Relaxing the Constraints of Clustered Planarity

Patrizio Angelini[a], Giordano Da Lozzo[a], Giuseppe Di Battista[a], Fabrizio Frati[b], Maurizio Patrignani[a], Vincenzo Roselli[a]

[a] *Dipartimento di Informatica e Automazione, Roma Tre University, Italy*
[b] *School of Information Technologies, The University of Sydney, Australia*

Abstract

In a drawing of a clustered graph vertices and edges are drawn as points and curves, respectively, while clusters are represented by simple closed regions. A drawing of a clustered graph is c-planar if it has no edge-edge, edge-region, or region-region crossings. Determining the complexity of testing whether a clustered graph admits a c-planar drawing is a long-standing open problem in the Graph Drawing research area. An obvious necessary condition for c-planarity is the planarity of the graph underlying the clustered graph. However, this condition is not sufficient and the consequences on the problem due to the requirement of not having edge-region and region-region crossings are not yet fully understood.

In order to shed light on the c-planarity problem, we consider a relaxed version of it, where some kinds of crossings (either edge-edge, edge-region, or region-region) are allowed even if the underlying graph is planar. We investigate the relationships among the minimum number of edge-edge, edge-region, and region-region crossings for drawings of the same clustered graph. Also, we consider drawings in which only crossings of one kind are admitted. In this setting, we prove that drawings with only edge-edge or with only edge-region crossings always exist, while drawings with only region-region crossings may not. Further, we provide upper and lower bounds for the number of such crossings. Finally, we give a polynomial-time algorithm to test whether a drawing with only region-region crossings exists for biconnected graphs, hence identifying a first non-trivial necessary condition for c-planarity that can be tested in polynomial time for a noticeable class of graphs.

Keywords: graph drawing, clustered planarity, planar graphs, NP-hardness

Email addresses: angelini@dia.uniroma3.it (Patrizio Angelini), dalozzo@dia.uniroma3.it (Giordano Da Lozzo), gdb@dia.uniroma3.it (Giuseppe Di Battista), brillo@it.usyd.edu.au (Fabrizio Frati), patrigna@dia.uniroma3.it (Maurizio Patrignani), roselli@dia.uniroma3.it (Vincenzo Roselli)

1. Introduction

Clustered planarity is a classical Graph Drawing topic (see [6] for a survey). A *clustered graph* $C(G, T)$ consists of a graph G and of a rooted tree T whose leaves are the vertices of G. Such a structure is used to enrich the vertices of the graph with hierarchical information. In fact, each internal node μ of T represents the subset, called *cluster*, of the vertices of G that are the leaves of the subtree of T rooted at μ. Tree T, which defines the inclusion relationships among clusters, is called *inclusion tree*, while G is the *underlying graph* of $C(G, T)$.

In a *drawing* of a clustered graph $C(G, T)$ vertices and edges of G are drawn as points and open curves, respectively, and each node μ of T is represented by a simple closed region $R(\mu)$ containing exactly the vertices of μ. Also, if μ is a descendant of a node ν, then $R(\nu)$ contains $R(\mu)$.

A drawing of C can have three types of crossings. *Edge-edge crossings* are crossings between edges of G. Algorithms to produce drawings allowing edge-edge crossings have already been proposed (see, for example, [11] and Fig. 2(a)). Two kinds of crossings involve regions, instead. Consider an edge e of G and a node μ of T. If e intersects the boundary of $R(\mu)$ only once, this is not considered as a crossing since there is no way of connecting the endpoints of e without intersecting the boundary of $R(\mu)$. On the contrary, if e intersects the boundary of $R(\mu)$ more than once, we have *edge-region crossings*. An example of this kind of crossings is provided by Fig. 2(b), where edge (u, w) traverses $R(\mu)$ and edge (u, v) exits and enters $R(\nu)$. Finally, consider two nodes μ and ν of T; if the boundary of $R(\mu)$ intersects the boundary of $R(\nu)$, we have a *region-region crossing* (see Fig. 2(c) for an example).

A drawing of a clustered graph is *c-planar* if it does not have any edge-edge, edge-region, or region-region crossing. A clustered graph is *c-planar* if it admits a c-planar drawing.

In the last decades c-planarity has been deeply studied. While the complexity of deciding if a clustered graph is c-planar is still an open problem in the general case, polynomial-time algorithms have been proposed to test c-planarity and produce c-planar drawings under several kinds of restrictions, such as:

- Assuming that each cluster induces a small number of connected components ([5, 7, 10, 16, 17, 21, 22, 24, 25]). In particular, the case in which the graph is *c-connected*, that is, for each node ν of T the graph induced by the vertices of ν is connected, has been deeply investigated.
- Considering only *flat* hierarchies, i.e., the height of T is two, namely no cluster different from the root contains other clusters ([8, 9, 12]).
- Focusing on particular families of underlying graphs ([8, 9, 26]).
- Fixing the embedding of the underlying graph ([12, 24]).

This huge body of research can be read as a collection of polynomial-time testable sufficient conditions for c-planarity.

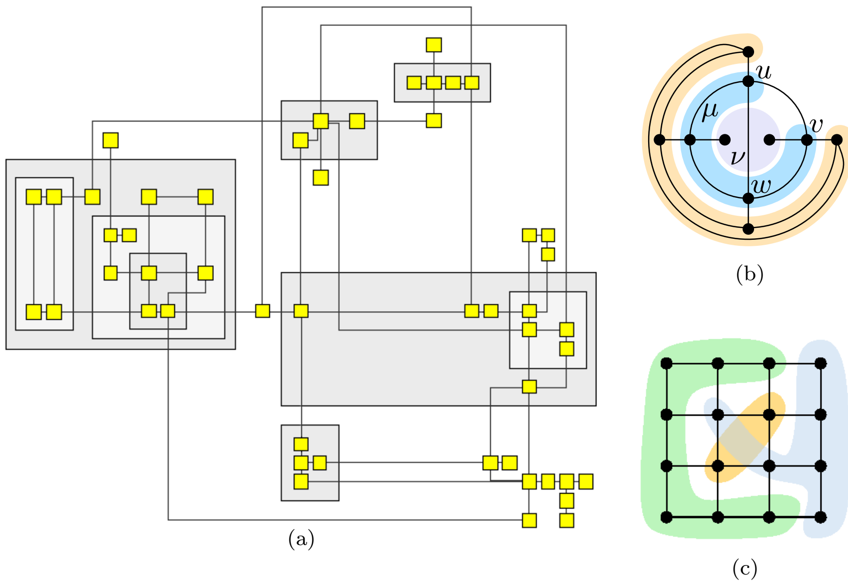

Figure 1: Examples of crossings in drawings of clustered graphs. (a) A drawing obtained with the planarization algorithm described in [11] and containing three edge-edge crossings. (b) A drawing with two edge-region crossings. (c) A drawing with a region-region crossing.

In contrast, the planarity of the underlying graph is the only polynomial-time testable necessary condition that has been found so far for c-planarity in the general case. Such a condition, however, is not sufficient and the consequences on the problem due to the requirement of not having edge-region and region-region crossings are not yet fully understood.

Other known necessary conditions are either trivial (i.e., satisfied by all clustered graphs) or of unknown complexity as the original problem is. An example of the first kind is the existence of a c-planar clustered graph obtained by splitting some cluster into sibling clusters [2]. An example of the second kind, which is also a sufficient condition, is the existence of a set of edges that, if added to the underlying graph, make the clustered graph c-connected and c-planar [16].

In this paper we study a relaxed model of c-planarity. Namely, we study $\langle\alpha,\beta,\gamma\rangle$-drawings of clustered graphs. In an $\langle\alpha,\beta,\gamma\rangle$-drawing the number of edge-edge, edge-region, and region-region crossings is equal to α, β, and γ, respectively. Figs. 2(a), 2(b), and 2(c) show examples of a $\langle 3,0,0\rangle$-drawing, a $\langle 0,2,0\rangle$-drawing, and a $\langle 0,0,1\rangle$-drawing, respectively. Notice that this model provides a generalization of c-planarity, as the traditional c-planar drawing is a special case of an $\langle\alpha,\beta,\gamma\rangle$-drawing where $\alpha=\beta=\gamma=0$. Hence, we can say that the existence of an $\langle\alpha,\beta,\gamma\rangle$-drawing, for some values of α, β, and γ, is a necessary condition for c-planarity.

In our study we focus on clustered graphs whose underlying graph is planar.

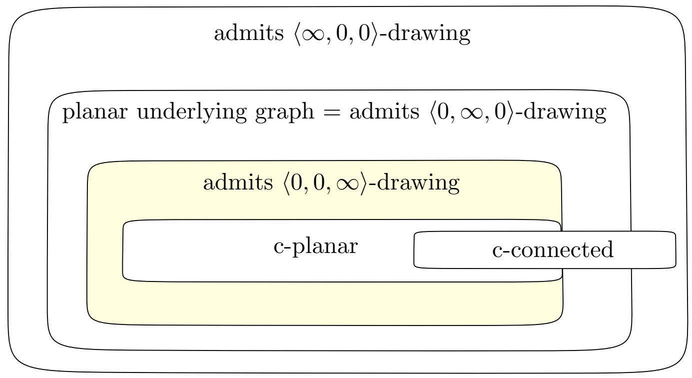

Figure 2: Containment relationships among instances of clustered planarity. The existence of a $\langle 0, 0, \infty\rangle$-drawing is a necessary condition for c-planarity. Note that any $\langle 0, 0, \infty\rangle$-drawing of a c-connected clustered graph $C(G, T)$ can be suitably modified to obtain a c-planar drawing of $C(G, T)$.

We mainly concentrate on the existence of drawings in which only one type of crossings is allowed. We call these drawings $\langle\infty, 0, 0\rangle$-, $\langle 0, \infty, 0\rangle$-, and $\langle 0, 0, \infty\rangle$-drawings, respectively. Our investigation uncovers that allowing different types of crossings has a different impact on the existence of drawings of clustered graphs (see Fig. 2). In particular, we prove that, while every clustered graph admits an $\langle\infty, 0, 0\rangle$-drawing (even if its underlying graph is not planar) and a $\langle 0, \infty, 0\rangle$-drawing (only if its underlying graph is planar), there exist clustered graphs not admitting any $\langle 0, 0, \infty\rangle$-drawing. Further, we provide a polynomial-time testing algorithm to decide whether a biconnected clustered graph admits a $\langle 0, 0, \infty\rangle$-drawing. From this fact we conclude that the existence of such a drawing is the first non-trivial necessary condition for the c-planarity of clustered graphs that can be tested efficiently. This allows us to further restrict the search for c-planar instances with respect to the obvious condition that the underlying graph is planar.

Also, we investigate the relationships among the minimum number of edge-edge, edge-region, and region-region crossings for drawings of the same clustered graph, showing that, in most of the cases, the fact that a clustered graph admits a drawing with few crossings of one type does not imply that such a clustered graph admits a drawing with few crossings of another type.

Finally, we show that minimizing the sum $\alpha + \beta + \gamma$ in a $\langle\alpha, \beta, \gamma\rangle$-drawing of a clustered graph is an NP-complete problem. Since in our construction it is possible to replace each crossing of any type with a crossing of a different type, this implies that the problems of minimizing crossings in $\langle\infty, 0, 0\rangle$-, $\langle 0, \infty, 0\rangle$-, and $\langle 0, 0, \infty\rangle$-drawings are also NP-complete. However, for the first two types of drawings we can prove NP-completeness even for simpler classes of clustered graphs.

We remark that drawings of clustered graphs where a few intersections are admitted may meet the requirements of many typical Graph Drawing appli-

c-c	flat	$\langle \alpha, 0, 0 \rangle$		$\langle 0, \beta, 0 \rangle$		$\langle 0, 0, \gamma \rangle$	
		α UB	α LB	β UB	β LB	γ UB	γ LB
NO	NO	$O(n^2)$ Th.1	$\Omega(n^2)$	$O(n^3)$ Th.2	$\Omega(n^2)$	$O(n^3)^{\maltese}$ Th.5	$\Omega(n^3)$ Th.12
NO	YES	$O(n^2)$	$\Omega(n^2)$	$O(n^2)$ Th.2	$\Omega(n^2)$ Cor.1	$O(n^2)^{\maltese}$ Th.5	$\Omega(n^2)$ Th.11
YES	NO	$O(n^2)$	$\Omega(n^2)$	$O(n^2)$ Th.3	$\Omega(n^2)$ Th.9	$0^{\maltese}$ [16]	$0^{\maltese}$ [16]
YES	YES	$O(n^2)$	$\Omega(n^2)$ Th.7	$O(n)$ Th.3	$\Omega(n)$ Th.10	$0^{\maltese}$ [16]	$0^{\maltese}$ [16]

Table 1: Upper and lower bounds for the number of crossings in $\langle \infty, 0, 0 \rangle$-, $\langle 0, \infty, 0 \rangle$-, and $\langle 0, 0, \infty \rangle$-drawings of clustered graphs. Flags *c-c* and *flat* mean that the clustered graph is *c-connected* and that the cluster hierarchy is *flat*, respectively. Results written in gray derive from those in black, while a "$\maltese$" means that there exist clustered graphs not admitting the corresponding drawings. A "0" occurs if the clustered graph is c-planar.

cations, and that their employment is encouraged by the fact that the class of c-planar instances might be too small to be relevant for some application contexts.

More in detail, we present the following results (recall that we assume the necessary condition that the underlying graph is planar to be always satisfied):

1. In Section 3 we provide algorithms to produce $\langle \infty, 0, 0 \rangle$-, $\langle 0, \infty, 0 \rangle$-, and $\langle 0, 0, \infty \rangle$-drawings of clustered graphs, if they exist. In particular, while $\langle \infty, 0, 0 \rangle$- and $\langle 0, \infty, 0 \rangle$-drawings always exist, we show that some clustered graphs do not admit any $\langle 0, 0, \infty \rangle$-drawing, and we present a polynomial-time algorithm to test whether a biconnected clustered graph admits a $\langle 0, 0, \infty \rangle$-drawing, which is a necessary condition for c-planarity. The algorithm, whose approach is reminiscent of [1], is based on a characterization of the planar embeddings that lead to $\langle 0, 0, \infty \rangle$-drawings, and on a subsequent structural characterization of the existence of a $\langle 0, 0, \infty \rangle$-drawing for any biconnected clustered graph $C(G, T)$, based on the SPQR-tree decomposition of G.

2. The above mentioned algorithms provide upper bounds on the number of crossings for the three kinds of drawings. We show that the majority of these upper bounds are tight by providing matching lower bounds in Section 4. These results are summarized in Tab. 1.

3. In Section 5 we show that there are clustered graphs admitting drawings with one crossing of a certain type but requiring many crossings in drawings where only different types of crossings are allowed. For example, there are clustered graphs that admit a $\langle 1, 0, 0 \rangle$-drawing and that require $\beta \in \Omega(n^2)$ in any $\langle 0, \beta, 0 \rangle$-drawing and $\gamma \in \Omega(n^2)$ in any $\langle 0, 0, \gamma \rangle$-drawing. See Tab. 2 for a summary of these results.

4. In Section 6 we present several complexity results. Namely, we show that:

    - minimizing $\alpha + \beta + \gamma$ in an $\langle \alpha, \beta, \gamma \rangle$-drawing is NP-complete even if the underlying graph is planar, namely a forest of star graphs;

→	$\langle \alpha, 0, 0 \rangle$	$\langle 0, \beta, 0 \rangle$	$\langle 0, 0, \gamma \rangle$
$\langle 1, 0, 0 \rangle$		$\Omega(n^2)$	$\Omega(n^2)$
$\langle 0, 1, 0 \rangle$	$\Omega(n)$		$\Omega(n^2)$
$\langle 0, 0, 1 \rangle$	$\Omega(n^2)$	$\Omega(n)$	

Table 2: Relationships between types of drawings proved in Theorem 14.

- minimizing α in an $\langle \alpha, 0, 0 \rangle$-drawing is NP-complete even if the underlying graph is a matching;
- minimizing β in a $\langle 0, \beta, 0 \rangle$-drawing is NP-complete (see also [18]) even for c-connected flat clustered graphs in which the underlying graph is a triconnected planar multigraph;

Section 2 gives definitions and preliminary lemmas, while Section 7 contains conclusions and open problems.

2. Preliminaries

We remark that every clustered graph $C(G, T)$ that is considered in this paper is such that G is planar.

Let $C(G, T)$ be a clustered graph. If μ is an internal node of T, we denote by $V(\mu)$ the leaves of the subtree of T rooted at μ. The subgraph of G induced by $V(\mu)$ is denoted by $G(\mu)$.

Some constraints are usually enforced on the crossings among the open curves representing edges in a drawing of a graph. Namely: $(\mathcal{C}_1)$ the intersections among curves form a set of isolated points; $(\mathcal{C}_2)$ no three curves intersect in the same point; and $(\mathcal{C}_3)$ two intersecting curves appear alternatingly in the circular order around their intersection point. Figure 3(a) shows a legal crossing, while Figures 3(b)-(d) show crossings violating Constraints $\mathcal{C}_1$, $\mathcal{C}_2$, and $\mathcal{C}_3$, respectively. These constraints naturally extend to encompass crossings involving regions representing clusters, by considering, for each region, the closed curve that forms its boundary.

Let Γ be a drawing of a clustered graph $C(G, T)$. First, we formally define the types of crossings of Γ and how to count them.

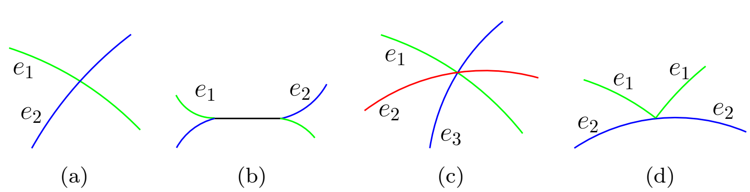

Figure 3: Allowed and forbidden crossings in a drawing of a graph. (a) A legal crossing. (b) A crossing violating Constraint $\mathcal{C}_1$. (c) A crossing violating Constraint $\mathcal{C}_2$. (d) A crossing violating Constraint $\mathcal{C}_3$.

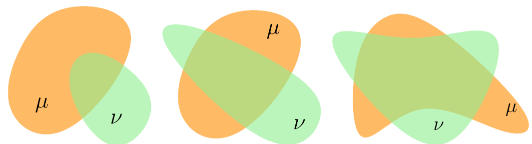

(a) (b) (c)

Figure 4: Examples of intersections between clusters generating (a) zero rr-crossings; (b) one rr-crossing; and (c) two rr-crossings.

Edge-edge crossings. Each crossing between two edges of G is an *edge-edge crossing* (or *ee-crossing* for short) of Γ.

Edge-region crossings. An *edge-region crossing* (*er-crossing*) is a crossing involving an edge e of G and a region $R(\mu)$ representing a cluster μ of T; namely, if e crosses the boundary of $R(\mu)$ k times, the number of er-crossings between e and $R(\mu)$ is $\lfloor \frac{k}{2} \rfloor$. Note that, if e intersects the boundary of $R(\mu)$ exactly once, then such an intersection does not count as an er-crossing, as in the traditional c-planarity literature.

Region-region crossings. A *region-region crossing* (*rr-crossing*) is a crossing involving two regions $R(\mu)$ and $R(\nu)$ representing clusters μ and ν of T, respectively, and such that μ is not an ancestor of ν and vice-versa. In fact, if μ is an ancestor of ν, then $R(\nu)$ is contained into $R(\mu)$ by the definition of drawing of a clustered graph. The number of rr-crossings between $R(\mu)$ and $R(\nu)$ is equal to the number of the topologically connected regions resulting from the relative complement of $R(\mu)$ in $R(\nu)$ (i.e., $R(\mu) \setminus R(\nu)$) minus one. Observe that, due to Constraints $\mathcal{C}_1$, $\mathcal{C}_2$, and $\mathcal{C}_3$, the number of rr-crossings between $R(\mu)$ and $R(\nu)$ is equal to the number of rr-crossings between $R(\nu)$ and $R(\mu)$. Also, as region $R(\mu)$ contains all and only the vertices of μ, intersections between regions cannot contain vertices of G. Figure 4 provides examples of region-region crossings.

Definition 1. *An $\langle \alpha, \beta, \gamma \rangle$-drawing of a clustered graph is a drawing with α ee-crossings, β er-crossings, and γ rr-crossings.*

2.1. Connectivity and SPQR-trees

A graph is *connected* if every two vertices are joined by a path. A graph G is *biconnected* (*triconnected*) if removing any vertex (any two vertices) leaves G connected.

To handle the decomposition of a biconnected graph into its triconnected components, we use *SPQR-trees*, a data structure introduced by Di Battista and Tamassia (see, e.g., [13, 14]).

A graph is *st-biconnectible* if adding edge (s, t) to it yields a biconnected graph. Let G be an st-biconnectible graph. A *separation pair* of G is a pair of vertices whose removal disconnects the graph. A *split pair* of G is either a separation pair or a pair of adjacent vertices. A *maximal split component* of G with respect to a split pair $\{u, v\}$ (or, simply, a maximal split component of $\{u, v\}$) is either an edge (u, v) or a maximal subgraph G' of G such that G' contains u and v, and $\{u, v\}$ is not a split pair of G'. A vertex $w \neq u, v$ belongs to exactly one maximal split component of $\{u, v\}$. We call *split component* of $\{u, v\}$ the union of any number of maximal split components of $\{u, v\}$.

We consider SPQR-trees that are rooted at one edge of the graph, called the *reference edge*.

The rooted SPQR-tree $\mathcal{T}$ of a biconnected graph G, with respect to a reference edge e, describes a recursive decomposition of G induced by its split pairs. The nodes of $\mathcal{T}$ are of four types: S, P, Q, and R. Their connections are called *arcs*, in order to distinguish them from the edges of G.

Each node τ of $\mathcal{T}$ has an associated st-biconnectible multigraph, called the *skeleton* of τ and denoted by $\mathrm{sk}(\tau)$. Skeleton $\mathrm{sk}(\tau)$ shows how the children of τ, represented by "virtual edges", are arranged in τ. The virtual edge in $\mathrm{sk}(\tau)$ associated with a child node σ, is called the *virtual edge of σ in sk(τ)*.

For each virtual edge e_i of $\mathrm{sk}(\tau)$, recursively replace e_i with the skeleton $\mathrm{sk}(\tau_i)$ of its corresponding child τ_i. The subgraph of G that is obtained in this way is the *pertinent graph* of τ and is denoted by $\mathrm{pert}(\tau)$.

Given a biconnected graph G and a reference edge $e = (u', v')$, tree $\mathcal{T}$ is recursively defined as follows. At each step, a split component G^*, a pair of vertices $\{u, v\}$, and a node σ in $\mathcal{T}$ are given. A node τ corresponding to G^* is introduced in $\mathcal{T}$ and attached to its parent σ. Vertices u and v are the *poles* of τ and denoted by $u(\tau)$ and $v(\tau)$, respectively. The decomposition possibly recurs on some split components of G^*. At the beginning of the decomposition $G^* = G - \{e\}$, $\{u, v\} = \{u', v'\}$, and σ is a Q-node corresponding to e.

Base Case: If G^* consists of exactly one edge between u and v, then τ is a Q-node whose skeleton is G^* itself.

Parallel Case: If G^* is composed of at least two maximal split components $G_1, \dots, G_k$ ($k \geq 2$) of G with respect to $\{u, v\}$, then τ is a P-node. Graph $\mathrm{sk}(\tau)$ consists of k parallel virtual edges between u and v, denoted by $e_1, \dots, e_k$ and corresponding to $G_1, \dots, G_k$, respectively. The decomposition recurs on $G_1, \dots, G_k$, with $\{u, v\}$ as pair of vertices for every graph, and with τ as parent node.

Series Case: If G^* is composed of exactly one maximal split component of G with respect to $\{u, v\}$ and if G^* has cutvertices $c_1, \dots, c_{k-1}$ ($k \geq 2$), appearing in this order on a path from u to v, then τ is an S-node. Graph $\mathrm{sk}(\tau)$ is the path $e_1, \dots, e_k$, where virtual edge e_i connects c_{i-1} with c_i ($i = 2, \dots, k-1$), e_1 connects u with c_1, and e_k connects c_{k-1} with v. The decomposition recurs on the split components corresponding to each of $e_1, e_2, \dots, e_{k-1}, e_k$ with τ as parent node, and with $\{u, c_1\}, \{c_1, c_2\}, \dots, \{c_{k-2}, c_{k-1}\}, \{c_{k-1}, v\}$ as pair of vertices, respectively.

Rigid Case: If none of the above cases applies, the purpose of the decomposition step is that of partitioning G^* into the minimum number of split components and recurring on each of them. We need some further definition. Given a maximal split component G' of a split pair $\{s,t\}$ of G^*, a vertex $w \in G'$ *properly belongs* to G' if $w \neq s,t$. Given a split pair $\{s,t\}$ of G^*, a maximal split component G' of $\{s,t\}$ is *internal* if neither u nor v (the poles of G^*) properly belongs to G', *external* otherwise. A *maximal split pair* $\{s,t\}$ of G^* is a split pair of G^* that is not contained into an internal maximal split component of any other split pair $\{s',t'\}$ of G^*. Let $\{u_1,v_1\},\ldots,\{u_k,v_k\}$ be the maximal split pairs of G^* ($k \geq 1$) and, for $i = 1,\ldots,k$, let G_i be the union of all the internal maximal split components of $\{u_i,v_i\}$. Observe that each vertex of G^* either properly belongs to exactly one G_i or belongs to some maximal split pair $\{u_i,v_i\}$. Node τ is an R-node. Graph $\mathrm{sk}(\tau)$ is the graph obtained from G^* by replacing each subgraph G_i with the virtual edge e_i between u_i and v_i. The decomposition recurs on each G_i with μ as parent node and with $\{u_i,v_i\}$ as pair of vertices.

