.aviz.fr/Research/Nodetrix>.
- [3] NodeTriX Representations: a proof-of-concept editor. URL: <http://www.dia.uniroma3.it/~dalozzo/projects/matrix>.
- [4] P. Angelini, E. Colasante, G. Di Battista, F. Frati, and M. Patrignani. Monotone drawings of graphs. *Journal of Graph Algorithms and Applications*, 16(1):5–35, 2012. doi:10.7155/jgaa.00249.
- [5] P. Angelini, G. Da Lozzo, G. Di Battista, F. Frati, M. Patrignani, and V. Roselli. Relaxing the constraints of clustered planarity. *Comput. Geom.*, 48(2):42–75, 2015. doi:10.1016/j.comgeo.2014.08.001.
- [6] P. Angelini, G. Da Lozzo, G. Di Battista, F. Frati, M. Patrignani, and I. Rutter. Intersection-link representations of graphs. *Journal of Graph Algorithms and Applications*, 21(4):731–755, 2017. doi:10.7155/jgaa.00437.
- [7] P. Angelini, G. Da Lozzo, G. Di Battista, F. Frati, and V. Roselli. The importance of being proper: (In clustered-level planarity and T-level planarity). *Theor. Comput. Sci.*, 571:1–9, 2015. doi:10.1016/j.tcs.2014.12.019.
- [8] P. Angelini, G. Da Lozzo, and D. Neuwirth. Advancements on SEFE and partitioned book embedding problems. *Theor. Comput. Sci.*, 575:71–89, 2015. doi:10.1016/j.tcs.2014.11.016.
- [9] B. Aspvall, M. F. Plass, and R. E. Tarjan. A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Inf. Process. Lett.*, 8(3):121–123, 1979. doi:10.1016/0020-0190(79)90002-4.
- [10] V. Batagelj, F. Brandenburg, W. Didimo, G. Liotta, P. Palladino, and M. Patrignani. Visual analysis of large graphs using (X,Y)-clustering and hybrid visualizations. *IEEE Trans. Vis. Comput. Graph.*, 17(11):1587–1598, 2011. doi:10.1109/TVCG.2010.265.
- [11] M. Chimani, G. Di Battista, F. Frati, and K. Klein. Advances on testing c-planarity of embedded flat clustered graphs. In C. Duncan and A. Symvonis, editors, *22nd International Symposium on Graph Drawing (GD '14)*, volume 8871 of *LNCS*, pages 416–427, 2014. doi:10.1007/978-3-662-45803-7_35.
- [12] P. Cortese, G. Di Battista, F. Frati, M. Patrignani, and M. Pizzonia. C-planarity of c-connected clustered graphs. *Journal of Graph Algorithms and Applications*, 12(2):225–262, 2008. doi:10.7155/jgaa.00165.

- [13] G. Di Battista and F. Frati. Efficient c-planarity testing for embedded flat clustered graphs with small faces. *Journal of Graph Algorithms and Applications*, 13(3):349–378, 2009. doi:10.7155/jgaa.00191.
- [14] P. Doreian, V. Batagelj, and A. Ferligoj. *Generalized Blockmodeling*. Cambridge University Press, New York, NY, USA, 2005.
- [15] Q. Feng, R. F. Cohen, and P. Eades. Planarity for clustered graphs. In P. G. Spirakis, editor, *Algorithms - ESA '95, Third Annual European Symposium, Corfu, Greece, September 25-27, 1995, Proceedings*, volume 979 of *LNCS*, pages 213–226. Springer, 1995. doi:10.1007/3-540-60313-1_145.
- [16] R. Fulek, M. J. Pelsmayer, M. Schaefer, and D. Štefankovič. Hanani–tutte, monotone drawings, and level-planarity. In J. Pach, editor, *Thirty Essays on Geometric Graph Theory*, pages 263–287. Springer New York, 2013. doi:10.1007/978-1-4614-0110-0_14.
- [17] C. Gutwenger, K. Klein, and P. Mutzel. Planarity testing and optimal edge insertion with embedding constraints. *Journal of Graph Algorithms and Applications*, 12(1):73–95, 2008. doi:10.7155/jgaa.00160.
- [18] N. Henry, A. Bezerianos, and J. Fekete. Improving the readability of clustered social networks using node duplication. *IEEE Trans. Vis. Comput. Graph.*, 14(6):1317–1324, 2008. doi:10.1109/TVCG.2008.141.
- [19] N. Henry, J. Fekete, and M. J. McGuffin. NodeTrix: a Hybrid Visualization of Social Networks. *IEEE Trans. Vis. Comput. Graph.*, 13(6):1302–1309, 2007. doi:10.1109/TVCG.2007.70582.
- [20] J. Opatrny. Total ordering problem. *SIAM J. Comput.*, 8(1):111–114, 1979. doi:10.1137/0208008.
- [21] J. Pach and G. Tóth. Monotone drawings of planar graphs. *Journal of Graph Theory*, 46(1):39–47, 2004. doi:10.1002/jgt.10168.
- [22] J. P. Spinrad. Recognition of circle graphs. *J. Algorithms*, 16(2):264–282, 1994. doi:10.1006/jagm.1994.1012.
- [23] W. Unger. On the k-colouring of circle-graphs. In R. Cori and M. Wirsing, editors, *5th Annual Symposium on Theoretical Aspects of Computer Science (STACS '88)*, volume 294 of *LNCS*, pages 61–72. Springer, 1988. doi:10.1007/BFb0035832.