For each node τ of $\mathcal{T}$, the construction of $\mathrm{sk}(\tau)$ is completed by adding a virtual edge (u,v) representing the rest of the graph.

In the reminder of the paper, even when not explicitly mentioned, we will always assume the considered SPQR-trees to be rooted at an edge of the graph.

The SPQR-tree $\mathcal{T}$ of a graph G with n vertices and m edges has m Q-nodes and $O(n)$ S-, P-, and R-nodes. Also, the total number of vertices of the skeletons stored at the nodes of $\mathcal{T}$ is $O(n)$. Finally, SPQR-trees can be constructed and handled efficiently. Namely, given a biconnected planar graph G, the SPQR-tree $\mathcal{T}$ of G can be computed in linear time [13, 14, 23].

3. Drawings of Clustered Graphs with Crossings

The following three sections deal with $\langle\infty,0,0\rangle$-, $\langle 0,\infty,0\rangle$- and $\langle 0,0,\infty\rangle$-drawings, respectively.

3.1. Drawings with Edge-Edge Crossings

In this section, we show a simple algorithm to construct an $\langle\alpha,0,0\rangle$-drawing of any clustered graph, in which α is asymptotically optimal in the worst case, as proved in Section 4.

Theorem 1. *Let $C(G,T)$ be a clustered graph. There exists an algorithm to compute an $\langle\alpha,0,0\rangle$-drawing of $C(G,T)$ with $\alpha \in O(n^2)$.*

Proof: Let $\sigma = v_1,\ldots,v_n$ be an ordering of the vertices of G such that vertices of the same cluster are consecutive in σ. A drawing of G can be constructed as follows. Place the vertices of G along a convex curve in the order they appear in σ. Draw the edges of G as straight-line segments. Since vertices belonging to the same cluster are consecutive in σ, drawing each cluster as the convex hull of the points assigned to its vertices yields a drawing without region-region and edge-region crossings (see Fig. 5). Further, since G has $O(n)$ edges, and since

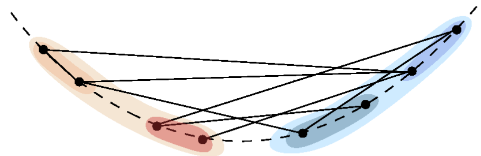

Figure 5: Illustration for Theorem 1.

edges are drawn as straight-line segments, such a construction produces $O(n^2)$ edge-edge crossings. □

Observe that, using the same construction used in the proof of Theorem 1, it can be proved that every clustered graph (even if its underlying graph is not planar) admits an $\langle \alpha, 0, 0 \rangle$-drawing with $\alpha \in O(n^4)$.

3.2. Drawings with Edge-Region Crossings

In this section, we show two algorithms for constructing a $\langle 0, \beta, 0 \rangle$-drawing of any clustered graph $C(G,T)$, in which β is asymptotically optimal in the worst case if $C(G,T)$ is c-connected or if it is flat, as proved in Section 4. If $C(G,T)$ is a general clustered graph, then β is a linear factor apart from the lower bound presented in Section 4.

The two algorithms handle the case in which $C(G,T)$ is not c-connected (Theorem 2) and in which $C(G,T)$ is c-connected (Theorem 3), respectively. Both algorithms have three steps:

1. A spanning tree $\mathcal{T}$ of the vertices of G is constructed in such a way that, for each cluster $\mu \in T$, the subgraph of $\mathcal{T}$ induced by the vertices of μ is connected. The two algorithms construct $\mathcal{T}$ in two different ways; in particular, $\mathcal{T}$ is a subgraph of G if $C(G,T)$ is c-connected, while it is not necessarily a subgraph of G if $C(G,T)$ is not c-connected.

2. A simultaneous embedding of G and $\mathcal{T}$ is computed. A *simultaneous embedding* of two graphs $G_1(V, E_1)$ and $G_2(V, E_2)$, on the same set V of vertices, is a drawing of $G(V, E_1 \cup E_2)$ such that any crossing involves an edge from E_1 and an edge from E_2 [4].

3. A $\langle 0, \beta, 0 \rangle$-drawing of $C(G,T)$ is constructed by drawing each cluster μ as a region $R(\mu)$ slightly surrounding the edges of $\mathcal{T}(\mu)$ and the regions $R(\mu_1), \dots, R(\mu_k)$ representing the children $\mu_1, \dots, \mu_k$ of μ.

In the case in which $C(G,T)$ is not c-connected, we get the following:

Theorem 2. *Let $C(G,T)$ be a clustered graph. Then, there exists an algorithm to compute a $\langle 0, \beta, 0 \rangle$-drawing of $C(G,T)$ with $\beta \in O(n^3)$. If $C(G,T)$ is flat, then $\beta \in O(n^2)$.*

Proof: In the first step, the tree $\mathcal{T}$ is constructed by means of a bottom-up traversal of T. Whenever a node $\mu \in T$ is considered, a spanning tree $\mathcal{T}(\mu)$ of $V(\mu)$ is constructed as follows. Denote by $\mu_1, \dots, \mu_k$ the children of μ in T (observe that, for each $1 \leq i \leq k$, μ_i is either a cluster or a vertex). Assume that spanning trees $\mathcal{T}(\mu_1), \dots, \mathcal{T}(\mu_k)$ of $V(\mu_1), \dots, V(\mu_k)$ have been already computed. The spanning tree $\mathcal{T}(\mu)$ of $V(\mu)$ is constructed by connecting a vertex of μ_1 to a vertex of each of $\mathcal{T}(\mu_2), \dots, \mathcal{T}(\mu_k)$. Tree $\mathcal{T}$ coincides with $\mathcal{T}(\rho)$, where ρ is the root of T. Observe that some of the edges of $\mathcal{T}$ might not belong to G.

In the second step, we apply the algorithm by Kammer [27] (see also [15]) to construct a simultaneous embedding of G and $\mathcal{T}$ in which each edge has at most two bends, which implies that each pair of edges $\langle e_1 \in G,\, e_2 \in \mathcal{T} \rangle$ crosses a constant number of times.

In the third step, each cluster μ is drawn as a region $R(\mu)$ slightly surrounding the edges of $\mathcal{T}(\mu)$ and the regions $R(\mu_1), \dots, R(\mu_k)$ representing the children $\mu_1, \dots, \mu_k$ of μ. Hence, each crossing between an edge $e_1 \in G$ and an edge $e_2 \in \mathcal{T}$ determines two intersections (hence one edge-region crossing) between e_1 and the boundary of each cluster ν such that $e_2 \in \mathcal{T}(\nu)$. Further, each edge $(u, v) \in G$ such that $(u, v) \notin \mathcal{T}$ and u and v belong to the same cluster ν, has a *er*-crossing with the boundary of $R(\nu)$.

Note that, for each edge $e_2 \in \mathcal{T}$, there exist $O(n)$ clusters ν such that $e_2 \in \mathcal{T}(\nu)$; also there exist $O(n^2)$ pairs of edges $\langle e_1 \in G, e_2 \in \mathcal{T} \rangle$; further, there exist $O(n)$ edges not belonging to τ; finally, for each edge $e \notin \mathcal{T}$ there exists $O(n)$ clusters ν such that both endvertices of e belong to ν. Hence, the total number of *er*-crossings is $O(n^3)$.

If $C(G, T)$ is flat, then for each edge $e_2 \in \mathcal{T}$ there exists at most one cluster ν different from the root such that $e_2 \in \mathcal{T}(\nu)$; also, for each edge $e \notin \mathcal{T}$ there exists at most one cluster ν different from the root such that both endvertices of e belong to ν. Hence, the total number of *er*-crossings is $O(n^2)$. □

If $C(G, T)$ is c-connected, we can improve the bounds of Theorem 2 as follows:

Theorem 3. *Let $C(G, T)$ be a c-connected clustered graph. Then, there exists an algorithm to compute a $\langle 0, \beta, 0 \rangle$-drawing of $C(G, T)$ with $\beta \in O(n^2)$. If $C(G, T)$ is flat, $\beta \in O(n)$.*

Proof: In the first step, the tree $\mathcal{T}$ is constructed by means of a bottom-up traversal of T. When a node $\mu \in T$ is considered, a spanning tree $\mathcal{T}(\mu)$ of $V(\mu)$ is constructed as follows. Denote by $\mu_1, \dots, \mu_k$ the children of μ in T (note that, for each $1 \leq i \leq k$, μ_i is either a cluster or a vertex). Assume that spanning trees $\mathcal{T}(\mu_1), \dots, \mathcal{T}(\mu_k)$ of $V(\mu_1), \dots, V(\mu_k)$ have been already computed so that $\mathcal{T}(\mu_i)$ is a subgraph of $G(\mu_i)$, for $i = 1, \dots, k$. Tree $\mathcal{T}(\mu)$ contains all the edges in $\mathcal{T}(\mu_1), \dots, \mathcal{T}(\mu_k)$ plus a minimal set of edges of $G(\mu)$ connecting $\mathcal{T}(\mu_1), \dots, \mathcal{T}(\mu_k)$. The latter set of edges always exists since $G(\mu)$ is connected. Tree $\mathcal{T}$ coincides with $\mathcal{T}(\rho)$, where ρ is the root of T. Observe that, in contrast with the construction in the proof of Theorem 2, all edges of $\mathcal{T}$ belong to G.

In the second step, since each edge of $\mathcal{T}$ is also an edge of G, any planar drawing of G determines a simultaneous embedding of G and $\mathcal{T}$ in which no edge of G properly crosses an edge of $\mathcal{T}$.

In the third step, clusters are drawn in the same way as in the proof of Theorem 2.

Note that the only edge-region crossings that may occur are those between any edge of G not in $\mathcal{T}$ whose endvertices belong to the same cluster μ and the boundary of $R(\mu)$. Since there exist $O(n)$ edges not belonging to $\mathcal{T}$ and since for each edge $e \notin \mathcal{T}$ there exist $O(n)$ clusters ν such that both endvertices of e belong to ν, it follows that the total number of edge-region crossings is $O(n^2)$.

If $C(G,T)$ is flat, then for each edge $e \notin \mathcal{T}$ there exists at most one cluster ν different from the root such that both endvertices of e belong to ν, and hence the total number of edge-region crossings is $O(n)$. $\square$

3.3. Drawings with Region-Region Crossings

In this section, we study $\langle 0,0,\infty\rangle$-drawings of clustered graphs. First, we prove that there are clustered graphs that do not admit any $\langle 0,0,\infty\rangle$-drawings. Second, we provide a polynomial-time algorithm for testing whether a clustered graph $C(G,T)$ with G biconnected admits a $\langle 0,0,\infty\rangle$-drawing and to compute one if it exists. Third, we show an algorithm that constructs a $\langle 0,0,\gamma\rangle$-drawing Γ of any clustered graph $C(G,T)$ that admits such a drawing (the input of the algorithm is any $\langle 0,0,\infty\rangle$-drawing Γ' of $C(G,T)$) in which γ is worst-case asymptotically optimal.

To show that there exist clustered graphs not admitting any $\langle 0,0,\infty\rangle$-drawing, we give two examples. Let $C(G,T)$ be a clustered graph such that G is triconnected and has a cycle of vertices belonging to a cluster μ separating two vertices not in μ (see Fig. 6(a)). Note that, even in the presence of rr-crossings, one of the two vertices not in μ is enclosed by $R(\mu)$ in any $\langle 0,0,\infty\rangle$-drawing of $C(G,T)$. This example exploits the triconnectivity of the underlying graph. Next we show that even clustered graphs with series-parallel underlying graph may not admit any $\langle 0,0,\infty\rangle$-drawing. Namely, let $C(G,T)$ be a clustered graph such that G has eight vertices and is composed of parallel paths p_1, p_2, p_3, and p_4. Tree T is such that cluster μ_1 contains a vertex of p_1 and a vertex of p_2; cluster μ_2 contains a vertex of p_2 and a vertex of p_3; cluster μ_3 contains a vertex of p_2 and a vertex of p_4 (see Fig. 6(b)). Note that, in any $\langle 0,0,\infty\rangle$-drawing of $C(G,T)$, path p_2 should be adjacent to all the other paths in the order around the poles, and this is not possible by the planarity of the drawing of G.

Since some clustered graphs do not admit any $\langle 0,0,\infty\rangle$-drawing, we study the complexity of testing whether a clustered graph $C(G,T)$ admits one. In order to do that, we first give a characterization of the planar embeddings of G that allow for the realization of a $\langle 0,0,\infty\rangle$-drawing of $C(G,T)$. Namely, let $C(G,T)$ be a clustered graph and let Γ be a planar embedding of G. For each cluster $\mu \in T$ consider an auxiliary graph $H(\mu)$ with the same vertices as $G(\mu)$

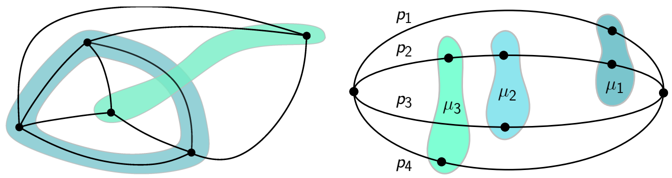

(a) (b)

Figure 6: Two clustered graphs not admitting any $\langle 0, 0, \infty \rangle$-drawing. The underlying graph of (a) is a triconnected planar graph, while the underlying graph of (b) is a series-parallel graph.

and such that there is an edge between two vertices of $H(\mu)$ if and only if the corresponding vertices of G are incident to the same face in Γ.

Lemma 1. *Let $C(G,T)$ be a clustered graph and let Γ be a planar embedding of G. Then, $C(G,T)$ admits a $\langle 0, 0, \infty \rangle$-drawing preserving Γ if and only if, for each cluster $\mu \in T$: (i) graph $H(\mu)$ is connected and (ii) there exists no cycle of G whose vertices belong to μ and whose interior contains in Γ a vertex not belonging to μ.*

Proof: We first prove the necessity of the conditions. For the necessity of Condition (i), suppose that $H(\mu)$ is not connected. Then, for any two distinct connected components $H_1(\mu)$ and $H_2(\mu)$ of $H(\mu)$, there exists a cycle $\mathcal{C}$ in G separating $H_1(\mu)$ and $H_2(\mu)$, as otherwise $H_1(\mu)$ and $H_2(\mu)$ would be incident to a common face, hence they would not be distinct connected components of $H(\mu)$. Therefore, the boundary of any region $R(\mu)$ representing μ intersects (at least) one of the edges of $\mathcal{C}$. For the necessity of Condition (ii), suppose that a cycle $\mathcal{C}$ exists in Γ whose vertices belong to μ and whose interior contains in Γ a vertex not belonging to μ. Then, in any drawing of $R(\mu)$ as a simple closed region containing all and only the vertices in μ, the border of $R(\mu)$ intersects (at least) one edge of $\mathcal{C}$.

We next prove the sufficiency of the conditions. Suppose that Conditions (i) and (ii) hold. Consider any subgraph $H'(\mu)$ of $H(\mu)$ such that: (a) $G(\mu) \subseteq H'(\mu)$; (b) $H'(\mu)$ is connected; and (c) for every cycle $\mathcal{C}$ in $H'(\mu)$, if any, all the edges of $\mathcal{C}$ belong to G. Observe that the fact that $H(\mu)$ satisfies conditions (i) and (ii) implies the existence of a graph $H'(\mu)$ satisfying (a), (b), and (c). Draw each edge of $H'(\mu)$ not in G inside the corresponding face. Represent μ as a region slightly surrounding the (possibly non-simple) cycle delimiting the outer face of $H'(\mu)$. Denote by Γ'_C the resulting drawing and denote by Γ_C the drawing of $C(G,T)$ obtained from Γ'_C by removing the edges not in G. We have that Γ_C contains no *ee*-crossing, since Γ is a planar embedding. Also, it contains no *er*-crossing, since the only edges crossing clusters in Γ'_C are those belonging to $H'(\mu)$ and not belonging to $G(\mu)$. □

Our next goal is to provide an algorithm that, given a clustered graph $C(G,T)$ such that G is biconnected, tests whether G admits a planar embedding allowing for a $\langle 0, 0, \infty \rangle$-drawing of $C(G,T)$.

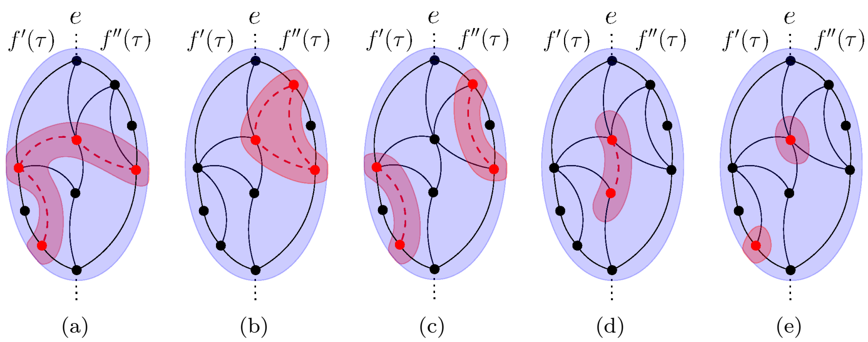

Figure 7: Examples of embeddings of $\Gamma(pert(\tau))$: (a) μ-traversable; (b) μ-sided; (c) μ-bisided; (d) μ-kernelized; (e) μ-infeasible. Dashed red edges belong to $H(\tau,\mu)$. The five drawings represent five embeddings $\Gamma(pert(\tau))$ for five different graphs $pert(\tau)$.

We start by giving some definitions. Let $C(G,T)$ be a clustered graph such that G is biconnected and consider the SPQR-tree $\mathcal{T}$ of G rooted at any Q-node ρ. The choice of rooting $\mathcal{T}$ at ρ corresponds to only consider planar embeddings of G in which the edge e_ρ of G corresponding to ρ is incident to the outer face.

Consider a node $\tau \in \mathcal{T}$, its pertinent graph $pert(\tau)$ (augmented with an edge e between the poles of τ, representing the parent of τ), and a planar embedding $\Gamma(pert(\tau))$ with e on the outer face. Observe that assuming e to be incident to the outer face of $\Gamma(pert(\tau))$ is not a loss of generality, given that e_ρ is assumed to be incident to the outer face of every considered planar embedding of G. Namely, consider any planar embedding Γ_G of G in which e_ρ is incident to the outer face and let $\Gamma^-(pert(\tau))$ be the embedding of $pert(\tau)$ (except for edge e) obtained by restricting Γ_G to $pert(\tau)$. Then, the subgraph of G not in $pert(\tau)$ (i.e., the "rest of the graph" with respect to τ) lies in Γ_G in the outer face of $\Gamma^-(pert(\tau))$, hence edge e can be inserted in the outer face of $\Gamma^-(pert(\tau))$, thus obtaining a planar embedding $\Gamma(pert(\tau))$ with e on the outer face.

Let $f'(\tau)$ and $f''(\tau)$ be the two faces of $\Gamma(pert(\tau))$ that are incident to e. For each cluster $\mu \in T$, we define an auxiliary graph $H(\tau,\mu)$ as the graph containing all the vertices of $pert(\tau)$ that belong to μ and such that two vertices of $H(\tau,\mu)$ are connected by an edge if and only if they are incident to the same face in $\Gamma(pert(\tau))$. Observe that $H(\rho,\mu)$ coincides with the above defined auxiliary graph $H(\mu)$. Also, observe that no two connected components of $H(\tau,\mu)$ exist both containing a vertex incident to $f'(\tau)$ or both containing a vertex incident to $f''(\tau)$.

Next we introduce some classifications of the nodes of $\mathcal{T}$ and of the embeddings of their pertinent graphs that will be used to find an embedding of G such that Conditions (i) and (ii) of Lemma 1 are satisfied for each cluster μ.

In order to keep track of the connectivity of $H(\mu)$ (Condition (i) of Lemma 1), for each node $\tau \in \mathcal{T}$ and for each cluster $\mu \in T$, we say that $\Gamma(pert(\tau))$ is:

***μ-traversable*:** if $H(\tau,\mu)$ is connected and contains at least one vertex incident to $f'(\tau)$ and one vertex incident to $f''(\tau)$ (see Fig. 8(a)).

***μ-sided*:** if $H(\tau, \mu)$ is connected and contains at least one vertex incident to $f'(\tau)$ and no vertex incident to $f''(\tau)$, or vice versa (see Fig. 8(b)).

***μ-bisided*:** if $H(\tau, \mu)$ consists of two connected components, one containing a vertex of $f'(\tau)$ and the other one containing a vertex of $f''(\tau)$ (see Fig. 8(c)).

***μ-kernelized*:** if $H(\tau, \mu)$ is connected and contains neither a vertex incident to $f'(\tau)$ nor a vertex incident to $f''(\tau)$ (see Fig. 8(d)).

***μ-infeasible*:** if $H(\tau, \mu)$ has at least two connected components of which one has no vertex incident to $f'(\tau)$ or $f''(\tau)$ (see Fig. 8(e)).

Note that, if τ contains at least one vertex of μ, then $\Gamma(pert(\tau))$ is exactly of one of the types of embedding defined above.

First, we derive an elementary condition that prevents an embedding from being μ-traversable.

Lemma 2. *Suppose that an embedding $\Gamma(pert(\tau))$ of $pert(\tau)$ with e incident to the outer face is neither μ-traversable nor μ-infeasible. Then, there exists a path in $pert(\tau)$ that connects the poles of τ, that is different from e, and none of whose vertices belongs to μ.*

Proof: Suppose first that $H(\tau, \mu)$ is connected in $\Gamma(pert(\tau))$. Then, $\Gamma(pert(\tau))$ is either μ-traversable, or μ-kernelized, or μ-sided, by definition. By assumption, $\Gamma(pert(\tau))$ is not μ-traversable, hence it is either μ-kernelized or μ-sided. In both cases, the path of $\Gamma(pert(\tau))$ delimiting $f'(\tau)$, different from e, and connecting the poles of τ, or the path of $\Gamma(pert(\tau))$ delimiting $f''(\tau)$, different from e, and connecting the poles of τ is such that none of its vertices belongs to μ, and the statement follows.

Suppose next that $H(\tau, \mu)$ is not connected in $\Gamma(pert(\tau))$. Then, $\Gamma(pert(\tau))$ is either μ-infeasible or μ-bisided, by definition. By assumption, $\Gamma(pert(\tau))$ is not μ-infeasible, hence it is μ-bisided. By definition of $H(\tau, \mu)$, the two connected components $H'(\tau, \mu)$ and $H''(\tau, \mu)$ of $H(\tau, \mu)$ are not incident to any common face. Thus, there exists a cycle $\mathcal{C}$ in $\Gamma(pert(\tau))$ containing such components on different sides and none of whose vertices belongs to μ. Cycle $\mathcal{C}$ contains edge e, given that $H'(\tau, \mu)$ and $H''(\tau, \mu)$ contain vertices incident to $f'(\tau)$ and $f''(\tau)$. It follows that the path obtained from $\mathcal{C}$ by removing edge e connects the poles of τ, is different from edge e, and is such that none of its vertices belongs to μ, thus proving the statement. $\square$

While the fact that $\Gamma(pert(\tau))$ is μ-sided, μ-bisided, μ-kernelized, or μ-infeasible does not rule the possibility that a different planar embedding of $pert(\tau)$ is of a different type, if $\Gamma(pert(\tau))$ is μ-traversable then any other embedding of $pert(\tau)$ is either μ-traversable or μ-infeasible, as proved in the following.

Lemma 3. *Let $\Gamma_1(pert(\tau))$ and $\Gamma_2(pert(\tau))$ be two planar embeddings of $pert(\tau)$, both having edge e incident to the outer face. If $\Gamma_1(pert(\tau))$ is μ-traversable, then $\Gamma_2(pert(\tau))$ is either μ-traversable or μ-infeasible.*

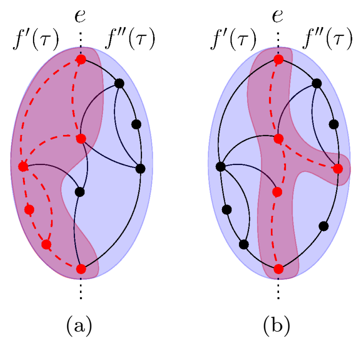

Figure 8: Examples of embeddings of $\Gamma(pert(\tau))$: (a) μ-side-spined; (b) μ-central-spined. Dashed red edges belong to $H(\tau, \mu)$. The two drawings represent two embeddings $\Gamma(pert(\tau))$ for two different graphs $pert(\tau)$.

Proof: Suppose, for a contradiction, that $\Gamma_2(pert(\tau))$ is neither μ-traversable nor μ-infeasible. By Lemma 2, there exists a path in $pert(\tau)$ that connects the poles of τ, that is different from e, and none of whose vertices belongs to μ. It follows that the auxiliary graph $H(\tau, \mu)$ associated with $\Gamma_2(pert(\tau))$ is disconnected, or does not contain a vertex incident to $f'(\tau)$, or does not contain a vertex incident to $f''(\tau)$. □

By Lemma 3, if an embedding of $pert(\tau)$ is μ-traversable, then every embedding of $pert(\tau)$ which is not μ-infeasible is μ-traversable. Hence, if a μ-traversable embedding of $pert(\tau)$ exists, we say that τ and the virtual edge representing τ in the skeleton of the parent of τ in $\mathcal{T}$ is μ-traversable.

Next, we introduce some definitions used to deal with Condition (*ii*) of Lemma 1. Namely, for each node $\tau \in \mathcal{T}$ and for each cluster $\mu \in T$, we say that τ (and the virtual edge representing τ in the skeleton of its parent) is:

***μ-touched*:** if there exists a vertex in $pert(\tau) \setminus \{u, v\}$ that belongs to μ.

***μ-full*:** if all the vertices in $pert(\tau)$ belong to μ.

***μ-spined*:** if there exists in $pert(\tau)$ a path P between the poles of τ different from e and containing only vertices of μ. Observe that, if τ is μ-spined and $pert(\tau)$ is not a single edge, then τ is μ-touched.

Given a μ-spined node τ, an embedding $\Gamma(pert(\tau))$ is **μ-side-spined** if at least one of the two paths different from edge e, connecting the poles of τ, and delimiting the outer face of $\Gamma(pert(\tau))$ has only vertices in μ (see Fig 9(a)). Otherwise, it is **μ-central-spined** (see Fig 9(b)).

Observe that, if τ is μ-spined, then it is also μ-traversable, since its poles belong to μ.

We say that an embedding $\Gamma(pert(\tau))$ of $pert(\tau)$ with e incident to the outer face is *extensible* if the following condition holds: If $C(G, T)$ admits a $\langle 0, 0, \infty \rangle$-drawing in which the edge e_ρ corresponding to ρ is incident to the outer face,

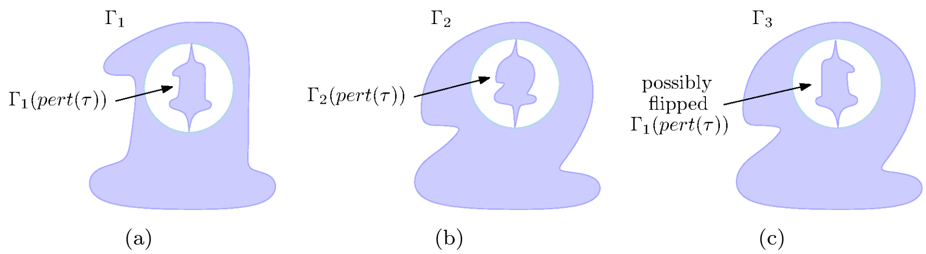

Figure 9: Illustration for the statement of Lemma 4. (a) Embedding Γ_1. (b) Embedding Γ_2. (c) Embedding Γ_3.

then it admits a $\langle 0, 0, \infty \rangle$-drawing in which e_ρ is incident to the outer face and in which the embedding of $pert(\tau)$ is $\Gamma(pert(\tau))$. Observe that, if an embedding $\Gamma(pert(\tau))$ of $pert(\tau)$ is μ-infeasible, for some $\mu \in T$, then $\Gamma(pert(\tau))$ is not extensible.

One key ingredient of our result is that if an embedding $\Gamma(pert(\tau))$ of $pert(\tau)$ is extensible, not only there exists a $\langle 0, 0, \infty \rangle$-drawing of $C(G,T)$ in which the embedding of $pert(\tau)$ is $\Gamma(pert(\tau))$, but for *every* $\langle 0, 0, \infty \rangle$-drawing Γ of $C(G,T)$ in which e_ρ is incident to the outer face, the embedding of $pert(\tau)$ in Γ can be modified to be $\Gamma(pert(\tau))$, without changing the rest of Γ, while maintaining the property that Γ is a $\langle 0, 0, \infty \rangle$-drawing. In the following we formalize such a claim.

Lemma 4. *Let $C(G,T)$ be a clustered graph, with G biconnected, that admits a $\langle 0, 0, \infty \rangle$-drawing in which the edge e_ρ corresponding to the root ρ of SPQR-tree $\mathcal{T}$ of G is incident to the outer face. Let Γ_1 and Γ_2 be two $\langle 0, 0, \infty \rangle$-drawings of $C(G,T)$ in which e_ρ is incident to the outer face. Let τ be a P-node or an R-node of $\mathcal{T}$. Let $\Gamma_i(pert(\tau))$ be the embedding of $pert(\tau)$ minus edge e in Γ_i, for $i = 1, 2$. Let Γ_3 be the drawing obtained from Γ_2 by replacing $\Gamma_2(pert(\tau))$ with $\Gamma_1(pert(\tau))$, possibly after performing a flip of $\Gamma_1(pert(\tau))$ (see Fig. 10(c)). Then, Γ_3 is a $\langle 0, 0, \infty \rangle$-drawing of $C(G,T)$ in which e_ρ is incident to the outer face.*

In the following we prove Lemma 4. Namely, we prove that, after the replacement of $\Gamma_2(pert(\tau))$ with $\Gamma_1(pert(\tau))$ and, possibly, a flip of $\Gamma_1(pert(\tau))$, the resulting embedding Γ_3 of G satisfies the conditions of Lemma 1. Observe that e_ρ is incident to the outer face of Γ_3, given that it is incident to the outer face of Γ_2.

We introduce some terminology. The *rest of the graph with respect to* τ is the graph $G(\overline{\tau})$ obtained from G by removing the vertices of $pert(\tau)$ different from its poles and by inserting a dummy edge e_τ between the poles of $pert(\tau)$. In terms of SPQR-trees, the rest of the graph can be equivalently defined as follows. Denote by τ' the parent of τ in $\mathcal{T}$. Then, the rest of the graph is the pertinent graph of τ' in any re-rooting of $\mathcal{T}$ in which τ' becomes a child of τ plus an edge e_τ between the poles of τ. Denote by $f'(\tau)$ and $f''(\tau)$ the faces of

$G(\overline{\tau})$ incident to e_τ. Recall that $f'(\tau)$ and $f''(\tau)$ also denote the faces incident to e in the embeddings $\Gamma_1(pert(\tau))$ and $\Gamma_2(pert(\tau))$ of $pert(\tau)$ in which edge e is added in the outer face. We further overload the notation $f'(\tau)$ and $f''(\tau)$ to let them represent the faces of Γ_i that are shared by $\Gamma_i(pert(\tau))$ and $\Gamma_i(G(\overline{\tau}))$, for $i = 1, 2, 3$. For any node μ of T, the *auxiliary graph* $H(\overline{\tau}, \mu)$ *of* $G(\overline{\tau})$ in Γ_i, for any $i = 1, 2, 3$, is the graph containing all the vertices of $G(\overline{\tau})$ that belong to μ and such that two vertices of $H(\overline{\tau}, \mu)$ are connected by an edge if and only if they are incident to the same face in Γ_i. Let $\Gamma_i(G(\overline{\tau}))$ be the embedding of $G(\overline{\tau})$ in Γ_i, for $i = 1, 2, 3$.

The definitions of μ-traversable, μ-sided, μ-bisided, μ-kernelized, μ-infeasible, μ-touched, μ-full, and μ-spined apply to the rest of the graph with respect to τ analogously as to the pertinent graph of τ. For example, $G(\overline{\tau})$ is *μ-traversable* in Γ_i if $H(\overline{\tau}, \mu)$ is connected and contains at least one vertex incident to $f'(\tau)$ and one vertex incident to $f''(\tau)$.

For any $i = 1, 2$, if Γ_i is such that $\Gamma_i(pert(\tau))$ is μ-sided or μ-side-spined, then, we denote by $p(\Gamma_i, \tau, \mu)$ the path that (i) connects the poles of τ, (ii) belongs to $pert(\tau)$, (iii) delimits $f'(\tau)$ or $f''(\tau)$ in Γ_i, and (iv) contains vertices of μ (if $\Gamma_i(pert(\tau))$ is μ-sided, see Fig. 11(a)), or entirely belongs to μ (if $\Gamma_i(pert(\tau))$ is μ-side-spined, see Fig. 11(b)).

Similarly, for any $i = 1, 2$, suppose Γ_i is such that $\Gamma_i(G(\overline{\tau}))$ is μ-sided, or is μ-side-spined, or is μ-full and both the poles of τ belong to the outer face of Γ_i (refer to Figs. 11(c), 11(d), and 11(e) for examples of the three cases, respectively). Then, we denote by $p(\Gamma_i, \overline{\tau}, \mu)$ the path that (i) connects the poles of τ, (ii) belongs to $G(\overline{\tau})$, (iii) delimits $f'(\tau)$ or $f''(\tau)$ in Γ_i, and (iv) contains vertices of μ (if $\Gamma_i(G(\overline{\tau}))$ is μ-sided, see Fig. 11(c)), or entirely belongs to μ (if $\Gamma_i(G(\overline{\tau}))$ is μ-side-spined, see Fig. 11(d)), or is not entirely incident to the outer face of Γ_i (if $\Gamma_i(G(\overline{\tau}))$ is μ-full and both the poles of τ belong to the outer face of Γ_i, see Fig. 11(e)).

We show a simple algorithm to determine a flip of $\Gamma_1(pert(\tau))$; observe that the choice of such a flip completely determines Γ_3, as the embedding of $G(\overline{\tau})$ in Γ_3 coincides with $\Gamma_2(G(\overline{\tau}))$. Consider any cluster $\mu \in T$ such that one of the following holds:

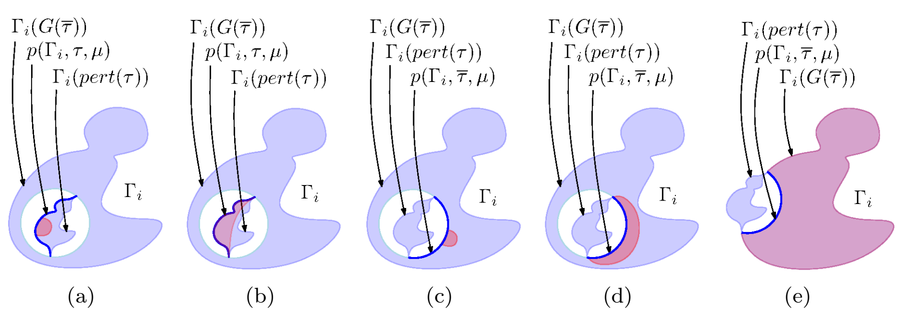

Figure 10: Illustration for the definition of paths $p(\Gamma_i, \tau, \mu)$ and $p(\Gamma_i, \overline{\tau}, \mu)$ of Γ_i, with $i = 1, 2$.

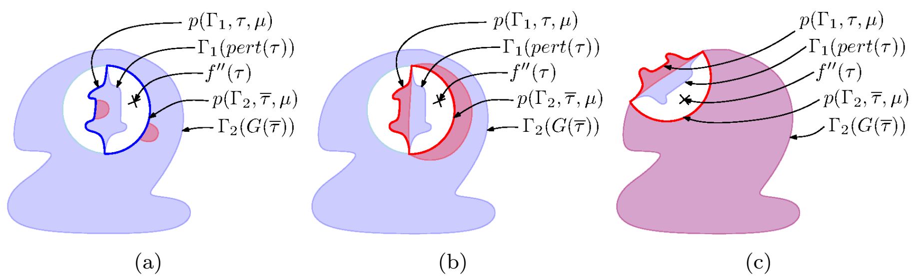

Figure 11: The three cases of the proof of Lemma 4. (a) Case 1. (b) Case 2. (c) Case 3.

Case 1: $\Gamma_1(pert(\tau))$ and $\Gamma_2(G(\overline{\tau}))$ are both μ-sided (see Fig. 12(a)).

Case 2: $\Gamma_1(pert(\tau))$ and $\Gamma_2(G(\overline{\tau}))$ are both μ-side-spined and not μ-full (see Fig. 12(b)).

Case 3: $G(\overline{\tau})$ is μ-full, $\Gamma_1(pert(\tau))$ is μ-side-spined and not μ-full, and both the poles of τ belong to the outer face of Γ_2 (see Fig. 12(c)).

Then, flip $\Gamma_1(pert(\tau))$ so that $p(\Gamma_1, \tau, \mu)$ and $p(\Gamma_2, \overline{\tau}, \mu)$ delimit a face in Γ_3.

If no cluster $\mu \in T$ exists that determines the flip of $\Gamma_1(pert(\tau))$, that is, if for all clusters of T none of the above cases applies, then arbitrarily choose a flip for $\Gamma_1(pert(\tau))$. We will prove later that no two clusters $\mu \neq \nu \in T$ exist that determine different flips for $\Gamma_1(pert(\tau))$.

We now proceed with the proof of Lemma 4, showing that the embedding Γ_3 of G satisfies the conditions of Lemma 1.

Claim 1. *Embedding Γ_3 satisfies Condition (i) of Lemma 1.*

Proof: Assume, for a contradiction, that Condition (i) of Lemma 1 is not satisfied by Γ_3. Consider any node $\mu \in T$ such that $H(\mu)$ is not connected in Γ_3.

If $pert(\tau)$ or $G(\overline{\tau})$ is not μ-touched, then the connectivity of $H(\mu)$ in Γ_3 descends from the connectivity of $H(\mu)$ in Γ_2 or in Γ_1, respectively, a contradiction.

If $\Gamma_1(pert(\tau))$ is μ-infeasible, then Γ_1 is not a $\langle 0, 0, \infty\rangle$-drawing, a contradiction. Analogously, if $\Gamma_2(G(\overline{\tau}))$ is μ-infeasible, then Γ_2 is not a $\langle 0, 0, \infty\rangle$-drawing, a contradiction.

If $\Gamma_1(pert(\tau))$ is μ-kernelized, then, since $H(\mu)$ is connected in Γ_1, it follows that $G(\overline{\tau})$ is not μ-touched, a case that has been already addressed. Analogous considerations apply if $\Gamma_2(G(\overline{\tau}))$ is μ-kernelized.

If $\Gamma_1(pert(\tau))$ is μ-traversable, then by Lemma 3 we have that $\Gamma_2(pert(\tau))$ is either μ-infeasible or μ-traversable. In the former case, we have that Γ_2 is not a $\langle 0, 0, \infty\rangle$-drawing, a contradiction. In the latter case, if $H(\mu)$ is not connected in Γ_3, then $\Gamma_2(G(\overline{\tau}))$ is μ-kernelized or μ-infeasible, hence $H(\mu)$ is not connected in Γ_2, a contradiction. Analogous considerations apply if $\Gamma_2(G(\overline{\tau}))$ is μ-traversable.

If $\Gamma_1(pert(\tau))$ is μ-bisided, then $\Gamma_1(G(\overline{\tau}))$ is μ-traversable. Then, by Lemma 3 we have that either $\Gamma_2(G(\overline{\tau}))$ is either μ-infeasible or μ-traversable. In both cases, we have already shown how to derive a contradiction. Analogous considerations apply if $\Gamma_2(G(\overline{\tau}))$ is μ-bisided.

Hence, the only case that remains to be considered is the one in which both $\Gamma_1(pert(\tau))$ and $\Gamma_2(G(\overline{\tau}))$ are μ-sided. Observe that Case 1 for the determination of the flip of $\Gamma_1(pert(\tau))$ applies, hence $\Gamma_1(pert(\tau))$ is flipped in such a way that the connected component of $H(\mu)$ induced by the vertices in $\Gamma_1(pert(\tau))$ and the connected component of $H(\mu)$ induced by the vertices in $\Gamma_2(G(\overline{\tau}))$ both contain vertices incident to either $f'(\tau)$ or $f''(\tau)$, hence they are connected in Γ_3. It follows that $H(\mu)$ is connected in Γ_3, a contradiction. □

Claim 2. *Embedding Γ_3 satisfies Condition (ii) of Lemma 1.*

Proof: Assume for a contradiction that Condition (ii) of Lemma 1 is not satisfied by Γ_3. Then, for some cluster $\mu \in T$, we have that Γ_3 contains a cycle $\mathcal{C}$ whose vertices belong to μ and whose interior contains in Γ_3 a vertex not belonging to μ.

Suppose first that all the edges of $\mathcal{C}$ belong to $G(\overline{\tau})$. Since the embedding $\Gamma_3(G(\overline{\tau}))$ of $G(\overline{\tau})$ in Γ_3 coincides with the embedding $\Gamma_2(G(\overline{\tau}))$ of $G(\overline{\tau})$ in Γ_2, it follows that $\mathcal{C}$ contains a vertex not belonging to μ in its interior in Γ_2. This contradicts Lemma 1.

Suppose next that all the edges of $\mathcal{C}$ belong to $pert(\tau)$. Since the embedding $\Gamma_3(pert(\tau))$ of $pert(\tau)$ in Γ_3 coincides with the embedding $\Gamma_1(pert(\tau))$ of $pert(\tau)$ in Γ_1, up to a flip, it follows that $\mathcal{C}$ contains a vertex not belonging to μ in its interior in Γ_1. This contradicts Lemma 1.

We can hence assume that $\mathcal{C}$ contains both edges of $pert(\tau)$ and edges of $G(\overline{\tau})$. That is, $\mathcal{C}$ is composed of a path $q(\tau,\mu)$ in $pert(\tau)$ connecting the poles of τ and of a path $q(\overline{\tau},\mu)$ in $G(\overline{\tau}) - e_\tau$ connecting the poles of τ. Thus, each of $\Gamma_1(pert(\tau))$ and $\Gamma_2(G(\overline{\tau}))$ is either μ-full, or μ-central-spined, or μ-side-spined and not μ-full.

If $\Gamma_1(pert(\tau))$ is μ-central-spined, then there exist two vertices x and y that belong to $pert(\tau)$, that do not belong to μ, and that lie on different sides of $\mathcal{C}$ in Γ_1. This contradicts Lemma 1. Analogously, if $\Gamma_2(G(\overline{\tau}))$ is μ-central-spined, then there exist two vertices x and y that belong to $G(\overline{\tau})$, that do not belong to μ, and that lie on different sides of $\mathcal{C}$ in Γ_2. This contradicts Lemma 1.

If $G(\overline{\tau})$ is μ-full and $pert(\tau)$ is μ-full, then $\mathcal{C}$ does not contain any vertex not belonging to μ in its interior, a contradiction.

If $pert(\tau)$ is μ-full and $\Gamma_2(G(\overline{\tau}))$ is μ-side-spined and not μ-full, then any vertex not belonging to μ belongs to $G(\overline{\tau})$. Since the embedding $\Gamma_3(G(\overline{\tau}))$ of $G(\overline{\tau})$ in Γ_3 coincides with the embedding $\Gamma_2(G(\overline{\tau}))$ of $G(\overline{\tau})$ in Γ_2, it follows that $\mathcal{C}$ contains a vertex not belonging to μ in its interior in Γ_2. This contradicts Lemma 1.

If $\Gamma_1(pert(\tau))$ and $\Gamma_2(G(\overline{\tau}))$ are both μ-side-spined and not μ-full, or if $G(\overline{\tau})$ is μ-full, $\Gamma_1(pert(\tau))$ is μ-side-spined and not μ-full, and both the poles of τ belong to the outer face of Γ_2, then Case 2 or Case 3 for the determination of the flip of $\Gamma_1(pert(\tau))$ applies, respectively, hence $\Gamma_1(pert(\tau))$ is flipped in such a way that $p(\Gamma_1, \tau, \mu)$ and $p(\Gamma_2, \overline{\tau}, \mu)$ are both incident to either $f'(\tau)$ or $f''(\tau)$ in Γ_3. Since Γ_1 is a $\langle 0, 0, \infty \rangle$-drawing of $C(G,T)$, by Lemma 1 the cycle composed of paths $p(\Gamma_1, \tau, \mu)$ and $q(\tau, \mu)$ does not contain any vertex not belonging to μ in its interior. Hence, $\mathcal{C}$ contains a vertex not belonging to μ in its interior if and only if the cycle $\mathcal{C}'$ composed of $q(\overline{\tau}, \mu)$ and $p(\Gamma_1, \tau, \mu)$ contains a vertex not belonging to μ in its interior. Also, since $p(\Gamma_1, \tau, \mu)$ and $p(\Gamma_2, \overline{\tau}, \mu)$ both delimit a face of Γ_3, we have that $\mathcal{C}'$ contains a vertex not belonging to μ in its interior if and only if the cycle $\mathcal{C}''$ composed of $q(\overline{\tau}, \mu)$ and $p(\Gamma_2, \overline{\tau}, \mu)$ contains a vertex not belonging to μ in its interior in Γ_3. Observe that $\mathcal{C}''$ is a cycle in $G(\overline{\tau})$. Hence, if $\mathcal{C}''$ contains a vertex not belonging to μ in its interior in Γ_3, then it contains a vertex not belonging to μ in its interior in Γ_2, given that the embedding $\Gamma_3(G(\overline{\tau}))$ of $G(\overline{\tau})$ in Γ_3 coincides with the embedding $\Gamma_2(G(\overline{\tau}))$ of $G(\overline{\tau})$ in Γ_2. This contradicts Lemma 1.

Finally, if $G(\overline{\tau})$ is μ-full, $pert(\tau)$ is not μ-full, and at least one of the poles of τ does not belong to the outer face of Γ_2, then there exists a cycle $\mathcal{C}'$ in $G(\overline{\tau})$ that contains all the vertices of $pert(\tau)$, except possibly for its poles, in its interior. All the vertices of $\mathcal{C}'$ belong to μ; further, at least one vertex of $pert(\tau)$ does not belong to μ. This contradicts Lemma 1. □

In order to conclude the proof of Lemma 4, it remains to prove that no two clusters $\mu \neq \nu \in T$ exist that determine different flips for $\Gamma_1(pert(\tau))$.

Claim 3. *No two clusters $\mu \neq \nu \in T$ exist that determine different flips for $\Gamma_1(pert(\tau))$ when constructing Γ_3.*

Proof: Assume, for a contradiction, that two distinct clusters μ and ν determine different flips for $\Gamma_1(pert(\tau))$.

- Suppose that Case 2 or Case 3 applies to μ and that Case 2 or 3 applies to ν to determine a different flip for $\Gamma_1(pert(\tau))$. Since the poles of τ are both contained in μ and in ν, it follows that μ is an ancestor of ν or ν is an ancestor of μ. Assume the former, the discussion for the latter case being analogous.
  - Assume first that $p(\Gamma_1, \tau, \mu)$ and $p(\Gamma_1, \tau, \nu)$ are distinct. Since μ is an ancestor of ν, either all vertices of $pert(\tau)$ belong to μ, hence τ is μ-full, thus contradicting the fact that Case 2 or Case 3 applies to μ, or the cycle composed of $p(\Gamma_1, \tau, \mu)$ and $p(\Gamma_1, \tau, \nu)$ entirely belongs to μ and contains in its interior a vertex not belonging to μ. This contradicts Lemma 1.
  - Assume next that $p(\Gamma_1, \tau, \mu)$ and $p(\Gamma_1, \tau, \nu)$ are the same path.
    - If Case 2 applies to μ, then Case 2 applies to ν as well. In fact, $G(\overline{\tau})$ is not ν-full, given that μ is an ancestor of ν and given that $G(\overline{\tau})$ is not μ-full. If $p(\Gamma_2, \overline{\tau}, \mu)$ and $p(\Gamma_2, \overline{\tau}, \nu)$ are distinct, then either all vertices of $pert(\tau)$ belong to μ, hence τ

is μ-full, thus contradicting the fact that Case 2 applies to μ, or the cycle composed of $p(\Gamma_2, \overline{\tau}, \mu)$ and $p(\Gamma_2, \overline{\tau}, \nu)$ entirely belongs to μ and contains in its interior a vertex not belonging to μ. This contradicts Lemma 1. Hence, $p(\Gamma_2, \overline{\tau}, \mu)$ and $p(\Gamma_2, \overline{\tau}, \nu)$ are the same path. However, since $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \overline{\tau}, \mu) = p(\Gamma_2, \overline{\tau}, \nu)$, clusters μ and ν determine the same flip for $\Gamma_1(pert(\tau))$, a contradiction to the assumptions.

* If Case 3 applies to μ and Case 2 applies to ν, then we argue as follows. If $p(\Gamma_2, \overline{\tau}, \mu)$ and $p(\Gamma_2, \overline{\tau}, \nu)$ are distinct, then $p(\Gamma_2, \overline{\tau}, \nu)$ delimits the outer face of Γ_2, given that Case 3 applies to μ and given that $p(\Gamma_2, \overline{\tau}, \mu)$ does not delimit the outer face of Γ_2. Thus, the cycle $\mathcal{C}$ composed of $p(\Gamma_2, \overline{\tau}, \nu)$ and of any path in $pert(\tau)$ connecting the poles of τ and entirely belonging to ν (such a path exists given that $\Gamma_1(pert(\tau))$ is ν-spined) contains in its interior in Γ_2 all the vertices of $G(\overline{\tau})$ not in $p(\Gamma_2, \overline{\tau}, \nu)$. Since $G(\overline{\tau})$ is not ν-full, $\mathcal{C}$ entirely belongs to ν and contains in its interior a vertex not belonging to ν. This contradicts Lemma 1. If $p(\Gamma_2, \overline{\tau}, \mu)$ and $p(\Gamma_2, \overline{\tau}, \nu)$ are the same path, then, since $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \overline{\tau}, \mu) = p(\Gamma_2, \overline{\tau}, \nu)$, clusters μ and ν determine the same flip for $\Gamma_1(pert(\tau))$, a contradiction to the assumptions.
* If Case 3 applies to both μ and ν, then $p(\Gamma_2, \overline{\tau}, \mu) = p(\Gamma_2, \overline{\tau}, \nu)$ is the path that connects the poles of τ, belongs to $G(\overline{\tau})$, and does not delimit the outer face of Γ_2. Since $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \overline{\tau}, \mu) = p(\Gamma_2, \overline{\tau}, \nu)$, clusters μ and ν determine the same flip for $\Gamma_1(pert(\tau))$, a contradiction to the assumptions.

- Suppose next that Case 1 applies both to μ and ν to determine a different flip for $\Gamma_1(pert(\tau))$.

  By assumption, $\Gamma_1(pert(\tau))$ and $\Gamma_2(G(\overline{\tau}))$ are μ-sided. It follows that $\Gamma_1(G(\overline{\tau}))$ and $\Gamma_2(pert(\tau))$ are μ-sided, as well. In fact, they are not μ-traversable by Lemma 3. Also, if they are μ-kernelized, or μ-infeasible, or μ-bisided, then $H(\mu)$ would not be connected in Γ_1 or in Γ_2. This contradicts Lemma 1. Analogously, $\Gamma_1(pert(\tau))$, $\Gamma_2(G(\overline{\tau}))$, $\Gamma_1(G(\overline{\tau}))$, and $\Gamma_2(pert(\tau))$ are ν-sided.

  Assume, w.l.o.g. up to renaming $f'(\tau)$ with $f''(\tau)$ in Γ_i, that $p(\Gamma_i, \overline{\tau}, \nu)$ is incident to $f''(\tau)$, for $i = 1, 2$.

  Our strategy is to show that either $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \tau, \mu) = p(\Gamma_2, \tau, \nu)$ (Condition A) or that $p(\Gamma_1, \tau, \mu) \neq p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$ (Condition B).

  We first show that, if Condition A or Condition B holds, then the assumption that ν and μ determine different flips for $\Gamma_1(pert(\tau))$ is contradicted. Suppose first that Condition A holds. Then, $p(\Gamma_1, \overline{\tau}, \mu) = p(\Gamma_1, \overline{\tau}, \nu)$ and $p(\Gamma_2, \overline{\tau}, \mu) = p(\Gamma_2, \overline{\tau}, \nu)$ and such paths are all incident to $f''(\tau)$, as otherwise $H(\mu)$ or $H(\nu)$ would not be connected in Γ_1 or in Γ_2, which is a contradiction by Lemma 1. Hence, the flip determined for $\Gamma_1(pert(\tau))$ by μ and ν is the same, a contradiction. Suppose next that Condition B holds. Then, $p(\Gamma_1, \overline{\tau}, \mu) \neq p(\Gamma_1, \overline{\tau}, \nu)$ and $p(\Gamma_2, \overline{\tau}, \mu) \neq p(\Gamma_2, \overline{\tau}, \nu)$, where $p(\Gamma_1, \overline{\tau}, \mu)$ and $p(\Gamma_2, \overline{\tau}, \mu)$ are incident to $f'(\tau)$, while $p(\Gamma_1, \overline{\tau}, \nu)$ and $p(\Gamma_2, \overline{\tau}, \nu)$ are incident to $f''(\tau)$, as otherwise $H(\mu)$ or $H(\nu)$ would not be connected in Γ_1 or in Γ_2, which is a contradiction by Lemma 1. Hence, the flip determined for $\Gamma_1(pert(\tau))$ by μ and ν is the same, a contradiction.

We now prove that either Condition A or Condition B holds. Namely, we prove by contradiction that $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$ do not hold simultaneously. The proof that $p(\Gamma_1, \tau, \mu) \neq p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \tau, \mu) = p(\Gamma_2, \tau, \nu)$ do not hold simultaneously is symmetrical. Since $p(\Gamma_i, \overline{\tau}, \nu)$ is incident to $f''(\tau)$, for $i = 1, 2$, we have that $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ is incident to $f''(\tau)$ in Γ_1, as otherwise $H(\nu)$ would not be connected in Γ_1, that $p(\Gamma_1, \overline{\tau}, \mu)$ is incident to $f''(\tau)$ in Γ_1, as otherwise $H(\mu)$ would not be connected in Γ_1, that $p(\Gamma_2, \tau, \nu)$ is incident to $f''(\tau)$ in Γ_2, as otherwise $H(\nu)$ would not be connected in Γ_2, that $p(\Gamma_2, \tau, \mu)$ is incident to $f'(\tau)$ in Γ_1, since $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$, and that $p(\Gamma_2, \overline{\tau}, \mu)$ is incident to $f'(\tau)$ in Γ_2, as otherwise $H(\mu)$ would not be connected in Γ_2.

We distinguish two cases.

Suppose that τ is an R-node. Denote by $V_\tau(\mu)$ and by $V_\tau(\nu)$ the set of vertices in $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ belonging to μ and ν, respectively (see Fig. 12(a)). Since $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$ and since $\Gamma_2(pert(\tau))$ is μ-sided and ν-sided, it follows that none of the vertices in $V_\tau(\mu)$ is incident to $f''(\tau)$ in Γ_2 and none of the vertices in $V_\tau(\nu)$ is incident to $f'(\tau)$ in Γ_2. Since τ is an R-node, all the vertices in $V_\tau(\mu)$ are incident to internal faces of $\Gamma_2(pert(\tau))$ or all the vertices in $V_\tau(\nu)$ are incident to internal faces of $\Gamma_2(pert(\tau))$. Suppose that all the vertices in $V_\tau(\mu)$ are incident to internal faces of $\Gamma_2(pert(\tau))$, the other case being analogous (see Fig. 12(b)). Then, in order for $H(\mu)$ to be connected and for $\Gamma_2(pert(\tau))$ to be μ-sided, a child τ_i of τ in $\mathcal{T}$ that is incident to $f'(\tau)$ is μ-traversable in Γ_2. By Lemma 3, we have that τ_i is either μ-infeasible or μ-traversable in Γ_1. In the former case, a contradiction to the fact that Γ_1 is a $\langle 0, 0, \infty \rangle$-drawing of $C(G, T)$ is obtained. In the latter case, since τ_i is incident to $f'(\tau)$ in Γ_1, $\Gamma_1(pert(\tau))$ is not μ-sided, a contradiction.

Next, suppose that τ is a P-node. Consider the sequence $\tau_1, \tau_2, \dots, \tau_p$ of children of τ that have vertices belonging to ν as they appear in $\Gamma_1(pert(\tau))$, where τ_1 is incident to $f''(\tau)$. Since $H(\nu)$ is connected in Γ_1, it follows that τ_i is ν-traversable in Γ_1, for $i = 1, \dots, p-1$, and that τ_p is either ν-traversable or ν-sided. Analogous considerations hold for the sequence $\tau_1, \tau_2, \dots, \tau_q$ of children of τ that have vertices belonging to μ. We distinguish two cases: Either $p \neq q$ (see Fig. 13(a)) or $p = q$ (see Fig. 13(b)).

In the first case suppose, without loss of generality, that $p > q$. Then, for

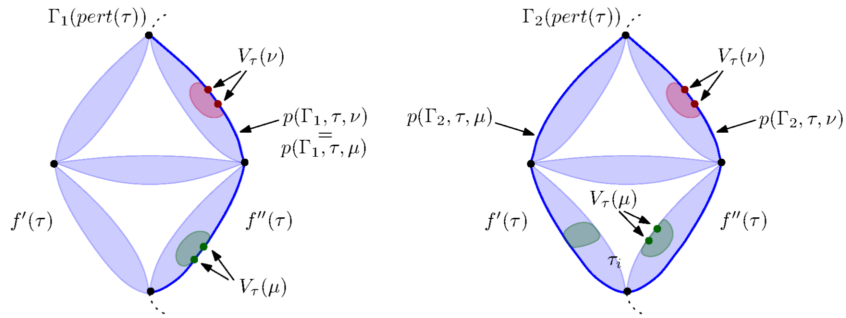

(a) (b)

Figure 12: Illustration for the proof that $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$ do not simultaneously hold if τ is an R-node. (a) Drawing Γ_1 with $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$. (b) Drawing Γ_2 with $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$.

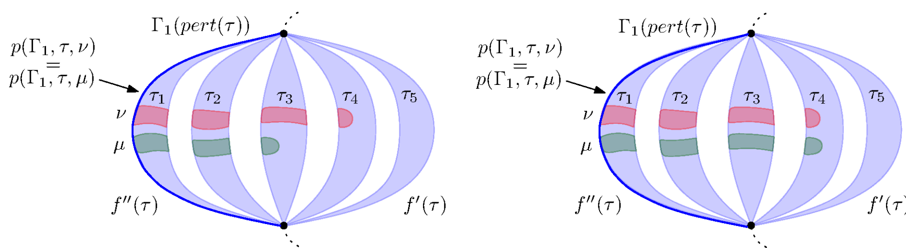

(a) (b)

Figure 13: Illustration for the proof that $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$ do not simultaneously hold if τ is a P-node. (a) The case $p > q$, with $p = 4$ and $q = 3$. (b) The case $p = q = 4$.

every $1 \leq i \leq q$, τ_i is ν-traversable in Γ_1. By Lemma 3, the embedding of $pert(\tau_i)$ in Γ_2 is either ν-infeasible or ν-traversable. In the former case, a contradiction to the fact that Γ_2 is a $\langle 0, 0, \infty \rangle$-drawing of $C(G, T)$ is obtained. In the latter case, the child of τ that is incident to $f'(\tau)$ in Γ_2 is ν-traversable, thus contradicting the assumption that $\Gamma_2(pert(\tau))$ is ν-sided. Suppose next that $p = q$. If τ_p is not incident to $f'(\tau)$ in Γ_2, then the child of τ that is incident to $f'(\tau)$ in Γ_2 is one of $\tau_1, \tau_2, \ldots, \tau_{p-1}$, hence, by Lemma 3, it is either ν-infeasible in Γ_2, thus contradicting the fact that Γ_2 is a $\langle 0, 0, \infty \rangle$-drawing of $C(G, T)$, or it is ν-traversable in Γ_2, thus contradicting the assumption that $\Gamma_2(pert(\tau))$ is ν-sided. Analogously, if τ_p is not incident to $f''(\tau)$ in Γ_2, then the child of τ that is incident to $f''(\tau)$ is one of $\tau_1, \tau_2, \ldots, \tau_{q-1}$, hence, by Lemma 3, it is either μ-infeasible in Γ_2, thus contradicting the fact that Γ_2 is a $\langle 0, 0, \infty \rangle$-drawing of $C(G, T)$, or it is μ-traversable in Γ_2, thus contradicting the assumption that $\Gamma_2(pert(\tau))$ is μ-sided. We have a contradiction as τ_p can not be at the same time incident to both $f'(\tau)$ and $f''(\tau)$.

- Finally, suppose that Case 2 or Case 3 applies to μ and that Case 1 applies to ν to determine a different flip for $\Gamma_1(pert(\tau))$.

  We show how to restrict to the case in which μ is an ancestor of ν. First, if ν is an ancestor of μ, then $p(\Gamma_1, \tau, \mu)$ entirely belongs to ν, hence τ is ν-traversable, a contradiction to the fact that Case 1 applies to ν. Second, since Case 2 or Case 3 applies to μ, it follows that $p(\Gamma_1, \tau, \mu)$ and $p(\Gamma_2, \overline{\tau}, \mu)$ are well-defined. Also, since Case 1 applies to ν, it follows that $p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \overline{\tau}, \nu)$ are also well-defined. Now suppose, for a contradiction, that μ is not an ancestor of ν. Since ν is not an ancestor of μ, it follows that μ and ν do not share vertices. Hence, $p(\Gamma_1, \tau, \mu) \neq p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \overline{\tau}, \mu) \neq p(\Gamma_2, \overline{\tau}, \nu)$. However, this implies that μ and ν determine the same flip for $\Gamma_1(pert(\tau))$, a contradiction to the assumptions. We can hence assume that μ is an ancestor of ν.

  Since Case 1 applies to ν, we have that $\Gamma_1(pert(\tau))$ is ν-sided. It follows that $\Gamma_2(pert(\tau))$ is ν-sided, as well. Namely, it is not ν-traversable by Lemma 3. Also, if it is ν-kernelized, or ν-infeasible, or ν-bisided, then $H(\nu)$ would not be connected in Γ_2, given that $\Gamma_2(G(\overline{\tau}))$ is ν-sided and hence not ν-traversable. This contradicts Lemma 1. An analogous argument proves that $\Gamma_1(G(\overline{\tau}))$ is ν-sided.

  - Suppose that Case 2 applies to μ. Then, $\Gamma_1(pert(\tau))$ and $\Gamma_2(G(\overline{\tau}))$ are both μ-side-spined and not μ-full. If $\Gamma_1(G(\overline{\tau}))$ is μ-central-spined, then Γ_1 contains a cycle whose vertices belong to μ containing in its interior a vertex in $G(\overline{\tau})$ not belonging to μ. This contradicts Lemma 1. It follows that $\Gamma_1(G(\overline{\tau}))$ is μ-side-spined and not μ-full. An analogous argument proves that $\Gamma_2(pert(\tau))$ is μ-side-spined and not μ-full.

    Assume, without loss of generality up to renaming $f'(\tau)$ with $f''(\tau)$ that $p(\Gamma_i, \overline{\tau}, \mu)$ is incident to $f''(\tau)$ in Γ_i, for $i = 1, 2$. Then, for $i = 1, 2$, path $p(\Gamma_i, \tau, \mu)$ is incident to $f''(\tau)$ in Γ_i, as otherwise $p(\Gamma_i, \overline{\tau}, \mu)$ together with $p(\Gamma_i, \tau, \mu)$ forms a cycle that entirely belongs to μ and that contains in its interior a vertex not belonging to μ, which by Lemma 1 contradicts the assumption that Γ_i is a $\langle 0, 0, \infty \rangle$-drawing of $C(G, T)$.

    Our strategy is to show that either $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \tau, \mu) = p(\Gamma_2, \tau, \nu)$ (Condition A) or that $p(\Gamma_1, \tau, \mu) \neq p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$ (Condition B).

    We first show that, if Condition A or Condition B holds, then the assumption that ν and μ determine different flips for $\Gamma_1(pert(\tau))$ is contradicted. Suppose first that Condition A holds. Then, for $i = 1, 2$, we have that $p(\Gamma_i, \overline{\tau}, \mu) = p(\Gamma_i, \overline{\tau}, \nu)$, as otherwise $H(\nu)$ would not be connected in Γ_i, which is a contradiction by Lemma 1. Hence, the flip determined for $\Gamma_1(pert(\tau))$ by μ and ν is the same, a contradiction. Suppose next that Condition B holds. Then, for $i = 1, 2$, we have that $p(\Gamma_i, \overline{\tau}, \mu) \neq p(\Gamma_i, \overline{\tau}, \nu)$, as otherwise $H(\nu)$

would not be connected in Γ_i, which is a contradiction by Lemma 1. Hence, the flip determined for $\Gamma_1(pert(\tau))$ by μ and ν is the same, a contradiction.

We now prove that either Condition A or Condition B holds. Namely, we prove by contradiction that $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$ do not simultaneously hold. The proof that $p(\Gamma_1, \tau, \mu) \neq p(\Gamma_1, \tau, \nu)$ and $p(\Gamma_2, \tau, \mu) = p(\Gamma_2, \tau, \nu)$ do not simultaneously hold is symmetrical. Hence, $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ is incident to $f''(\tau)$ in Γ_1, $p(\Gamma_1, \overline{\tau}, \nu)$ is incident to $f''(\tau)$ in Γ_1, as otherwise $H(\nu)$ would not be connected in Γ_1, $p(\Gamma_2, \tau, \nu) \neq p(\Gamma_2, \tau, \mu)$ is incident to $f'(\tau)$ in Γ_2, and $p(\Gamma_2, \overline{\tau}, \nu)$ is incident to $f'(\tau)$ in Γ_2, as otherwise $H(\nu)$ would not be connected in Γ_2.

We distinguish two cases.

Suppose that τ is an R-node. Denote by $V_\tau(\nu)$ the set of vertices in $p(\Gamma_1, \tau, \mu) = p(\Gamma_1, \tau, \nu)$ belonging to ν. Since $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$ and since $\Gamma_2(pert(\tau))$ is ν-sided, it follows that none of the vertices in $V_\tau(\nu)$ is incident to $f''(\tau)$ in Γ_2. Hence, all the vertices in $V_\tau(\nu)$ are incident to internal faces of $\Gamma_2(pert(\tau))$. Then, in order for $H(\nu)$ to be connected and for $\Gamma_2(pert(\tau))$ to be ν-sided, a child τ_i of τ in $\mathcal{T}$ that is incident to $f'(\tau)$ in Γ_2 is ν-traversable. By Lemma 3, we have that τ_i is either μ-infeasible or μ-traversable in Γ_1. In the former case, a contradiction to the fact that Γ_1 is a $\langle 0, 0, \infty \rangle$-drawing of $C(G, T)$ is obtained. In the latter case, since τ is an R-node, τ_i is incident to $f'(\tau)$ in Γ_1. This contradicts the fact that $\Gamma_1(pert(\tau))$ is ν-sided.

Next, suppose that τ is a P-node. Let k be the number of children of τ in $\mathcal{T}$. Consider the sequence $\tau_1, \tau_2, \ldots, \tau_p$ of children of τ that have vertices belonging to ν as they appear in $\Gamma_1(pert(\tau))$, where τ_1 is incident to $f''(\tau)$. Since $H(\nu)$ is connected in Γ_1, it follows that τ_i is ν-traversable in Γ_1, for $i = 1, \ldots, p-1$, and that τ_p is either ν-traversable or ν-sided. Also, consider the sequence $\tau_1, \tau_2, \ldots, \tau_q$ of children of τ that are μ-spined, as they appear in $\Gamma_1(pert(\tau))$, where τ_1 is incident to $f''(\tau)$. Since $\Gamma_1(pert(\tau))$ is μ-side-spined, it follows that τ_i is μ-full, for $i = 1, \ldots, q-1$, and that τ_q is either μ-full or μ-side-spined. Observe that $p, q \geq 1$, since $\Gamma_1(pert(\tau))$ is ν-sided and μ-side-spined. Also observe that $q < k$ or τ_q is not μ-full, as otherwise τ would be μ-full, which contradicts the assumptions.

We distinguish some cases.

* Suppose first that $p \leq q < k$. Since $\Gamma_2(pert(\tau))$ is μ-side-spined and $p(\Gamma_2, \tau, \mu)$ is incident to $f''(\tau)$ in Γ_2, it follows that $\tau_1, \tau_2, \ldots, \tau_q$ are the first q children of τ as they appear in $\Gamma_2(pert(\tau))$, possibly in a different relative order with respect to their order in $\Gamma_1(pert(\tau))$, where one of $\tau_1, \tau_2, \ldots, \tau_q$ is incident to $f''(\tau)$. Hence, the child of τ incident to $f'(\tau)$ does not contain any vertex belonging to ν, which contradicts either $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$

or the fact that $\Gamma_2(pert(\tau))$ is ν-sided.

* Suppose next that $p = q = k$. Then, τ_q is not μ-full. Since $\Gamma_2(pert(\tau))$ is μ-side-spined and $p(\Gamma_2, \tau, \mu)$ is incident to $f''(\tau)$ in Γ_2, it follows that τ_q is the child of τ incident to $f'(\tau)$ in Γ_2. Thus, by Lemma 3, the child of τ incident to $f''(\tau)$ in Γ_2 is either ν-infeasible, thus contradicting the fact that Γ_2 is a $\langle 0, 0, \infty \rangle$-drawing of $C(G, T)$, or it is ν-traversable, which contradicts either $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$ or the fact that $\Gamma_2(pert(\tau))$ is ν-sided.
* Suppose next that $p < q = k$. Then, τ_q is not μ-full. Since $\Gamma_2(pert(\tau))$ is μ-side-spined and $p(\Gamma_2, \tau, \mu)$ is incident to $f''(\tau)$ in Γ_2, it follows that τ_q is the child of τ incident to $f'(\tau)$ in Γ_2. Hence, no vertex in ν is incident to $f'(\tau)$ in Γ_2, which contradicts either $p(\Gamma_2, \tau, \mu) \neq p(\Gamma_2, \tau, \nu)$ or the fact that $\Gamma_2(pert(\tau))$ is ν-sided.
* Suppose finally that $p > q$. Then, all of $\tau_1, \tau_2, \dots, \tau_q$ are ν-traversable. Since $\tau_1, \tau_2, \dots, \tau_q$ are the first q children of τ as they appear in $\Gamma_2(pert(\tau))$, possibly in a different relative order, where one of $\tau_1, \tau_2, \dots, \tau_q$ is incident to $f''(\tau)$, by Lemma 3 and since no child of τ is ν-infeasible in Γ_2, as otherwise Γ_2 would not be a $\langle 0, 0, \infty \rangle$-drawing of $C(G, T)$, it follows that $\Gamma_2(pert(\tau))$ has a vertex belonging to ν and incident to $f''(\tau)$, thus contradicting the assumption that $p(\Gamma_1, \tau, \nu)$ is incident to $f''(\tau)$ or the assumption that $\Gamma_1(pert(\tau))$ is ν-sided.

– Suppose that Case 3 applies to μ. Then, $G(\overline{\tau})$ is μ-full. By assumption, $\Gamma_1(pert(\tau))$ is μ-side-spined and not μ-full. If at least one of the poles of τ is not incident to the outer face of Γ_1, then Γ_1 contains a cycle in $G(\overline{\tau})$ whose vertices belong to μ containing in its interior a vertex in $pert(\tau)$ not belonging to μ. This contradicts Lemma 1. Assume then that both poles of τ are incident to the outer face of Γ_1. If $p(\Gamma_1, \tau, \mu)$ delimits the outer face of Γ_1, then $p(\Gamma_1, \tau, \mu)$ together with $p(\Gamma_1, \overline{\tau}, \mu)$ forms a cycle whose vertices belong to μ containing in its interior a vertex in $pert(\tau)$ not belonging to μ, again contradicting the assumption that Γ_1 is a $\langle 0, 0, \infty \rangle$-drawing of $C(G, T)$, by Lemma 1. Also, recall that, by assumption, both the poles of τ are incident to the outer face of Γ_2.

  Assume, without loss of generality up to renaming $f'(\tau)$ with $f''(\tau)$ with that $p(\Gamma_i, \overline{\tau}, \mu)$ is incident to $f''(\tau)$ in Γ_i, for $i = 1, 2$. Hence, $f'(\tau)$ is the outer face of Γ_i, for $i = 1, 2$, and $p(\Gamma_i, \overline{\tau}, \mu)$ delimits $f''(\tau)$ in Γ_i, for $i = 1, 2$, by definition.

  The reminder of the proof is exactly the same as when Case 2 applies to μ.

This concludes the proof of the claim. □

By Claim 3, no two clusters $\mu \neq \nu \in T$ exist that determine different flips for $\Gamma_1(pert(\tau))$ when constructing Γ_3. Then, by Claims 1 and 2, the constructed embedding Γ_3 of $C(G,T)$ satisfies Conditions (i) and (ii) of Lemma 1, for each cluster μ, hence Γ_3 is a $\langle 0,0,\infty \rangle$-drawing. This concludes the proof of Lemma 4.

We now determine conditions for the nodes of the SPQR-tree $\mathcal{T}$ of G in order for $C(G,T)$ to admit a $\langle 0,0,\infty \rangle$-drawing.

We say that an embedding $\Gamma(sk(\tau))$ of $sk(\tau)$ in which the edge e representing the parent of τ is incident to the outer face is *extensible* if the following condition holds: If $C(G,T)$ admits a $\langle 0,0,\infty \rangle$-drawing in which e_ρ is incident to the outer face, then it admits a $\langle 0,0,\infty \rangle$-drawing in which e_ρ is incident to the outer face and in which the embedding of $sk(\tau)$ is $\Gamma(sk(\tau))$.

For an embedding $\Gamma(sk(\tau))$ of $sk(\tau)$ and a cluster μ in T, we define an auxiliary graph $G'(\tau,\mu)$ as follows. Graph $G'(\tau,\mu)$ has one vertex v_f for each face f of $\Gamma(sk(\tau))$ containing a μ-traversable virtual edge on its boundary; two vertices of $G'(\tau,\mu)$ are connected by an edge if they share a μ-traversable virtual edge. We are now ready to prove the following main lemma.

Lemma 5. *Let $C(G,T)$ be a clustered graph, with G biconnected, that admits a $\langle 0,0,\infty \rangle$-drawing in which the edge e_ρ representing the root ρ of the SPQR-tree $\mathcal{T}$ of G is incident to the outer face. Then, an embedding $\Gamma(sk(\tau))$ in which the edge e representing the parent of τ is incident to the outer face is extensible if and only if the following properties hold. For each cluster $\mu \in T$:*

(A) There exists no cycle in $\Gamma(sk(\tau))$ that is composed of μ-spined virtual edges $e_1,\dots,e_h$ containing in its interior a virtual edge that is not μ-full;

(B) $G'(\tau,\mu)$ is connected; and

(C) if $G'(\tau,\mu)$ contains at least one vertex, then each virtual edge of $sk(\tau)$ which is μ-touched and not μ-traversable shares a face with a μ-traversable virtual edge in $\Gamma(sk(\tau))$. Otherwise (that is, if $G'(\tau,\mu)$ contains no vertex), all the μ-touched virtual edges are incident to the same face of $\Gamma(sk(\tau))$.

In the following we prove Lemma 5. We first prove the necessity. Suppose that an embedding $\Gamma(sk(\tau))$ of $sk(\tau)$ is extensible. That is, $C(G,T)$ admits a $\langle 0,0,\infty \rangle$-drawing in which e_ρ is incident to the outer face and in which the embedding of $sk(\tau)$ is $\Gamma(sk(\tau))$. We prove that $\Gamma(sk(\tau))$ satisfies Properties (A), (B), and (C) by means of suitable claims.

Claim 4. *$\Gamma(sk(\tau))$ satisfies Property (A).*

Proof: No cycle $(e_1,\dots,e_h)$ in $\Gamma(sk(\tau))$ such that virtual edges $e_1,\dots,e_h$ are μ-spined contains in its interior a virtual edge that is not μ-full, as otherwise, in any drawing Γ of $C(G,T)$ in which e_ρ is incident to the outer face and in which the embedding of $sk(\tau)$ is $\Gamma(sk(\tau))$, there would exist a cycle whose vertices all belong to μ enclosing a vertex not belonging to μ, thus implying that $R(\mu)$ is not simple or that Γ contains an edge-region crossing. $\square$

Claim 5. $\Gamma(sk(\tau))$ *satisfies Property (B).*

Proof: We have that $G'(\tau, \mu)$ is connected, as otherwise, in any drawing Γ of $C(G, T)$ in which e_ρ is incident to the outer face and in which the embedding of $sk(\tau)$ is $\Gamma(sk(\tau))$, there would exist a cycle $\mathcal{C}$ such that none of the vertices of $\mathcal{C}$ belongs to μ and such that $\mathcal{C}$ separates vertices belonging to μ, thus implying that $R(\mu)$ is not simple or that Γ contains an edge-region crossing. $\square$

Claim 6. $\Gamma(sk(\tau))$ *satisfies Property (C).*

Proof: We distinguish the case in which $G'(\tau, \mu)$ contains at least one vertex and the case in which $G'(\tau, \mu)$ contains no vertex.

- Suppose that $G'(\tau, \mu)$ contains at least one vertex. Refer to Fig. 15(a). Then, we prove that each virtual edge of $sk(\tau)$ which is μ-touched and not μ-traversable shares a face with a μ-traversable virtual edge in $\Gamma(sk(\tau))$.

  Suppose, for a contradiction, that there exists a virtual edge e of $sk(\tau)$ that is μ-touched, that is not μ-traversable, and that does not share any face with a μ-traversable virtual edge in $\Gamma(sk(\tau))$. Consider the two faces f_e^1 and f_e^2 of $\Gamma(sk(\tau))$ incident to e. Denote by $\mathcal{C}_s$ the cycle of virtual edges composed of the edges delimiting f_e^1 and f_e^2, except for e.

  By assumption, for each edge e_i of $\mathcal{C}_s$, the node τ_i of $\mathcal{T}$ corresponding to e_i is such that any embedding of $pert(\tau_i)$ is not μ-traversable. Also, in any $\langle 0, 0, \infty \rangle$-drawing of $C(G, T)$, the embedding of $pert(\tau_i)$ is not μ-infeasible. Hence, by Lemma 2, there exists a path in $pert(\tau_i)$ that connects the poles of τ_i, that is different from the edge connecting the poles of τ_i, and none of whose vertices belongs to μ. Concatenating all such paths for all the nodes of $\mathcal{T}$ corresponding to the edges of $\mathcal{C}_s$ results in a cycle $\mathcal{C}$ such that none of the vertices of $\mathcal{C}$ belongs to μ and such that $\mathcal{C}$ passes through all the vertices of $\mathcal{C}_s$. Observe that $\mathcal{C}$ contains vertices of μ on both sides, namely vertices of μ in $pert(e)$ and vertices of μ in the pertinent graph of a virtual edge e' of $sk(\tau)$ which is μ-traversable; such an edge e' exists since $G'(\tau, \mu)$ contains at least one vertex. Hence, in any drawing Γ of $C(G, T)$ in which the embedding of $sk(\tau)$ is $\Gamma(sk(\tau))$, there exists a cycle $\mathcal{C}$ such that none of the vertices of $\mathcal{C}$ belongs to μ and such that $\mathcal{C}$ separates vertices belonging to μ, thus implying that $R(\mu)$ is not simple or that Γ contains an edge-region crossing, a contradiction.

- Suppose that $G'(\tau, \mu)$ contains no vertex. Refer to Fig. 15(b). We prove that all the μ-touched virtual edges are incident to the same face of $\Gamma(sk(\tau))$.

  Suppose, for a contradiction, that there exists no face of $\Gamma(sk(\tau))$ such that all the μ-touched virtual edges of $\Gamma(sk(\tau))$ are incident to such a face. Consider the two faces f_e^1 and f_e^2 of $\Gamma(sk(\tau))$ incident to any μ-touched virtual edge e. Denote by $\mathcal{C}_s^1$ and $\mathcal{C}_s^2$ the cycles delimiting f_e^1 and f_e^2, respectively.

By assumption, for each edge e_i of $\mathcal{C}_s^1$ (of $\mathcal{C}_s^2$), the node τ_i of $\mathcal{T}$ corresponding to e_i is such that any embedding of $pert(\tau_i)$ is not μ-traversable. Also, in any $\langle 0,0,\infty\rangle$-drawing of $C(G,T)$, the embedding of $pert(\tau_i)$ is not μ-infeasible. Hence, by Lemma 2, there exists a path in $pert(\tau_i)$ that connects the poles of τ_i, that is different from the edge connecting the poles of τ_i, and none of whose vertices belongs to μ. Concatenating all such paths for all the nodes of $\mathcal{T}$ corresponding to the edges of $\mathcal{C}_s^1$ (resp. of $\mathcal{C}_s^2$) results in a cycle $\mathcal{C}^1$ (resp. $\mathcal{C}^2$) such that none of the vertices of $\mathcal{C}^1$ (resp. of $\mathcal{C}^2$) belongs to μ and such that $\mathcal{C}^1$ (resp. $\mathcal{C}^2$) passes through all the vertices of $\mathcal{C}_s^1$ (resp. $\mathcal{C}_s^2$). Then, $\mathcal{C}^1$ or $\mathcal{C}^2$ contains vertices of μ on both sides, namely vertices of μ in $pert(e)$ and vertices of μ in the pertinent graph of a virtual edge e' of $sk(\tau)$ which is not incident to f_e^1 or to f_e^2, respectively; such an edge e' exists since not all the μ-touched virtual edges are incident to the same face of $\Gamma(sk(\tau))$. This implies that $R(\mu)$ is not simple or that Γ contains an edge-region crossing.

This concludes the proof of the claim. □

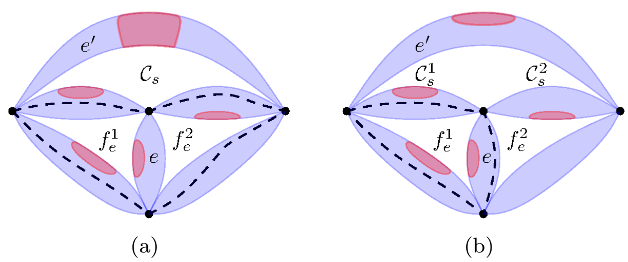

Figure 14: Proof that Property (C) is satisfied when $C(G,T)$ admits a $\langle 0,0,\infty\rangle$-drawing in which e_ρ is incident to the outer face and in which the embedding of $sk(\tau)$ is $\Gamma(sk(\tau))$. (a) $G'(\tau,\mu)$ contains at least one vertex and (b) $G'(\tau,\mu)$ contains no vertex. In both figures, cycle $\mathcal{C}$ is represented by a dashed curve.

We now prove the sufficiency. Namely, suppose that Properties (A), (B), and (C) hold for an embedding $\Gamma(sk(\tau))$ of $sk(\tau)$. We prove that $\Gamma(sk(\tau))$ is extensible, that is, we show that a $\langle 0,0,\infty\rangle$-drawing of $C(G,T)$ exists in which e_ρ is incident to the outer face and in which the embedding of $sk(\tau)$ is $\Gamma(sk(\tau))$.

Let Γ' be any $\langle 0,0,\infty\rangle$-drawing of $C(G,T)$ in which e_ρ is incident to the outer face (clustered graph $C(G,T)$ admits such a drawing by hypothesis). Let $\Gamma'(sk(\tau))$ be the embedding of $sk(\tau)$ in Γ'. If $\Gamma'(sk(\tau))$ coincides with $\Gamma(sk(\tau))$, then there is nothing to prove. Otherwise, assume that the two embeddings of $sk(\tau)$ do not coincide. Observe that this implies that τ is not an S-node, as the skeleton of an S-node is a cycle, which has a unique embedding.

Next, suppose that τ is an R-node. Then, since $sk(\tau)$ has exactly two embeddings, which are one the flip of the other, $\Gamma(sk(\tau))$ is the flip of $\Gamma'(sk(\tau))$. Consider the drawing Γ of $C(G,T)$ obtained by flipping Γ' around the poles of the root ρ of $\mathcal{T}$, that is, by reverting the adjacency list of every vertex of $C(G,T)$. Observe that Γ is a $\langle 0,0,\infty\rangle$-drawing since Γ' is. Also, e_ρ is incident to the outer face of Γ since it is incident to the outer face of Γ'. Moreover, the embedding of $sk(\tau)$ in Γ is the flip of the embedding $\Gamma'(sk(\tau))$ of $sk(\tau)$ in Γ', hence it coincides with $\Gamma(sk(\tau))$. Thus a $\langle 0, 0, \infty \rangle$-drawing Γ of $C(G, T)$ in which e_ρ is incident to the outer face and in which the embedding of $sk(\tau)$ is $\Gamma(sk(\tau))$ exists.

It remains to consider the case in which τ is a P-node. We show how to construct a $\langle 0, 0, \infty \rangle$-drawing Γ^C in which e_ρ is incident to the outer face and in which the embedding of $sk(\tau)$ is $\Gamma(sk(\tau))$. For every neighbor τ' of τ in $\mathcal{T}$ (including its parent), denote by $\Gamma^1(pert(\tau'))$ the embedding of $pert(\tau')$ in Γ'. Moreover, denote by $\Gamma^2(pert(\tau'))$ the embedding of $pert(\tau')$ obtained by flipping $\Gamma^1(pert(\tau'))$ around the poles of τ'. Drawing Γ^C is such that, for every neighbor τ' of τ in $\mathcal{T}$, the embedding of $pert(\tau')$ is either $\Gamma^1(pert(\tau'))$ or $\Gamma^2(pert(\tau'))$.

The remainder of the proof is devoted to the P-node case and is structured into three parts as follows.

In the first part, we describe how to obtain Γ^C. In particular, we show how to choose the embedding of $pert(\tau')$ to be either $\Gamma^1(pert(\tau'))$ or $\Gamma^2(pert(\tau'))$ based on the constraints imposed by the clusters containing vertices of $pert(\tau')$.

In the second part, we show that the algorithm we describe in the first part univocally determines the embedding of $pert(\tau')$ to be either $\Gamma^1(pert(\tau'))$ or $\Gamma^2(pert(\tau'))$.

In the third part, we show that the performed choices actually lead to a $\langle 0, 0, \infty \rangle$-drawing of $C(G, T)$ in which e_ρ is incident to the outer face and in which the embedding of $sk(\tau)$ is $\Gamma(sk(\tau))$.

First part. We show an algorithm to determine Γ^C. We fix the embedding of $sk(\tau)$ to be $\Gamma(sk(\tau))$ with an outer face that is any of the two faces incident to the virtual edge of $sk(\tau)$ representing the parent of τ in $\mathcal{T}$. Later, we will possibly modify the choice of the outer face. We now show how to choose the embedding of $pert(\tau')$. This is done according to rules that aim at satisfying Conditions (i) and (ii) of Lemma 1. Namely, for each cluster $\mu \in T$ apply one of the following rules.

- *Rules to satisfy Condition (i) of Lemma 1.*

  Consider any neighbor τ' of τ in $\mathcal{T}$ such that $\Gamma^1(pert(\tau'))$ is μ-sided. Also, consider the neighbors τ'' and τ''' of τ in $\mathcal{T}$ following and preceding τ' in the circular order of the neighbors of τ determined by $\Gamma(sk(\tau))$, respectively. Without loss of generality assume that if the embedding of $pert(\tau')$ is $\Gamma^1(pert(\tau'))$, then a vertex in $pert(\tau')$ and in μ is incident to the face of $\Gamma(sk(\tau))$ to which τ'' is incident; and if the embedding of $pert(\tau')$ is $\Gamma^2(pert(\tau'))$, then a vertex in $pert(\tau')$ and in μ is incident to the face of $\Gamma(sk(\tau))$ to which τ''' is incident.

  - **Rule RI-1** If τ'' is μ-traversable and τ''' is not (if τ''' is μ-traversable and τ'' is not), then choose the embedding of $pert(\tau')$ to be $\Gamma^1(pert(\tau'))$ (resp. $\Gamma^2(pert(\tau'))$), see Fig. 16(a).
  - **Rule RI-2** If none of τ'' and τ''' is μ-traversable, if τ'' is μ-sided and τ''' is not (if τ''' is μ-sided and τ'' is not), then choose the embedding of $pert(\tau')$ to be $\Gamma^1(pert(\tau'))$ (resp. $\Gamma^2(pert(\tau'))$), see Fig. 16(b).

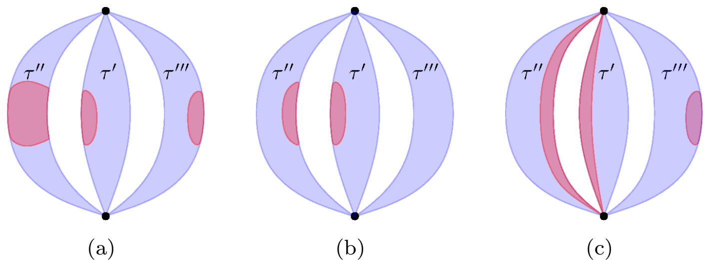

Figure 15: Choosing the embedding of $pert(\tau')$ in order to satisfy Conditions (i) and (ii) of Lemma 1 when τ is a P-node. (a) If τ'' is μ-traversable and τ''' is not, then the embedding of $pert(\tau')$ is $\Gamma^1(pert(\tau'))$; (b) if none of τ'' and τ''' is μ-traversable and if τ'' is μ-sided, then the embedding of $pert(\tau')$ is $\Gamma^1(pert(\tau'))$; (c) if τ'' is μ-spined, then the embedding of $pert(\tau')$ is $\Gamma^1(pert(\tau'))$.

- *Rules to satisfy Condition* (ii) *of Lemma 1.*

  Consider any neighbor τ' of τ in $\mathcal{T}$ such that $\Gamma^1(pert(\tau'))$ is μ-side-spined and not μ-full. Also, consider the neighbors τ'' and τ''' of τ in $\mathcal{T}$ following and preceding τ' in the circular order of the neighbors of τ determined by $\Gamma(sk(\tau))$, respectively. Without loss of generality assume that if the embedding of $pert(\tau')$ is $\Gamma^1(pert(\tau'))$, then the path $P(\mu)$ delimiting the outer face of $\Gamma^1(pert(\tau'))$ and composed only of vertices in μ is incident to the face of $\Gamma(sk(\tau))$ to which τ'' is incident; and if the embedding of $pert(\tau')$ is $\Gamma^2(pert(\tau'))$, then $P(\mu)$ is incident to the face of $\Gamma(sk(\tau))$ to which τ''' is incident.

  - **Rule RII-1** If τ'' is μ-spined and τ''' is not (resp. if τ''' is μ-spined and τ'' is not), and the face shared by τ' and τ'' in $\Gamma(sk(\tau))$ (τ' and τ''' in $\Gamma(sk(\tau))$) is different from the outer face of $\Gamma(sk(\tau))$, then choose the embedding of $pert(\tau')$ to be $\Gamma^1(pert(\tau'))$ (resp. $\Gamma^2(pert(\tau'))$), see Fig. 16(c).

- Suppose that the embedding of $pert(\tau')$, for some neighbor τ' of τ, has not been determined by rules RI-1, RI-2, and RII-1 over all clusters $\mu \in T$.

  - **Rule R0** If τ' is the parent of τ in $\mathcal{T}$, then set the embedding of $pert(\tau')$ to be $\Gamma^1(pert(\tau'))$ or $\Gamma^2(pert(\tau'))$ so that e_ρ is incident to the outer face. Otherwise, arbitrarily set the embedding of $pert(\tau')$ to be $\Gamma^1(pert(\tau'))$ or $\Gamma^2(pert(\tau'))$.

Denote by Γ the drawing constructed by the described algorithm. In order to complete the construction of Γ^C, we (possibly) modify the outer face of Γ. Namely, if e_ρ is incident to the outer face of Γ, then we let $\Gamma^C = \Gamma$. Otherwise, we choose as outer face of Γ^C the face of Γ that is incident to e_ρ and that is delimited by paths belonging to the pertinent graphs of different neighbors of τ in $\mathcal{T}$. Observe that one of such neighbors is the parent of τ in $\mathcal{T}$. Also, the choice of the outer face of Γ^C does not alter the embedding $\Gamma(sk(\tau))$ of $sk(\tau)$. Namely, a different outer face is chosen as outer face in $\Gamma(sk(\tau))$, however the circular ordering of the virtual edges around the poles stays the same; moreover, the virtual edge representing the parent of τ in $\mathcal{T}$ is still incident to the outer face of $\Gamma(sk(\tau))$ in Γ^C.

Second part. We now prove that the embedding choices performed when considering two distinct clusters do not conflict.

Claim 7. *For any two clusters μ and ν ($\mu \neq \nu$), the application of the above rules when μ is considered does not produce an embedding choice for $pert(\tau')$ that is conflicting with the one that is produced when ν is considered.*

Proof: The proof is independent of the modification of the outer face of Γ to obtain Γ^C, hence we will refer to drawing Γ rather than to Γ^C. The proof distinguishes several cases.

Case 1: Suppose that the embedding of $pert(\tau')$ has been determined to be $\Gamma^1(pert(\tau'))$ by Rule RI-1 because τ' is μ-sided, because τ'' is μ-traversable, and because τ''' is not μ-traversable, for some cluster $\mu \in T$, or by Rule RI-2 because τ' is μ-sided, because τ'' is μ-sided, and because τ''' is not μ-traversable and not μ-sided.

Case 1A: Suppose, for a contradiction, that the embedding of $pert(\tau')$ has been determined to be $\Gamma^2(pert(\tau'))$ by Rule RI-1 because τ' is ν-sided, because τ''' is ν-traversable, and because τ'' is not ν-traversable, or by Rule RI-2 because τ' is ν-sided, because τ''' is ν-sided, and because τ'' is not ν-traversable and not ν-sided, for some cluster $\nu \in T$ with $\nu \neq \mu$.

We have that Condition (i) of Lemma 1 is not satisfied by Γ', a contradiction. Namely, one of the two paths between the poles of τ' delimiting the outer face of $\Gamma^1(pert(\tau'))$, say $P(\mu, \nu)$, contains vertices in μ and in ν, while the other path does not. Hence, if Γ' is such that τ'' is found before τ''' when traversing the neighbors of τ in $\mathcal{T}$ starting from τ' in the direction "defined by $P(\mu, \nu)$", then $H(\nu)$ is not connected, given that τ' and τ'' are not ν-traversable. Otherwise $H(\mu)$ is not connected, given that τ' and τ''' are not μ-traversable.

Case 1B: Suppose, for a contradiction, that the embedding of $pert(\tau')$ has been determined by Rule RII-1 to be $\Gamma^2(pert(\tau'))$ because τ' is ν-side-spined and not ν-full, and because τ''' is ν-spined and τ'' is not, for some cluster $\nu \in T$ with $\nu \neq \mu$. Then one of the two paths between the poles of τ' delimiting the outer face of $\Gamma^1(pert(\tau'))$, say $P(\nu)$, entirely belongs to ν and the same path also contains a vertex in μ.

This gives rise to a contradiction if μ is not a descendant of ν in T. Hence, assume that μ is a descendant of ν in T. If Γ' is such that τ''' is found before τ'' when traversing the neighbors of τ in $\mathcal{T}$ starting from τ' in the direction "defined by $P(\nu)$", then $H(\mu)$ is not connected, given that τ' and τ''' are not μ-traversable, thus Condition (i) of Lemma 1 is not satisfied by Γ', a contradiction. Otherwise, there exists a cycle in Γ' whose vertices belong to ν and whose interior contains in Γ' a vertex not belonging to ν, thus implying that Condition (ii) of Lemma 1 is not satisfied by Γ', a contradiction. Namely, such a cycle is composed of $P(\nu)$ and of any path belonging to ν and connecting the poles of τ''' in $pert(\tau''')$; the vertex not in ν in the interior of this cycle is either a vertex in $pert(\tau')$ not in ν (which exists since τ' is not ν-full), or a vertex in $pert(\tau'')$ not in ν (which exists since τ'' is not ν-spined and hence not ν-full), depending on the "position" of the outer face in Γ'.

Case 2: Suppose that the embedding of $pert(\tau')$ has been determined to be $\Gamma^1(pert(\tau'))$ by Rule RII-1 because τ' is μ-side-spined and not μ-full, and because τ'' is μ-spined and τ''' is not, for some cluster $\mu \in T$.

Case 2A: Suppose that the embedding of $pert(\tau')$ has been determined to be $\Gamma^2(pert(\tau'))$ by Rule RI-1 (by Rule RI-2) because τ' is ν-sided, because τ''' is ν-traversable (ν-sided, respectively), and because τ'' is not ν-traversable (not ν-traversable and not ν-sided), for some cluster $\nu \in T$ with $\nu \neq \mu$. Such a case can be discussed analogously to Case 1B.

Case 2B: Suppose that the embedding of $pert(\tau')$ has been determined to be $\Gamma^2(pert(\tau'))$ by Rule RII-1 because τ' is ν-side-spined and not ν-full, and because τ''' is ν-spined and τ'' is not, for some cluster $\nu \in T$ with $\nu \neq \mu$. Observe that, since μ and ν share vertices, one of them is the ancestor of the other one. Assume without loss of generality that μ is an ancestor of ν. Since τ''' is ν-spined, it is also μ-spined, a contradiction.

This concludes the proof of the claim. $\square$

Third part. We now prove that the drawing Γ^C resulting from the above described algorithm is a $\langle 0,0,\infty \rangle$-drawing of $C(G,T)$ by exploiting Lemma 1. Observe that edge e_ρ is incident to the outer face of Γ^C by construction.

Claim 8. *Drawing Γ^C satisfies Condition (i) of Lemma 1.*

Proof: Consider any cluster μ. Observe that, by Lemma 1, $H(\mu)$ is connected in Γ' since Γ' is a $\langle 0,0,\infty \rangle$-drawing. We prove that $H(\mu)$ is connected in Γ^C. Observe that $H(\mu)$ is connected in Γ^C if and only if it is connected in Γ. Hence, we will refer to drawing Γ rather than to Γ^C. Suppose, for a contradiction, that $H(\mu)$ is not connected in Γ.

For any neighbor τ' of τ in $\mathcal{T}$, we can assume that $\Gamma^1(pert(\tau'))$ is neither μ-infeasible, nor μ-kernelized, nor μ-bisided. Namely:

- If $\Gamma^1(pert(\tau'))$ is μ-infeasible, we have that $H(\mu)$ is not connected in Γ', given that the embedding of $pert(\tau')$ in Γ' is the same as in Γ, up to a flip, thus leading to a contradiction.

- If $\Gamma^1(pert(\tau'))$ is μ-kernelized, then either there exists a neighbor of τ in $\mathcal{T}$ different from τ' containing a vertex in μ, thus implying that $H(\mu)$ is not connected in Γ', given that the embedding of $pert(\tau')$ in Γ' is the same as in Γ, up to a flip, or there exists no neighbor of τ in $\mathcal{T}$ different from τ' containing a vertex in μ, thus implying that $H(\mu)$ is connected in Γ; in both cases this leads to a contradiction.

- If $\Gamma^1(pert(\tau'))$ is μ-bisided, then, in order for $H(\mu)$ to be connected in Γ', we have that $\Gamma^1(pert(\tau''))$ is μ-traversable for each neighbor τ'' of τ in $\mathcal{T}$

different from τ'. Since the embedding of $pert(\tau'')$ in Γ' is the same as in Γ, up to a flip, it follows that τ'' is μ-traversable in Γ, as well, thus $H(\mu)$ is connected in Γ, thus leading to a contradiction.

Hence, we can assume that, for each neighbor τ' of τ in $\mathcal{T}$, $\Gamma^1(pert(\tau'))$ (and hence $\Gamma^2(pert(\tau'))$) is either μ-sided, or μ-traversable, or it contains no vertex of μ. Consider a maximal sequence $\tau_1, \tau_2, \ldots, \tau_k$ of neighbors of τ in $\mathcal{T}$, ordered as in $\Gamma(sk(\tau))$, such that τ_i contains a vertex in μ. We distinguish two cases.

- If there exists no neighbor τ' of τ in $\mathcal{T}$ such that $\Gamma^1(pert(\tau'))$ is μ-traversable, then, by Property (C) of $\Gamma(sk(\tau))$, there exist at most two neighbors τ_1 and τ_2 of τ in $\mathcal{T}$ such that $\Gamma^1(pert(\tau_1))$ and $\Gamma^1(pert(\tau_2))$ are μ-sided, and moreover τ_1 and τ_2 are adjacent in $\Gamma(sk(\tau))$. By construction, the embedding of $pert(\tau_1)$ is chosen to be $\Gamma^1(pert(\tau_1))$ or $\Gamma^2(pert(\tau_1))$ so that a vertex in $pert(\tau_1)$ and in μ is incident to the face of $\Gamma(sk(\tau))$ to which τ_2 is incident. Analogously, the embedding of $pert(\tau_2)$ is chosen to be $\Gamma^1(pert(\tau_2))$ or $\Gamma^2(pert(\tau_2))$ so that a vertex in $pert(\tau_2)$ and in μ is incident to the face of $\Gamma(sk(\tau))$ to which τ_1 is incident. Hence, the subgraph of $H(\mu)$ in Γ induced by the vertices in τ_1 and the subgraph of $H(\mu)$ in Γ induced by the vertices in τ_2 are connected by an edge; moreover, both such subgraphs are connected. It follows that $H(\mu)$ is connected in Γ, a contradiction.

- If there exists at least one neighbor of τ in $\mathcal{T}$ whose embedding in Γ' is μ-traversable, then, by Properties (B) and (C) of $\Gamma(sk(\tau))$, the order in $\Gamma(sk(\tau))$ of the neighbors of τ in $\mathcal{T}$ is $\tau_1, \tau_2, \ldots, \tau_k$, where $\tau_2, \ldots, \tau_{k-1}$ are μ-traversable, where τ_1 and τ_k are μ-sided (and they might not exist), and where $k \geq 3$. By construction, the embedding of $pert(\tau_1)$, if τ_1 exists, is chosen to be $\Gamma^1(pert(\tau_1))$ or $\Gamma^2(pert(\tau_1))$ so that a vertex in $pert(\tau_1)$ and in μ is incident to the face of $\Gamma(sk(\tau))$ to which τ_2 is incident. Analogously, the embedding of $pert(\tau_k)$, if τ_k exists, is chosen to be $\Gamma^1(pert(\tau_k))$ or $\Gamma^2(pert(\tau_k))$ so that a vertex in $pert(\tau_k)$ and in μ is incident to the face of $\Gamma(sk(\tau))$ to which τ_{k-1} is incident. Hence, the subgraph of $H(\mu)$ in Γ induced by the vertices in τ_1 is connected by an edge to the subgraph of $H(\mu)$ in Γ induced by the vertices in $\tau_2, \ldots, \tau_{k-1}$; moreover, the subgraph of $H(\mu)$ in Γ induced by the vertices in τ_k is connected by an edge to the subgraph of $H(\mu)$ in Γ induced by the vertices in $\tau_2, \ldots, \tau_{k-1}$; furthermore, these three subgraphs are connected. It follows that $H(\mu)$ is connected in Γ, a contradiction.

This concludes the proof of the claim. □

Claim 9. *Drawing Γ^C satisfies Condition (ii) of Lemma 1.*

Proof: We prove that Γ^C contains no cycle whose vertices all belong to the same cluster μ and whose interior contains a vertex v not in μ. We first prove this statement for drawing Γ, and we will later show how modifying the outer face of Γ to obtain Γ^C does not create a cycle whose vertices all belong to the same cluster μ and whose interior contains a vertex not in μ.

Suppose, for a contradiction, that Γ contains a cycle $\mathcal{C}$ whose vertices all belong to the same cluster μ and whose interior contains a vertex v not in μ. First, we discuss the existence of a cycle $\mathcal{C}$ violating Condition (ii) of Lemma 1 in the drawing constructed before.

If $\mathcal{C}$ entirely belongs to $pert(\tau')$, for some neighbor τ' of τ in $\mathcal{T}$, then $\mathcal{C}$ contains v in its interior also in Γ', given that the embedding of $pert(\tau')$ in Γ is the same as in Γ', up to a flip. However, by Lemma 1, this implies that Γ' is not a $\langle 0, 0, \infty\rangle$-drawing, a contradiction.

Otherwise, $\mathcal{C}$ is composed of two paths connecting the poles of τ, where the first path $P_\mu(\tau'')$ entirely belongs to $pert(\tau'')$ and the second path $P_\mu(\tau''')$ entirely belongs to $pert(\tau''')$, for two distinct neighbors τ'' and τ''' of τ in $\mathcal{T}$. Hence, $\Gamma^1(pert(\tau''))$ and $\Gamma^1(pert(\tau'''))$ (and hence $\Gamma^2(pert(\tau''))$ and $\Gamma^2(pert(\tau''')))$ are μ-spined. Moreover, they are both μ-side-spined (and possibly μ-full). Namely, if one of them, say $\Gamma^1(pert(\tau''))$, is μ-central-spined, then $\mathcal{C}$ contains a vertex not in μ in Γ', thus implying that Γ' does not satisfy Condition (ii) of Lemma 1 and hence that Γ' is not a $\langle 0, 0, \infty\rangle$-drawing, a contradiction.

Assume that $\Gamma^1(pert(\tau''))$ and $\Gamma^1(pert(\tau'''))$ (and hence $\Gamma^2(pert(\tau''))$ and $\Gamma^2(pert(\tau''')))$ are both μ-side-spined (possibly μ-full). By Property (A) of $\Gamma(sk(\tau))$, all the virtual edges, if any, that lie inside the cycle composed of the virtual edges representing τ'' and τ''' in $\Gamma(sk(\tau))$ are μ-full. It follows that, if $\mathcal{C}$ contains a vertex not in μ in its interior, then such a vertex belongs to $pert(\tau'')$ or to $pert(\tau''')$. Suppose the former, the discussion for the latter case being analogous. Observe that this implies that $\Gamma^1(pert(\tau''))$ is not μ-full. Then, denote by $Q_\mu(\tau'')$ the path that is composed only of vertices of μ, that connects the poles of τ, and that delimits the outer face of $\Gamma^1(pert(\tau''))$.

Consider the two neighbors σ_1 and σ_2 of τ in $\mathcal{T}$ such that the virtual edges representing σ_1 and σ_2 are consecutive with the virtual edge representing τ'' in $\Gamma(sk(\tau))$. Assume, w.l.o.g., that the virtual edge representing σ_1 is internal to the cycle composed of the virtual edges representing τ'' and τ''' in $\Gamma(sk(\tau))$ (if any virtual edge internal to such a cycle exists) or that σ_1 coincides with τ''' (otherwise).

If σ_2 is not μ-spined, then by Rule RII-1, the embedding of $pert(\tau'')$ is chosen in such a way that $Q_\mu(\tau'')$ is incident to the face of $\Gamma(sk(\tau))$ the virtual edge representing σ_1 is incident to. However, this implies that $\mathcal{C}$ does not contain any vertex not belonging to μ in its interior, a contradiction.

Otherwise, σ_2 is μ-spined, and Rule RII-1 was applied to choose the embedding of $pert(\tau'')$ in such a way that $Q_\mu(\tau'')$ is incident to the face of $\Gamma(sk(\tau))$ the virtual edge representing σ_2 is incident to.

It follows that the virtual edges representing σ_2 and τ'' delimit the outer face of $\Gamma(sk(\tau))$. Namely, if that's not the case, then the cycle that is composed of the virtual edges representing σ_2 and τ''' would contain the virtual edge representing τ'', which is not μ-full, in its interior in $\Gamma(sk(\tau))$, thus contradicting the fact that $\Gamma(sk(\tau))$ satisfies Property (A). However, that the virtual edges representing σ_2 and τ'' delimit the outer face of $\Gamma(sk(\tau))$ contradicts the assumption that Rule

RII-1 was applied to choose the embedding of $pert(\tau'')$ in such a way that $Q_\mu(\tau'')$ is incident to the face of $\Gamma(sk(\tau))$ the virtual edge representing σ_2 is incident to, thus obtaining a contradiction.

This completes the proof that Γ does not contain a cycle $\mathcal{C}$ whose vertices all belong to the same cluster μ and whose interior contains a vertex v not in μ.

Now we deal with Γ^C. If Γ^C has the same outer face as Γ, then there is nothing to prove. Otherwise, suppose, for a contradiction, that Γ^C contains a cycle $\mathcal{C}$ whose vertices all belong to the same cluster μ and whose interior contains a vertex v not in μ. Since Γ and Γ^C coincide when restricted to the pertinent graphs of the children of τ in $\mathcal{T}$, it follows that $\mathcal{C}$ is composed of a path $P_\mu(\sigma)$ in the pertinent graph of the parent σ of τ in $\mathcal{T}$ and of a path $P_\mu(\sigma')$ in the pertinent graph of a neighbor $\sigma' \neq \sigma$ of τ in $\mathcal{T}$. Also, since rule R0 sets the embedding of $pert(\sigma)$ so that e_ρ is incident to the outer face, it follows that either rule RI-1, or rule RI-2, or rule RII-1 was applied to determine the embedding of $pert(\sigma)$ to be either $\Gamma^1(pert(\sigma))$ or $\Gamma^2(pert(\sigma))$, so that e_ρ is not incident to the outer face of Γ. We distinguish two cases based on which rule was applied to determine the embedding of $pert(\sigma)$.

Suppose first that rule RII-1 was applied to determine the embedding of $pert(\sigma)$. By the assumptions of rule RII-1, $\Gamma^1(pert(\sigma))$ is ν-side-spined and not ν-full, for some cluster $\nu \in T$ (possibly $\nu = \mu$), and the embedding of the pertinent graph of a neighbor σ' of σ in $\Gamma(sk(\tau))$ is ν-spined. Also, since e_ρ is not incident to the outer face of Γ, it follows that the path $P_\nu(\sigma)$ in $pert(\sigma)$ that connects the poles of σ, that delimits the outer face of $\Gamma^1(pert(\sigma))$, and that entirely belongs to ν contains e_ρ. However, since the embedding of $pert(\sigma)$ in Γ^C coincides with the embedding of $pert(\sigma)$ in Γ' and since e_ρ is incident to the outer face of Γ', it follows that the cycle $\mathcal{C}'$ composed of $P_\nu(\sigma)$ and of any path in $pert(\sigma')$ that connects the poles of σ' and that entirely belongs to ν (such a path exists because σ' is ν-spined) contains in its interior in Γ' a vertex not belonging to ν (namely any vertex in $pert(\sigma)$ not in ν; such a vertex exists because $pert(\sigma)$ is not ν-full). Thus, Γ' does not satisfy Condition (ii) of Lemma 1, a contradiction.

Suppose next that rule RI-1 or rule RI-2 was applied to determine the embedding of $pert(\sigma)$. By the assumptions of rules RI-1 and RI-2, $\Gamma^1(pert(\sigma))$ is ν-sided, for some cluster $\nu \in T$, and the embedding of the pertinent graph of a neighbor σ' of σ in $\Gamma(sk(\tau))$ is ν-traversable or ν-sided.

- If $\Gamma^1(pert(\sigma))$ is μ-central-spined, then Γ and Γ' contain cycles that entirely belong to μ and that contain in their interior vertices not belonging to μ (namely vertices in $pert(\sigma)$), thus contradicting the fact that Γ and Γ' are $\langle 0, 0, \infty\rangle$-drawings of $C(G,T)$.
- If $\Gamma^1(pert(\sigma))$ is μ-side-spined and the path that connects the poles of σ, that delimits the outer face of $\Gamma^1(pert(\sigma))$, and that entirely belongs to μ contains e_ρ, then Γ' contains a cycle that entirely belongs to μ and that contains in its interior a vertex not belonging to μ (namely a vertex in $pert(\sigma)$), thus contradicting the fact that Γ' is a $\langle 0, 0, \infty\rangle$-drawing of $C(G,T)$.

- If $\Gamma^1(pert(\sigma))$ is μ-side-spined and the path that connects the poles of σ, that delimits the outer face of $\Gamma^1(pert(\sigma))$, and that entirely belongs to μ does not contain e_ρ, then Γ contains a cycle that entirely belongs to μ and that contains in its interior a vertex not belonging to μ (namely a vertex in $pert(\sigma)$), thus contradicting the fact that Γ is a $\langle 0, 0, \infty \rangle$-drawing of $C(G,T)$.

We can hence assume that $pert(\sigma)$ is μ-full. Denote by $\sigma_1, \sigma_2, \dots, \sigma_l$ the order of the children of τ as in $\Gamma(sk(\tau))$, where σ_1 is incident to the outer face of $\Gamma(sk(\tau))$ in Γ^C, while σ_l is incident to the outer face of $\Gamma(sk(\tau))$ in Γ.

If any μ-full child σ_j of τ in $\mathcal{T}$ exists, then all of $\sigma_1, \sigma_2, \dots, \sigma_j$ are μ-full, as otherwise Γ would not satisfy Condition (ii) of Lemma 1. Denote by f the largest index such that $\sigma_1, \sigma_2, \dots, \sigma_f$ are μ-full. Also, if any ν-traversable child σ_j of τ in $\mathcal{T}$ exists, then all of $\sigma_1, \sigma_2, \dots, \sigma_j$ are ν-traversable, as otherwise Γ would not satisfy Condition (i) of Lemma 1. Denote by t the largest index such that $\sigma_1, \sigma_2, \dots, \sigma_t$ are ν-traversable.

- If $t \geq f > 0$, then every μ-full child of τ is ν-traversable. All the μ-full children of τ occur consecutively in $\Gamma'(sk(\tau))$, as otherwise Γ' does not satisfy Condition (ii) of Lemma 1. Since $f > 0$, at least one μ-full and ν-traversable child of τ exists, and it is next to σ in $\Gamma'(sk(\tau))$. Hence, either a child of τ that is not ν-traversable exists, thus implying that $H(\nu)$ is not connected in Γ' and hence that Γ' does not satisfy Condition (i) of Lemma 1, a contradiction, or every child of τ is ν-traversable, thus implying that rules RI-1 and RI-2 were not applied to determine the embedding of $pert(\sigma)$, a contradiction.

- If $f \geq t > 0$, then every ν-traversable child of τ is μ-full. All the μ-full children of τ occur consecutively in $\Gamma'(sk(\tau))$, as otherwise Γ' does not satisfy Condition (ii) of Lemma 1. Since $t > 0$, at least one μ-full and ν-traversable child of τ exists. Hence, either a child of τ exists that is not ν-traversable and not μ-full exists, thus implying that $H(\nu)$ is not connected in Γ' and hence that Γ' does not satisfy Condition (i) of Lemma 1, or every child of τ is μ-full, thus contradicting the assumption that $\mathcal{C}$ contains a vertex not in μ in its interior.

- If $t = 0$, then σ_1 is ν-sided, as otherwise neither rule RI-1 nor rule RII-2 would be applied to determine the embedding of $pert(\sigma)$. It follows that σ_1 is incident to the outer face of $\Gamma'(sk(\tau))$. Thus, if σ_1 is μ-spined, then Γ' does not satisfy Condition (ii) of Lemma 1, a contradiction, while if σ_1 is not μ-spined (and some other child of τ is), then Γ does not satisfy Condition (ii) of Lemma 1, a contradiction.

- If $f = 0$, then there exists exactly one child of τ that is μ-spined. If σ_1 is not μ-spined, then Γ does not satisfy Condition (ii) of Lemma 1, a contradiction. Otherwise σ_1 is μ-spined. Hence, σ_1 is next to σ in $\Gamma'(sk(\tau))$. Thus, if σ_1 is ν-traversable, then either every child of τ is

ν-traversable, thus implying that rules RI-1 and RI-2 were not applied to determine the embedding of $pert(\sigma)$, a contradiction, or a child of τ that is not ν-traversable exists, thus implying that $H(\nu)$ is not connected in Γ' and hence that Γ' does not satisfy Condition (i) of Lemma 1, a contradiction. Finally, if σ_1 is ν-sided, then $H(\nu)$ is not connected either in Γ or in Γ', a contradiction.

This concludes the proof of the claim. $\square$

Claims 4–6 prove that Properties (A), (B), and (C) of Lemma 5 are necessary in order for an embedding $\Gamma(sk(\tau))$ in which the edge e representing the parent of τ is incident to the outer face to be extensible. The sufficiency of Properties (A), (B), and (C) of Lemma 5 is easily proved as described above if τ is an S-node or an R-node; Claims 7–9 prove the sufficiency of Properties (A), (B), and (C) of Lemma 5 if τ is a P-node. This completes the proof of Lemma 5.

We are now ready to give an algorithm for testing whether a given clustered graph admits a $\langle 0, 0, \infty \rangle$-drawing.

Theorem 4. *Let $C(G,T)$ be a clustered graph such that G is biconnected. There exists a polynomial-time algorithm to test whether $C(G,T)$ admits a $\langle 0, 0, \infty \rangle$-drawing.*

Proof: Let $\mathcal{T}$ be the SPQR-tree of G. Consider any Q-node ρ of $\mathcal{T}$ and root $\mathcal{T}$ at ρ. Such a choice corresponds to assuming that any considered planar embedding of G has the edge e_ρ of G corresponding to ρ incident to the outer face. In the following, we describe how to test whether $C(G,T)$ admits a $\langle 0, 0, \infty \rangle$-drawing under the above assumption. The repetition of such a test for all possible choices of ρ results in a test of whether $C(G,T)$ admits a $\langle 0, 0, \infty \rangle$-drawing.

First, we perform a preprocessing step to compute the following information.

For each node $\tau \in \mathcal{T}$ and for each cluster $\mu \in T$, we label τ (and the virtual edge representing τ in the skeleton of the parent of τ) with flags stating whether τ is (i) μ-touched, (ii) τ is μ-full, and (iii) τ is μ-spined.

Observe that such information can be easily computed in polynomial time based on whether any vertex of $pert(\tau)$ different from its poles belongs to μ, on whether all the vertices of $pert(\tau)$ belong to μ, and on whether there exists a path in $pert(\tau)$ connecting the poles of τ and entirely belonging to μ, respectively. In particular, such information does not change if the embedding of $pert(\tau)$ varies.

Also, for each node $\tau \in \mathcal{T}$ and for each cluster $\mu \in T$, we label τ (and the virtual edge representing τ in the skeleton of the parent of τ) with a flag stating whether τ is μ-traversable in any planar embedding that is not μ-infeasible. Namely, by Lemma 3, if an embedding of $pert(\tau)$ is μ-traversable, then every embedding of $pert(\tau)$ that is not μ-infeasible is μ-traversable. To compute this information, we traverse $\mathcal{T}$ bottom-up while computing, for each encountered node τ' of $\mathcal{T}$, whether $H(\tau', \mu)$ has a connected component that contains $f'(\tau')$ and $f''(\tau')$ in any planar embedding of $pert(\tau')$. In particular:

- If τ' is a Q-node, then we label τ' as μ-traversable if and only if one of the poles of τ' belongs to μ;

- if τ' is an S-node, then we label τ' as μ-traversable if and only if at least one of the children of τ' is labeled as μ-traversable;
- if τ' is a P-node, then we label τ' as μ-traversable if and only if every child of τ' is labeled as μ-traversable;
- if τ' is an R-node, then we label τ' as μ-traversable if and only if, in the unique (up to a flip) planar embedding $\Gamma(sk(\tau'))$ of $sk(\tau')$, there exists a sequence $(\tau_1, \tau_2, \ldots, \tau_x)$ of children of τ such that: (1) τ_j is a child of τ that is labeled as μ-traversable, for every $1 \leq j \leq x$, (2) the virtual edges representing τ_1 and τ_x are incident to the two faces to which the virtual edge representing the parent of τ is incident to, and (3) the virtual edges representing τ_j and τ_{j+1} are incident to a common face in $\Gamma(sk(\tau'))$, for every $1 \leq j \leq x-1$.

Observe that, whether $pert(\tau)$ is μ-sided, μ-bisided, μ-kernelized, μ-infeasible, μ-side-spined, or μ-central-spined depends on the actual embedding of $pert(\tau)$.

Second, we traverse $\mathcal{T}$ bottom-up. For every P-node and every R-node τ of $\mathcal{T}$, the *visible nodes* of τ are the children of τ that are not S-nodes plus the children of each S-node that is a child of τ. At each step, we consider either a P-node or an R-node τ with visible nodes $\tau_1, \ldots, \tau_k$. We inductively assume that, for each visible node τ_i, with $1 \leq i \leq k$, an extensible embedding $\Gamma(pert(\tau_i))$ has been computed, together with the information whether $\Gamma(pert(\tau_i))$ is μ-traversable, μ-sided, μ-bisided, μ-kernelized, μ-infeasible, μ-side-spined, or μ-central-spined, for each cluster μ in T.

We show how to test whether an extensible embedding $\Gamma(pert(\tau))$ of $pert(\tau)$ exists. Such a test consists of two phases. We first test whether $sk(\tau)$ admits an extensible embedding $\Gamma(sk(\tau))$. In the negative case, we can conclude that $C(G,T)$ has no $\langle 0, 0, \infty\rangle$-drawing in which e_ρ is incident to the outer face. In the positive case, we also test whether a flip of each $\Gamma(pert(\tau_i))$ exists such that the resulting embedding $\Gamma(pert(\tau))$ is extensible.

Extensible embedding of the skeleton of τ.

Suppose that τ is an R-node. Then, $sk(\tau)$ has a unique embedding, up to a flip. We hence test whether Properties (A), (B), and (C) of Lemma 5 are satisfied. Observe that such a test can be easily performed in polynomial time, based on the information on whether the visible nodes of μ (and hence the children of τ) are μ-spined, μ-full, μ-touched, and μ-traversable.

Suppose that τ is a P-node. We check whether there exists an extensible embedding $\Gamma(sk(\tau))$ of $sk(\tau)$ as follows. We impose constraints on the ordering of the virtual edges of τ.

A first set of constraints establishes that $\Gamma(sk(\tau))$ satisfies Property (A) of Lemma 5. Namely, for each cluster μ:

(a) We constrain all the μ-full virtual edges to be consecutive;

(b) if there exists no μ-full virtual edge, then we constrain each pair of μ-spined virtual edges to be consecutive; and

(c) if there exists at least one μ-full virtual edge, then, for each μ-spined virtual edge, we constrain such an edge and all the μ-full virtual edges to be consecutive.

A second set of constraints establishes that $\Gamma(sk(\tau))$ satisfies Properties (B) and (C) of Lemma 5. Namely, for each cluster μ:

(a) We constrain all the μ-traversable virtual edges to be consecutive;

(b) if there exists no μ-traversable virtual edge, then we constrain each pair of μ-touched virtual edges to be consecutive; and

(c) if there exists at least one μ-traversable virtual edge, then, for each μ-touched virtual edge, we constrain such an edge and all the μ-full virtual edges to be consecutive.

We check whether an ordering of the virtual edges of $sk(\tau)$ that enforces all these constraints exists by using the PQ-tree data structure [3]. If such an ordering does not exist, we conclude that $sk(\tau)$ admits no extensible embedding. Otherwise, we have an embedding $\Gamma(sk(\tau))$ of $sk(\tau)$ which satisfies Properties (A)–(C), hence it is extensible.

Extensible embedding of the pertinent graph of τ.

We now determine, for each node τ of $\mathcal{T}$ that is either an R- or a P-node, an extensible embedding $\Gamma(pert(\tau))$ of $pert(\tau)$, if one exists. This is done by choosing the flip of the embedding $\Gamma(pert(\tau_i))$ of each visible node τ_i of τ in such a way that $\Gamma(pert(\tau))$ satisfy Properties (i) and (ii) of Lemma 1. Observe that the choice of the flip of the embedding $\Gamma(pert(\tau_i))$ of each visible node τ_i of τ, together with the choice of the embedding of $sk(\tau)$ to be $\Gamma(sk(\tau))$, completely determines $\Gamma(pert(\tau))$.

We will construct a 2-SAT formula F such that $pert(\tau)$ admits an extensible embedding if and only if F is satisfiable. We initialize $F = \emptyset$. Then, for each visible node τ_i of τ, we assign an arbitrary flip to τ_i and define a boolean variable x_i that is positive if τ_i has the assigned flip and negative otherwise.

We first add some clauses to F in order to ensure that $\Gamma(pert(\tau))$ satisfies Property (ii) of Lemma 1.

For each cluster μ, we consider the embedded subgraph $\Gamma_\mu(sk(\tau))$ of $\Gamma(sk(\tau))$ containing all the μ-spined virtual edges. Note that, since $\Gamma(sk(\tau))$ satisfies Property (A) of Lemma 5, each edge of $\Gamma_\mu(sk(\tau))$ that is not incident to the outer face of $\Gamma_\mu(sk(\tau))$ is μ-full. Hence, no flip choice has to be done for these edges.

Consider any edge g in $\Gamma_\mu(sk(\tau))$ such that g is incident to the external face and to an internal face f_g of $\Gamma_\mu(sk(\tau))$. Consider each visible node τ_i of τ such that: (i) either g corresponds to τ_i or g corresponds to the S-node which is the parent of τ_i and (ii) $\Gamma(pert(\tau_i))$ is either μ-side-spined or μ-central-spined.

Then, if $\Gamma(pert(\tau_i))$ is μ-central-spined, then we conclude that $pert(\tau)$ has no extensible embedding (with e incident to the outer face). Otherwise, $\Gamma(pert(\tau_i))$ is μ-side-spined. In this case, add clause $\{x_i\}$ to F if the default flip of $\Gamma(pert(\tau_i))$

does not place any vertex not in μ on f_g, and add clause $\{\neg x_i\}$ to F if the default flip of $\Gamma(pert(\tau_i))$ places a vertex not in μ on f_g.

We next add some clauses to F in order to ensure that $\Gamma(pert(\tau))$ satisfy Property (i) of Lemma 1.

Suppose that there exists a visible node τ_i of τ such that: (i) $\Gamma(pert(\tau_i))$ is μ-bisided; and (ii) if τ_i is child of an S-node σ, then no child of σ is μ-traversable.

Then, we check:

(a) Whether all the visible nodes of τ that are children of τ, except for τ_i, are μ-traversable; and

(b) whether, for each S-node γ that is child of τ and that is not the parent of τ_i, at least one child of γ is μ-traversable.

If one of the checks fails, we conclude that $pert(\tau)$ has no extensible embedding (with e incident to the outer face), otherwise we do not add any clause to F and we continue as follows.

Suppose that there exists a visible node τ_i of τ such that:

(i) $\Gamma(pert(\tau_i))$ is μ-sided; (ii) if τ_i is child of an S-node σ, then no child of σ is μ-traversable; and (iii) τ_i shares exactly one face with a μ-traversable or μ-sided node τ_j.

Then, add clause $\{x_i\}$ to F if the default flip of $\Gamma(pert(\tau_i))$ places the vertices of $pert(\tau_i)$ belonging to μ on the face that τ_i shares with τ_j, and add clause $\{\neg x_i\}$ to F otherwise.

Suppose that there exists an S-node σ child of τ such that:

(i) no child of σ is μ-traversable; and (ii) no child of τ different from σ is μ-touched.

Consider each pair of visible nodes τ_i and τ_j children of σ that are both μ-sided.

Then, add clauses $(x_i \vee \neg x_j)$ and $(\neg x_i \vee x_j)$ to F if the default flips of $\Gamma(pert(\tau_i))$ and of $\Gamma(pert(\tau_j))$ place vertices of $pert(\tau_i)$ and vertices of $pert(\tau_j)$ belonging to μ on the same face. Otherwise, add clauses $(x_i \vee x_j)$ and $(\neg x_i \vee \neg x_j)$ to F.

Observe that all the described checks and embedding choices, and the construction and solution of the 2-SAT formula can be easily performed in polynomial time. Once an embedding $\Gamma(pert(\tau))$ of $pert(\tau)$ has been computed, by traversing $\Gamma(pert(\tau))$ it can be determined in polynomial time whether such an embedding is μ-sided, μ-bisided, μ-kernelized, μ-side-spined, or μ-central-spined. By Lemma 4, if a $\langle 0, 0, \infty \rangle$-drawing of $C(G,T)$ exists, then there exists a $\langle 0, 0, \infty \rangle$-drawing of $C(G,T)$ in which the embedding of $pert(\tau)$ is $\Gamma(pert(\tau))$ or its flip. This allows the bottom-up visit of $\mathcal{T}$ to go through. The correctness of the embedding choices performed in order to construct $\Gamma(pert(\tau))$ follows from Lemmata 1, 4, and 5. This concludes the proof of the theorem. □

We now turn our attention to establish bounds on the minimum value of γ in a $\langle 0, 0, \gamma \rangle$-drawing of a clustered graph.

Theorem 5. *Let $C(G,T)$ be a clustered graph. There exists an algorithm to compute a $\langle 0,0,\gamma\rangle$-drawing of $C(G,T)$ with $\gamma \in O(n^3)$, if any such drawing exists. If $C(G,T)$ is flat, then $\gamma \in O(n^2)$.*

Proof: Suppose that $C(G,T)$ admits a $\langle 0,0,\infty\rangle$-drawing. Then, consider the drawing Γ of the underlying graph G in any such a drawing. For each cluster μ, place a vertex $u_{\mu,f}$ inside any face f of Γ that contains at least one vertex belonging to μ, and connect $u_{\mu,f}$ to all the vertices of μ incident to f. Note that, the graph composed by the vertices of μ and by the added vertices u_{μ,f_i} is connected. Then, construct a spanning tree of such a graph and draw $R(\mu)$ slightly surrounding such a spanning tree. The cubic bound on γ comes from the fact that each of the $O(n)$ clusters crosses each of the $O(n)$ other clusters a linear number of times. On the other hand, if $C(G,T)$ is flat, then each of the $O(n)$ clusters crosses each of the $O(n)$ other clusters just once. □

4. Lower bounds

In this section we give lower bounds on the number of *ee*-, *er*-, and *rr*-crossings in $\langle\alpha,\beta,\gamma\rangle$-drawings of clustered graphs.

First, we prove an auxiliary lemma concerning the crossing number in graphs without a cluster hierarchy. Given a graph G, we define $G(m)$ as the multigraph obtained by replacing each edge of G with a set of m multiple edges. For each pair (u,v) of vertices, we denote by $S(u,v)$ the set of multiple edges connecting u and v.

Lemma 6. *Graph $G(m)$ has crossing number $cr(G(m)) \geq m^2 \cdot cr(G)$.*

Proof: Consider a drawing Γ of $G(m)$ with the minimum number $cr(G(m))$ of crossings.

First, observe that in Γ no edge intersects itself, no two edges between the same pair of vertices intersect, and each pair of edges crosses at most once. Namely, if any of these conditions does not hold, it is easy to modify Γ to obtain another drawing of $G(m)$ with a smaller number of crossings, which is not possible by hypothesis (see, e.g., [28]).

We show that there exists a drawing Γ' of $G(m)$ with $cr(G(m))$ crossings in which, for each pair of vertices u and v, all the edges between u and v cross the same set of edges in the same order. Let $e_{min}(u,v)$ be any edge with the minimum number of crossings among the edges of $S(u,v)$. Redraw all the edges in $S(u,v) \setminus e_{min}(u,v)$ so that they intersect the same set of edges as $e_{min}(u,v)$, in the same order as $e_{min}(u,v)$. Repeating this operation for each set $S(u,v)$ yields a drawing Γ' with the required property.

Starting from Γ', we construct a drawing Γ_G of G. For each set of edges $S(u,v)$ remove all edges except for one edge $e^*(u,v)$. The resulting drawing Γ_G of G has at least $cr(G)$ crossings, by definition. For any two edges $e^*(u,v)$ and $e^*(w,z)$ that cross in Γ_G, we have that each edge in $S(u,v)$ crosses each edge in

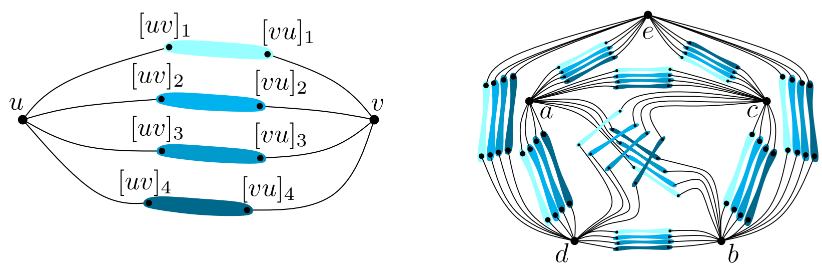

(a)

(b)

Figure 16: Illustrations for the proof of Theorem 6. (a) Edges and clusters in $M(u,v)$. (b) Clustered graph $C(G,T)$.

$S(w,z)$, by the properties of Γ'. Hence Γ' contains at least $m^2 \cdot cr(G)$ crossings. $\square$

We prove a lower bound on the total number of crossings in an $\langle\alpha,\beta,\gamma\rangle$-drawing of a clustered graph when all the three types of crossings are admitted.

Theorem 6. *There exists an n-vertex non-c-connected flat clustered graph $C(G,T)$ that admits $\langle\infty,0,0\rangle$-, $\langle 0,\infty,0\rangle$-, and $\langle 0,0,\infty\rangle$-drawings, and such that $\alpha+\beta+\gamma \in \Omega(n^2)$ in every $\langle\alpha,\beta,\gamma\rangle$-drawing of $C(G,T)$.*

Proof: Clustered graph $C(G,T)$ is as follows. Initialize graph G with five vertices a,b,c,d,e. For each two vertices $u,v \in \{a,b,c,d,e\}$, with $u \neq v$, and for $i = 1,\dots,m$, add to G vertices $[uv]_i$, $[vu]_i$, and edges $(u,[uv]_i)$ and $(v,[vu]_i)$, and add to T a cluster $\mu(u,v)_i = \{[uv]_i,[vu]_i\}$. Vertices a,b,c,d,e belong to clusters $\mu_a,\mu_b,\mu_c,\mu_d,\mu_e$, respectively. See Fig. 16. We denote by $M(u,v) = \{(u,[uv]_i),(v,[vu]_i),\mu(u,v)_i | i = 1,\dots,m\}$.

First, we prove that C admits $\langle\infty,0,0\rangle$-, $\langle 0,\infty,0\rangle$-, and $\langle 0,0,\infty\rangle$-drawings. Consider a drawing Γ^* of K_5 on vertices $\{a,b,c,d,e\}$ with one crossing. Assume, without loss of generality, that the crossing is on (a,b) and (c,d). For each pair of vertices $u,v \in \{a,b,c,d,e\}$, with $(u,v) \notin \{(a,b),(c,d)\}$, replace edge (u,v) in Γ^* with $M(u,v)$ in such a way that the drawing of the edges and clusters in $M(u,v)$ is arbitrarily close to the drawing of (u,v). It remains to draw edges and clusters in $M(a,b)$ and $M(c,d)$. This is done differently for $\langle\infty,0,0\rangle$-, $\langle 0,\infty,0\rangle$-, and $\langle 0,0,\infty\rangle$-drawings.

$\langle\infty,0,0\rangle$**-drawing** Replace (a,b) and (c,d) in Γ^* with $M(a,b)$ and $M(c,d)$ in such a way that the drawing of the edges and clusters in $M(a,b)$ (in $M(c,d)$) is arbitrarily close to the drawing of (a,b) (of (c,d)) and for each $1 \leq i,j \leq m$, edge $(a,[ab]_i)$ crosses edge $(c,[cd]_j)$, while edges $(b,[ba]_i)$ and $(d,[dc]_i)$, and regions $R(\mu(a,b)_i)$ and $R(\mu(c,d)_j)$ are not involved in any crossing. See Fig. 18(a).

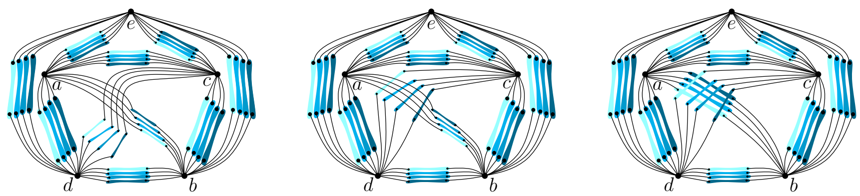

(a) (b) (c)

Figure 17: (a) $\langle \infty, 0, 0 \rangle$-drawing of C. (b) $\langle 0, \infty, 0 \rangle$-drawing of C. (c) $\langle 0, 0, \infty \rangle$-drawing of C.

$\langle 0, \infty, 0 \rangle$-drawing Replace (a, b) and (c, d) in Γ^* with $M(a, b)$ and $M(c, d)$ in such a way that the drawing of the edges and clusters in $M(a, b)$ (in $M(c, d)$) is arbitrarily close to the drawing of (a, b) (of (c, d)) and for each $1 \leq i, j \leq m$, edge $(a, [ab]_i)$ crosses region $R(\mu(c, d)_j)$, while edges $(b, [ba]_i)$, $(c, [cd]_i)$, and $(d, [dc]_j)$, and region $R(\mu(a, b)_i)$ are not involved in any crossing. See Fig. 18(b).

$\langle 0, 0, \infty \rangle$-drawing Replace (a, b) and (c, d) in Γ^* with $M(a, b)$ and $M(c, d)$ in such a way that the drawing of the edges and clusters in $M(a, b)$ (in $M(c, d)$) is arbitrarily close to the drawing of (a, b) (of (c, d)) and for each $1 \leq i, j \leq m$, region $R(\mu(a, b)_i)$ crosses region $R(\mu(c, d)_j)$, while edges $(a, [ab]_i)$, $(b, [ba]_i)$, $(c, [cd]_j)$, and $(d, [dc]_i)$ are not involved in any crossing. See Fig. 18(c).

Second, we show that $\alpha + \beta + \gamma \in \Omega(n^2)$ in every $\langle \alpha, \beta, \gamma \rangle$-drawing of $C(G, T)$. Consider any such a drawing Γ. Starting from Γ, we obtain a drawing Γ' of a subdivision of $K_5(m)$ as follows. For each $u, v \in \{a, b, c, d, e\}$, with $u \neq v$, and for each $i = 1, \ldots, m$, insert a drawing of edge $([uv]_i, [vu]_i)$ inside $R(\mu(u, v)_i)$ and remove region $R(\mu(u, v)_i)$. Further, remove regions $R(\mu_a)$, $R(\mu_b)$, $R(\mu_c)$, $R(\mu_d)$, and $R(\mu_e)$. The obtained graph is a subdivision of $K_5(m)$. Hence, by Lemma 6, Γ' has $\Omega(n^2)$ crossings. Moreover, each crossing in Γ' corresponds either to an edge-edge crossing, or to an edge-region crossing, or to a region-region crossing in Γ, thus proving the theorem. $\square$

We now turn our attention to drawings in which only one type of crossings is allowed. In this setting, we show that the majority of the upper bounds presented in the previous section are tight by giving lower bounds on the number of crossings of $\langle \infty, 0, 0 \rangle$-, $\langle 0, \infty, 0 \rangle$-, and $\langle 0, 0, \infty \rangle$-drawings. As a corollary of Theorem 6, there exists a clustered graph $C(G, T)$ such that $\alpha \in \Omega(n^2)$ in every $\langle \alpha, 0, 0 \rangle$-drawing of $C(G, T)$, such that $\beta \in \Omega(n^2)$ in every $\langle 0, \beta, 0 \rangle$-drawing of $C(G, T)$, and such that $\gamma \in \Omega(n^2)$ in every $\langle 0, 0, \gamma \rangle$-drawing of $C(G, T)$. However, in the following we present quadratic lower bounds on restricted classes of clustered graphs and a cubic lower bound for $\langle 0, 0, \infty \rangle$-drawings of clustered graphs.

We first consider $\langle \infty, 0, 0 \rangle$-drawings. We give two lower bounds, which deal with c-connected and non-c-connected clustered graphs, respectively.

Theorem 7. *There exists a c-connected flat clustered graph $C(G,T)$ such that $\alpha \in \Omega(n^2)$ in every $\langle \alpha, 0, 0 \rangle$-drawing of $C(G,T)$.*

Proof: We first describe $C(G,T)$. Graph G is a subdivision of $K_5(m)$, with $m = \frac{n-5}{9}$, where the set of edges $S(d,e)$ has been removed. Tree T is such that $\mu_2 = \{d\}$, $\mu_3 = \{e\}$, and all the other vertices belong to μ_1. See Fig. 19(a). Since, in any $\langle \infty, 0, 0 \rangle$-drawing Γ of $C(G,T)$, both d and e must be outside any cycle composed of vertices of μ_1 (as otherwise they would lie inside $R(\mu_1)$, see Fig. 19(b)), a set of m length-2 paths can be drawn in Γ between d and e without creating other crossings, thus obtaining a drawing of a subdivision of $K_5(m)$ in which the crossings are the same as in Γ (see Fig. 19(c)). Since $cr(K_5) = 1$ and since $m = \Omega(n)$, by Lemma 6, $\alpha \in \Omega(n^2)$. $\square$

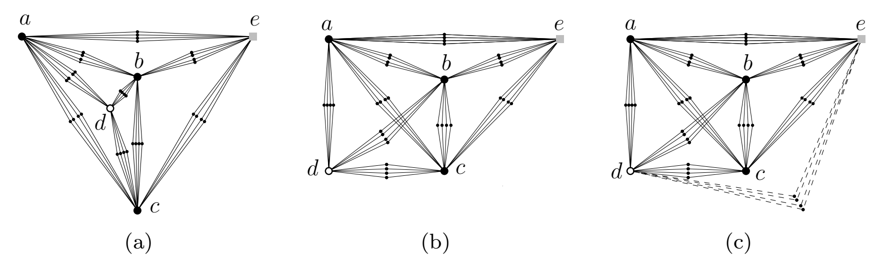

Figure 18: Illustration for the proof of Theorem 7. Vertices in μ_1 are black, vertices in μ_2 are white, and vertices in μ_3 are gray. (a) Graph G. (b) Vertices d and e must be outside all the cycles composed of vertices of μ_1. (c) Graph G'', where the length-2 paths connecting d and e are dashed.

Theorem 8. *There exists a non-c-connected flat clustered graph $C(G,T)$, where G is a matching, such that $\alpha \in \Omega(n^2)$ in every $\langle \alpha, 0, 0 \rangle$-drawing of $C(G,T)$.*

Proof: Clustered graph $C(G,T)$ is constructed as follows. Tree T is a star graph with five leaves $\mu_1, \dots, \mu_5$. For each $i \neq j$ with $1 \leq i, j \leq 5$, add $\frac{n}{20}$ vertices to μ_i and to μ_j and construct a matching between these two sets of vertices.

Consider any $\langle \alpha, 0, 0 \rangle$-drawing Γ of $C(G,T)$ such that α is minimum. We prove that Γ does not contain any edge-edge crossing inside the regions representing clusters. This implies that, contracting the regions to single points yields a drawing of a subdivision of $K_5(n/20)$, and Lemma 6 applies to obtain the claimed lower bound for α.

Assume, for a contradiction, that a crossing between two edges e_1 and e_2 occurs inside the region $R(\mu)$ representing a cluster μ. Since Γ has no edge-region crossings, both e_1 and e_2 connect a vertex in μ with a vertex not in μ. Then, one might place the endvertex of e_1 belonging to μ arbitrarily close to the boundary of $R(\mu)$ in such a way that it does not cross e_2 inside $R(\mu)$. Since this operation reduces the number of crossings, we have a contradiction to the fact that α is minimum.

Then, we add a vertex to each cluster μ_i and connect it to all the vertices of μ_i. Observe that, since no two edges cross inside the region representing a cluster, such vertices and edges can be added without creating any new crossings.

Finally, removing from Γ the drawings of the regions representing the clusters leads to a drawing of a subdivision of $K_5(n/20)$ with α crossings. By Lemma 6, $\alpha \in \Omega(n^2)$. □

We now prove some lower bounds on the number of *er*-crossings in $\langle 0, \infty, 0\rangle$-drawings of clustered graphs. In the case of non-c-connected flat clustered graphs, a quadratic lower bound directly follows from Theorem 6, as stated in the following.

Corollary 1. *There exists a non-c-connected flat clustered graph $C(G,T)$ such that $\beta \in \Omega(n^2)$ in every $\langle 0, \beta, 0\rangle$-drawing of $C(G,T)$.*

Next, we deal with the c-connected case and present a quadratic and a linear lower bound for non-flat and flat cluster hierarchies, respectively.

Theorem 9. *There exists a c-connected non-flat clustered graph $C(G,T)$ such that $\beta \in \Omega(n^2)$ in every $\langle 0, \beta, 0\rangle$-drawing of $C(G,T)$.*

Proof: Let G be an $(n+2)$-vertex triconnected planar graph such that for $i = 1, \dots, \frac{n}{3}$, G contains a 3-cycle $C_i = (a_i, b_i, c_i)$. Further, for $i = 1, \dots, \frac{n}{3} - 1$, G has edges (a_i, a_{i+1}), (b_i, b_{i+1}), (c_i, c_{i+1}). Finally, G contains two vertices v_a and v_b such that v_a is connected to a_1, b_1, c_1 and v_b is connected to $a_{\frac{n}{3}}, b_{\frac{n}{3}}, c_{\frac{n}{3}}$. Tree T is defined as follows: $\mu_1 = \{a_1, b_1, c_1\}$ and, for each $i = 2, \dots, \frac{n}{3}$, $\mu_i = \mu_{i-1} \cup \{a_i, b_i, c_i\}$; moreover $\mu_a = \{v_a\}$ and $\mu_b = \{v_b\}$. See Fig. 20(a).

Note that, in any planar embedding of G, there exists a set S of at least $\frac{n}{6}$ nested 3-cycles, and all such cycles contain either v_a or v_b, say v_b, in their interior. Let C_i be any of such cycles. For each cluster μ containing a_i, b_i, and c_i, not all the edges of C_i can be drawn entirely inside the region $R(\mu)$ representing μ in any $\langle 0, \infty, 0\rangle$-drawing of $C(G,T)$, as otherwise $R(\mu)$ would enclose v_b. This implies that C_i intersects the border of $R(\mu)$ twice, hence creating an edge-region crossing. Since there exist $\Omega(n)$ cycles in S, each of which is contained in $\Omega(n)$ clusters, we have that any $\langle 0, \beta, 0\rangle$-drawing of $C(G,T)$ has $\beta = \Omega(n^2)$ edge-region crossings. □

Theorem 10. *There exists a c-connected flat clustered graph $C(G,T)$ such that $\beta \in \Omega(n)$ in every $\langle 0, \beta, 0\rangle$-drawing of $C(G,T)$.*

Proof: The underlying graph G is defined as in the proof of Theorem 9. Tree T is such that, for $i = 1, \dots, n$, there exists a cluster μ_i containing vertices a_i, b_i, and c_i; moreover, $\mu_a = \{v_a\}$ and $\mu_b = \{v_b\}$. See Fig. 20(b).

In any planar embedding of G there exists a set S of at least $\frac{n}{6}$ nested 3-cycles, and all such cycles contain either v_a or v_b, say v_b, in their interior. Let C_i be any of such cycles. Not all the edges of C_i can be drawn entirely inside the region $R(\mu_i)$ representing μ_i in any $\langle 0, \infty, 0\rangle$-drawing of $C(G,T)$, as otherwise

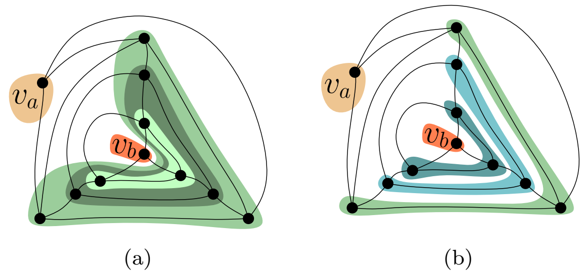

Figure 19: (a) Illustration for Theorem 9. (b) Illustration for Theorem 10.

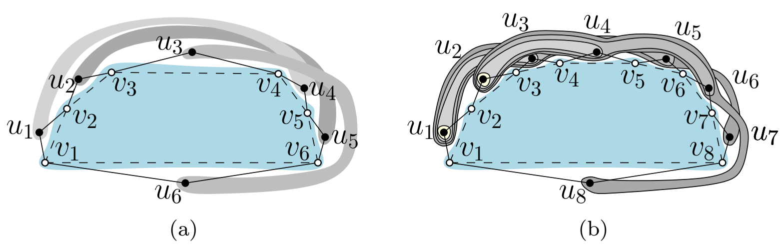

Figure 20: (a) Illustration for Theorem 11. (b) Illustration for Theorem 12.

$R(\mu_i)$ would enclose v_b. This implies that C_i intersects the border of $R(\mu_i)$ twice. Since there exist $\Omega(n)$ cycles in S, we have that any $\langle 0, \beta, 0\rangle$-drawing of $C(G,T)$ has $\beta = \Omega(n)$ edge-region crossings. □

Finally, we prove lower bounds on the number of rr-crossings in $\langle 0, 0, \infty\rangle$-drawings of clustered graphs. We only consider non-c-connected clustered graphs, since a c-connected clustered graph either does not admit any $\langle 0, 0, \infty\rangle$-drawing or is c-planar. We distinguish two cases based on whether the considered clustered graphs are flat or not.

Theorem 11. *There exists a non-c-connected flat clustered graph $C(G,T)$, where G is outerplanar, such that $\gamma \in \Omega(n^2)$ in every $\langle 0, 0, \gamma\rangle$-drawing of $C(G,T)$.*

Proof: We first describe $C(G,T)$. Refer to Fig. 21(a). Consider a cycle $\mathcal{C}$ of n vertices $v_1, \ldots, v_n$, such that n is even. For $i = 1, \ldots, n$, add to $\mathcal{C}$ a vertex u_i and connect it to v_i and v_{i+1}, where $v_{n+1} = v_1$. Denote by G the resulting outerplanar graph. Tree T is such that vertices $v_1, \ldots, v_n$ belong to the same cluster μ^* and, for $i = 1, \ldots, n/2$, vertices u_i and $u_{n/2+i}$ belong to μ_i.

Since all vertices $u_1, \ldots, u_n$ have to lie outside region $R(\mu^*)$ in any $\langle 0, 0, \infty\rangle$-drawing of $C(G,T)$, the embedding of G is outerplanar. Hence, for any $i \neq j \in \{1, \ldots, n/2\}$, cluster μ_i intersects cluster μ_j, thus proving the theorem. □

Theorem 12. *There exists a non-c-connected non-flat clustered graph $C(G,T)$, where G is outerplanar, such that $\gamma \in \Omega(n^3)$ in every $\langle 0, 0, \gamma\rangle$-drawing of $C(G,T)$.*

Proof: We first describe $C(G,T)$. Refer to Fig. 21(b). Graph G is an outerplanar graph constructed as in the proof of Theorem 11, such that n is a multiple of 4. Tree T is defined as follows. Set $\mu_1 = \{u_1\}$ and $\mu_2 = \{u_2\}$. Then, for each $i = 3, 4, \ldots, n$, set $\mu_i = \mu_{i-2} \cup \{u_i\}$. Finally, set $\mu^* = \{v_1, \ldots, v_n\}$.

Since all vertices $u_1, \ldots, u_n$ have to lie outside region $R(\mu^*)$, in any $\langle 0, 0, \infty \rangle$-drawing of $C(G,T)$ the embedding of G is outerplanar.

We claim that, for each $i \in \{\frac{n}{2}, \frac{n}{2}+2, \ldots, n\}$ and $j \in \{\frac{n}{2}+1, \frac{n}{2}+3, \ldots, n-1\}$, the border of region $R(\mu_i)$ intersects $\Omega(n)$ times the border of region $R(\mu_j)$. Observe that the claim implies the theorem.

We prove the claim. Consider the border $B(\mu_i)$ of $R(\mu_i)$, for any $i \in \{\frac{n}{2}, \frac{n}{2}+2, \ldots, n\}$. First, for each $2 \leq k \leq \frac{n}{2}$ such that k is even, $B(\mu_i)$ properly crosses edge (v_k, u_k) in a point p_k and edge (v_{k+1}, u_k) in a point p'_k, given that μ_i contains u_k and does not contain v_k and v_{k+1}. Second, for each $1 \leq h \leq \frac{n}{2}$ such that h is odd, $B(\mu_i)$ does not cross edges (v_h, u_h), given that μ_i contains neither u_h, nor v_h, nor v_{k+1}. Third, the intersection point of $B(\mu_i)$ with G that comes after p_k and p'_k is p_{k+2}, as otherwise $B(\mu_i)$ would not be a simple curve or an *er*-crossing would occur. Analogous considerations hold for each $j \in \{\frac{n}{2}+1, \frac{n}{2}+3, \ldots, n-1\}$. Hence, the part of $B(\mu_i)$ between p'_k and p_{k+2} not containing p_k intersects the part of $B(\mu_j)$ between p'_{k+1} and p_{k+3}. This concludes the proof of the theorem. □

5. Relationships between α, β and γ

In this section we discuss the interplay between *ee*-, *er*-, and *rr*-crossings for the realizability of $\langle \alpha, \beta, \gamma \rangle$-drawings of clustered graphs.

As a first observation in this direction, we note that the result proved in Theorem 6 shows that there exist *c*-graphs for which allowing *ee*-, *er*-, and *rr*-crossing at the same time does not reduce the total number of crossings with respect to allowing only one type of crossings.

Next, we study the following question: suppose that a clustered graph $C(G,T)$ admits a $\langle 1, 0, 0 \rangle$-drawing (resp. a $\langle 0, 1, 0 \rangle$-drawing, resp. a $\langle 0, 0, 1 \rangle$-drawing); does this imply that $C(G,T)$ admits a $\langle 0, \beta, 0 \rangle$-drawing and a $\langle 0, 0, \gamma \rangle$-drawing (resp. an $\langle \alpha, 0, 0 \rangle$-drawing and a $\langle 0, 0, \gamma \rangle$-drawing, resp. an $\langle \alpha, 0, 0 \rangle$-drawing and a $\langle 0, \beta, 0 \rangle$-drawing) with small number of crossings?

In the following, we prove that the answer to this question is often negative, as we can only prove (Theorem 13) that every graph admitting a drawing with one single *er*-crossing also admits a drawing with $O(n)$ *ee*-crossings, while in many other cases we can prove (Theorem 14) the existence of graphs that, even admitting a drawing with one single crossing of one type, require up to a quadratic number of crossings of a different type.

We first present Theorem 13. Observe that this theorem gives a stronger result than the one needed to answer the above question, as it proves that every $\langle \alpha, \beta, \gamma \rangle$-drawing of a clustered graph can be transformed into a $\langle \alpha + \beta \cdot O(n), 0, \gamma \rangle$-drawing.

Theorem 13. *Any n-vertex clustered graph admitting a $\langle 0, \beta, 0\rangle$-drawing also admits an $\langle \alpha, 0, 0\rangle$-drawing with $\alpha \in O(\beta n)$.*

Proof: Let Γ be a $\langle 0, \beta, 0\rangle$-drawing of a clustered graph $C(G,T)$. We construct an $\langle \alpha, 0, 0\rangle$-drawing of $C(G,T)$ with $\alpha \in O(\beta n)$ by modifying Γ as follows. For each cluster $\mu \in T$, consider the set of edges that cross the boundary of $R(\mu)$ at least twice. Partition this set into two sets E_{in} and E_{out} as follows. Each edge whose endvertices both belong to μ is in E_{in}; each edge none of whose endvertices belongs to μ is in E_{out}; all the other edges are arbitrarily placed either in E_{in} or in E_{out}. Fig. 22(a) represents a cluster μ and the corresponding set E_{out}. We describe the construction for E_{out}. For each edge $e \in E_{out}$ consider the set of curves obtained as $e \cap R(\mu)$, except for the curves having the endvertices of e as endpoints. Consider the set $\mathcal{S}$ that is the union of the sets of curves obtained from all the edges of E_{out}. Starting from any point of the boundary of $R(\mu)$, follow such a boundary in clockwise direction and assign increasing integer labels to the endpoints of all the curves in $\mathcal{S}$. See Fig. 22(b). Consider a curve $\zeta \in \mathcal{S}$ such that there exists no other curve $\zeta' \in \mathcal{S}$ whose both endpoints have a label that is between the labels of the two endpoints of ζ. Then, consider the edge e such that ζ is a portion of e. Consider two points p_1 and p_2 of e arbitrarily close to the two endpoints of ζ, respectively, and not contained into $R(\mu)$. Redraw the portion of e between p_1 and p_2 as a curve outside $R(\mu)$ following clockwise the boundary $B(\zeta, \mu)$ of $R(\mu)$ between the smallest and the largest endpoint of ζ, and arbitrarily close to $B(\zeta, \mu)$ in such a way that it crosses only the edges that cross $B(\zeta, \mu)$ and the edges that used to cross the portion of e between p_1 and p_2 before redrawing it. See Fig. 22(c), where the curve ζ between 6 and 8 is redrawn. Remove ζ from $\mathcal{S}$ and repeat this procedure until $\mathcal{S}$ is empty. Fig. 22(d) shows the final drawing obtained by applying the described procedure to the drawing in Fig. 22(a).

The construction for E_{in} is analogous, with the portion of e being redrawn inside $R(\mu)$. Observe that, every time the portion of an edge e between p_1 and p_2, corresponding to a curve $\zeta \in \mathcal{S}$, is redrawn, an *er*-crossing is removed from the drawing and at most $O(n)$ *ee*-crossings between e and the edges crossing $B(\zeta, \mu)$ are added to the drawing. This concludes the proof of the theorem. $\square$

In Theorem 14 we prove that there exist graphs that, even admitting a drawing with one single crossing of one type, require up to a quadratic number of crossings of a different type.

Theorem 14. *There exist clustered graphs C_1, C_2, and C_3 such that:*

(i) C_1 *admits a* $\langle 1, 0, 0\rangle$*-drawing,* $\beta \in \Omega(n^2)$ *in every* $\langle 0, \beta, 0\rangle$*-drawing of* C_1*, and* $\gamma \in \Omega(n^2)$ *in every* $\langle 0, 0, \gamma\rangle$*-drawing of* C_1*;*

(ii) C_2 *admits a* $\langle 0, 1, 0\rangle$*-drawing,* $\alpha \in \Omega(n)$ *in every* $\langle \alpha, 0, 0\rangle$*-drawing of* C_2*, and* $\gamma \in \Omega(n^2)$ *in every* $\langle 0, 0, \gamma\rangle$*-drawing of* C_2*;*

(iii) C_3 *admits* $\langle 0, 0, 1\rangle$*-drawing,* $\alpha \in \Omega(n^2)$ *in every* $\langle \alpha, 0, 0\rangle$*-drawing, and* $\beta \in \Omega(n)$ *in every* $\langle 0, \beta, 0\rangle$*-drawing of* C_3*.*

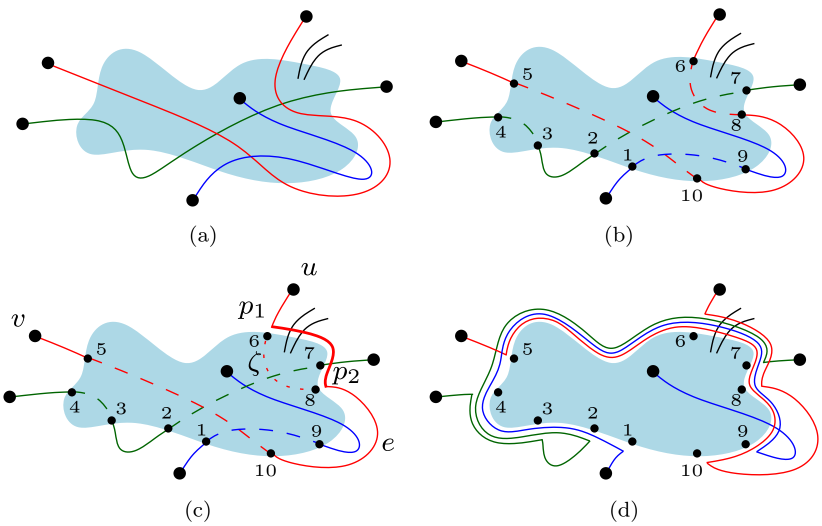

Figure 21: Illustration for Theorem 13. (a) A cluster μ with a set of edges crossing $R(\mu)$ at least twice and belonging to E_{out}. (b) The curves belonging to $\mathcal{S}$ are represented by dashed curve segments, while the other portions of the edges are represented by solid curve segments. The intersection points between curves in $\mathcal{S}$ and $R(\mu)$ are labeled with increasing integers. (c) The curve ζ between intersection points 6 and 8 that is a portion of edge $e = (u, v)$ is selected, since there exists no curve $\zeta' \in \mathcal{S}$ whose both endpoints have a label that is between 6 and 8. The old drawing of curve ζ is represented by a dotted curve segment, while the new drawing of ζ is represented by a fat solid curve. Note that the new drawing of ζ crosses all the edges that cross the boundary of $R(\mu)$ between 6 and 8. (d) The final drawing obtained by applying the described procedure to all the curves in $\mathcal{S}$.

Proof: We start by describing a clustered graph $C^*(G^*, T^*)$, that will be used as a template for the graphs in the proof. Graph G^* is obtained as follows.Refer to Fig. 23(a). Initialize $G^* = K_5(n)$ on vertices $\{a, b, c, d, e\}$. First, for each $u, v \in \{a, b, c, d, e\}$, with $u \neq v$, replace the set of n multiple edges $S(u, v)$ with a set $S'(u, v)$ of n length-2 paths between u and v. Then, remove from G^* sets $S'(a, d)$, $S'(c, e)$, $S'(a, e)$, and $S'(c, d)$. Finally, for $i = 1, \dots, n$, add to G^* vertices $[ae]_i, [ea]_i, [cd]_i, [dc]_i$, and edges $(a, [ae]_i)$, $(e, [ea]_i)$, $(c, [cd]_i)$, $(d, [dc]_i)$. For $i = 1, \dots, m$, T^* contains clusters $\mu(a, e)_i = \{[ae]_i, [ea]_i\}$ and $\mu(c, d)_i = \{[cd]_i, [dc]_i\}$. Denote by $M(a, e) = \{(a, [ae]_i), (e, [ea]_i), \mu(a, e)_i | i = 1, \dots, n\}$ and $M(c, d) = \{(c, [cd]_i), (d, [dc]_i), \mu(c, d)_i | i = 1, \dots, n\}$.

Clustered graph $C_1(G_1, T_1)$ is obtained by adding edges (a, d) and (c, e) to G^* and by setting $T_1 = T^*$. A $\langle 1, 0, 0\rangle$-drawing of C_1 is depicted in Fig. 23(b), where edges (a, d) and (c, e) cross.

Clustered graph $C_2(G_2, T_2)$ is obtained by adding edge (c, e) to G^* and by adding a cluster $\mu(a, d) = \{a, d\}$ to T^*. A $\langle 0, 1, 0\rangle$-drawing of C_2 is depicted in Fig. 23(c), where edge (c, e) and region $R(\mu(a, d))$ cross.

Clustered graph $C_3(G_3, T_3)$ is obtained by setting $G_3 = G^*$ and by adding clusters $\mu(a, d) = \{a, d\}$ and $\mu(c, e) = \{c, e\}$ to T^*. A $\langle 0, 0, 1\rangle$-drawing of C_3 is

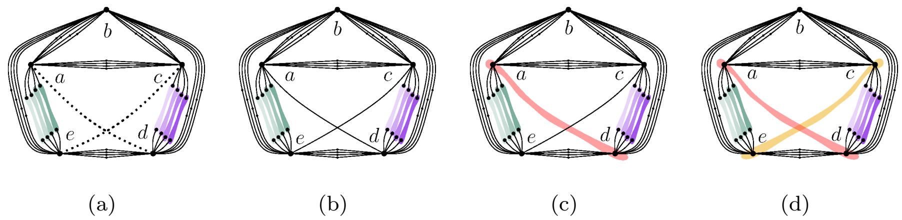

Figure 22: Illustration for the proof of Theorem 14. (a) C-graph C^*: Dotted lines are placeholders for gadgets. (b) A $\langle 1,0,0\rangle$-drawing of C_1, (c) a $\langle 0,1,0\rangle$-drawing of C_2, and (d) a $\langle 0,0,1\rangle$-drawing of C_3.

depicted in Fig. 23(d), where regions $R(\mu(a,d))$ and $R(\mu(c,e))$ cross.

The lower bounds claimed in the theorem can be obtained with arguments analogous to those used in the proof of Theorem 6.

For example, consider any $\langle 0,\beta,0\rangle$-drawing Γ_β of C_1 which minimizes β. For $i = 1,\dots,n$, draw an edge $([ae]_i,[ea_i])$ inside region $R(\mu(a,e)_i)$ and an edge $([cd]_i,[dc_i])$ inside region $R(\mu(c,d)_i)$, and remove such regions. Then, remove edges (a,d) and (c,e) and draw two sets $S(a,d)$ and $S(c,e)$ of n multiple edges arbitrarily close to the drawings of (a,d) and (c,e). The obtained drawing Γ'_β is a drawing of a subdivision of $K_5(n)$, and hence contain $\Omega(n^2)$ crossings. Since Γ_β is a $\langle 0,\infty,0\rangle$-drawing, each crossing in Γ'_β involves exactly one edge in $\{([ae]_i,[ea_i]),([cd]_i,[dc_i])\}$. Also, since Γ_β minimizes β, edges in $\{(a,d),(c,e)\}$ are not involved in any crossing in Γ'_β, since both such edges are adjacent to an edge belonging to $M(a,e)$ and to an edge belonging to $M(c,d)$ (recall that adjacent edges do not cross in any drawing of a graph whose number of crossings is minimum). Thus, each crossing in Γ'_β corresponds to an er-crossing in Γ_β, which implies that $\beta \in \Omega(n^2)$.

This concludes the proof of the theorem. □

6. Complexity

In this section we study the problem of minimizing the number of crossings in $\langle\alpha,\beta,\gamma\rangle$-drawings.

We define the problem (α,β,γ)-ClusterCrossingNumber ((α,β,γ)-CCN) as follows. Given a clustered graph $C(G,T)$ and an integer $k > 0$, problem (α,β,γ)-CCN asks whether $C(G,T)$ admits a $\langle\alpha,\beta,\gamma\rangle$-drawing with $\alpha+\beta+\gamma \leq k$.

First, we prove that problem (α,β,γ)-CCN belongs to class NP.

Lemma 7. *Problem (α,β,γ)-CCN is in NP.*

Proof: Similarly to the proof that the CrossingNumber problem is in NP [20], we need to "guess" a drawing of $C(G,T)$ with α ee-crossings, with β er-crossings, and with γ rr-crossings, for each choices of α, β, and γ satisfying $\alpha+\beta+\gamma \leq k$.

This is done as follows. Let m be the number of edge-cluster pairs $\langle e, \mu \rangle$ such that one end-vertex of e is in μ and the other one is not. Let $0 \leq p \leq \gamma$ be a guess on the number of pairs of clusters that intersect each other. Let $\mathcal{E}$ be a guess on the rotation schemes of the vertices of G. Arbitrarily orient each edge in G; also, arbitrarily fix a "starting point" on the boundary of each cluster in T and orient such a boundary in any way.

For each edge e, guess a sequence of crossings $x_1, x_2, \dots, x_{k(e)}$ occurring along e while traversing it according to its orientation. Each of such crossings x_i is associated with: (1) the edge e' that crosses e in x_i or the cluster μ' such that the boundary of $R(\mu')$ crosses e in x_i; and (2) a boolean value $b(x_i)$ stating whether e' (resp. the boundary of $R(\mu')$) crosses e from left to right according to the orientations of e and e' (resp. of e and the boundary of $R(\mu')$).

Analogously, for each cluster μ, guess a sequence of crossings $x_1, x_2, \dots, x_{k(\mu)}$ occurring along the boundary of $R(\mu)$ while traversing it from its starting point according to its orientation. Again, each of such crossings x_i is associated with: (1) the edge e' that crosses the boundary of $R(\mu)$ in x_i or the cluster μ' such that the boundary of $R(\mu')$ crosses the boundary of $R(\mu)$ in x_i; and (2) a boolean value $b(x_i)$ stating whether e' (resp. the boundary of $R(\mu')$) crosses the boundary of $R(\mu)$ from left to right according to the orientations of the boundary of $R(\mu)$ and e' (resp. of the boundary of $R(\mu)$ and the boundary of $R(\mu')$). Observe that the guessed crossings respect constraints $\mathcal{C}_1$, $\mathcal{C}_2$, and $\mathcal{C}_3$.

Crossings are guessed in such a way that there is a total number of α crossings between edge-edge pairs, a total number of $2\beta + 2m$ crossings between edge-cluster pairs, and a total number of $2\gamma + 2p$ crossings between cluster-cluster pairs, so that p pairs of clusters have a crossing.

We construct a graph G^* with a fixed rotation scheme around each vertex as follows. Start with G^* having the same vertex set of G and containing no edge. For each edge e in G, add to G^* a path starting at one end-vertex of e, ending at the other end-vertex of e, and containing a vertex for each crossing associated with e. For each cluster μ in T, add to G^* a cycle containing a vertex for each crossing associated with μ. This is done in such a way that one single vertex is introduced in G^* for each guessed crossing. The rotation scheme of each vertex in G^* that is also a vertex in G is the one in $\mathcal{E}$. The rotation scheme of each vertex in G^* corresponding to a crossing x_i is determined according to $b(x_i)$.

Check in linear time whether the constructed graph G^* with a fixed rotation scheme around each vertex is planar. For each cluster μ, check in linear time whether the cycle representing the boundary of $R(\mu)$ contains in its interior all and only the vertices of G and the clusters in T (that is, all the vertices of the cycles representing such clusters) it has to contain. Observe that, if the checks succeed and a planar drawing of G^* with the a fixed rotation scheme around each vertex can be constructed, the corresponding drawing of $C(G, T)$ is an $\langle \alpha, \beta, \gamma \rangle$-drawing. □

Second, we prove that (α, β, γ)-CCN is NP-complete, even if the underlying graph is planar, namely a forest of star graphs, by means of a reduction from the CrossingNumber problem.

Theorem 15. *Problem (α, β, γ)-CCN is NP-complete, even in the case in which the underlying graph is a forest of star graphs.*

Proof: The membership in NP is proved in Lemma 7.

The NP-hardness is proved by means of a polynomial-time reduction from the CrossingNumber problem, which has been proved to be NP-complete by Garey and Johnson [20]. Given a graph G^* and an integer $k^* > 0$, the CrossingNumber problem consists of deciding whether G^* admits a drawing with at most k^* crossings.

We describe how to construct an instance $\langle C(G,T), k\rangle$ of (α, β, γ)-CCN starting from an instance $\langle G^*, k^*\rangle$ of CrossingNumber.

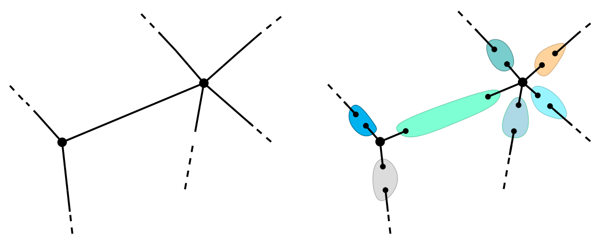

(a) (b)

Figure 23: Illustration for the proof of Theorem 15: A part of graph G^* (a) and the corresponding part of $C(G,T)$ (b).

Initialize $G = G^*$ and $T = \rho$. For each edge (u_i, u_j) of G^*, subdivide (u_i, u_j) with two subdivision vertices u'_i and u'_j, add a cluster $\mu_{i,j}$ to T containing u'_i and u'_j as a child of ρ, and remove edge (u'_i, u'_j) from G and from G^*. See Fig.23. Further, set $k = k^*$. Note that, graph G is a forest of star graphs. Also, instance $\langle$C(G,T) $, k\rangle$ can be constructed in polynomial time.

We show that instance $\langle C(G,T), k\rangle$ has a solution if and only if instance $\langle G^*, k^*\rangle$ has a solution. Both directions of the proof use techniques similar to those used in Section 4.

Suppose that $\langle G^*, k^*\rangle$ admits a solution, that is, G^* has a drawing Γ^* with at most k^* crossings. An $\langle \alpha, \beta, \gamma\rangle$-drawing Γ of $C(G,T)$ with $\alpha + \beta + \gamma \leq k$ can be constructed by subdividing twice each edge (u_i, u_j) of G^* and by replacing the central edge (u'_i, u'_j) of each length-3 path representing an edge (u_i, u_j) of G^* with a cluster whose drawing is arbitrarily close to the drawing of (u'_i, u'_j). By construction, each crossing between two edges in Γ^* corresponds to either an *ee*-crossing, or to an *er*-crossing, or to a *rr*-crossing in Γ. Hence, drawing Γ contains the same number of crossings as Γ^*, that is, at most $k^* = k$.

Suppose that $\langle C(G,T), k\rangle$ admits a solution, that is, $C(G,T)$ has an $\langle \alpha, \beta, \gamma\rangle$-drawing Γ with $\alpha + \beta + \gamma \leq k$. A drawing Γ^* of G^* with at most k^* crossings can be constructed by replacing each cluster $\mu_{i,j} = \{u'_i, u'_j\}$ with an edge between u'_i and u'_j inside $R(\mu_{i,j})$ and by replacing each length-3 path $P(i,j) = (u_i, u'_i, u'_j, u_j)$ with an edge (u_i, u_j) whose drawing is the same as the drawing of $P(i,j)$ in Γ. By construction, each crossing (that is either an *ee*-crossing, or an er-crossing, or an rr-crossing) in Γ corresponds to a crossing between two edges in Γ^*. Hence, drawing Γ^* contains the same number of crossings as Γ, that is, at most $k = k^*$.

This concludes the proof of the theorem. □

As for the problems considered in the previous sections, it is interesting to study the (α, β, γ)-CCN problem when only one out of α, β, and γ is allowed to be different from 0. We call α-CCN, β-CCN, and γ-CCN the corresponding decision problems.

We observe that the proof of Theorem 15 can be easily modified to show that all of α-CCN, β-CCN, and γ-CCN are NP-complete, even in the case in which the underlying graph is a forest of star graphs.

In the following we prove that stronger results can be found for α-CCN and β-CCN, by giving NP-hardness proofs for more restricted clustered graph classes.

Theorem 16. *Problem α-CCN is NP-complete even in the case in which the underlying graph is a matching.*

Proof: The membership in NP follows from Lemma 7.

The NP-hardness is proved by means of a polynomial-time reduction from the known CrossingNumber problem [20].

We describe how to construct an instance $\langle$C(G,T) $, k\rangle$ of α-CCN starting from an instance $\langle G^*, k^*\rangle$ of CrossingNumber. See Figs. 25(a)-(b).

Initialize $G = G^*$ and $T = \rho$. Subdivide each edge of G with two subdivision vertices. For each vertex v_i of G, add a cluster μ_i to T as a child of ρ containing all the neighbors of v_i, and remove from G vertex v_i and its incident edges.

Further, set $k = k^*$. Note that graph G is a matching. Also, instance $\langle C(G, T), k\rangle$ can be constructed in polynomial time.

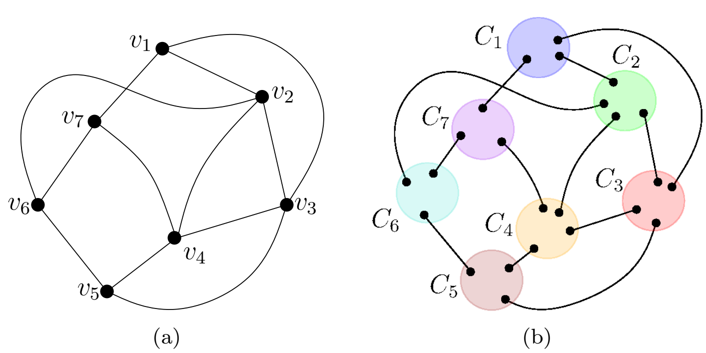

Figure 24: (a) Graph G^* in the proof of Theorem 16. (b) The clustered graph $C(G, T)$ corresponding to G^*.

We show that instance $\langle C(G, T), k\rangle$ has a solution if and only if instance $\langle G^*, k^*\rangle$ has a solution.

Suppose that $\langle G^*, k^* \rangle$ admits a solution, that is, G^* has a drawing Γ^* with at most k^* crossings. An $\langle \alpha, 0, 0 \rangle$-drawing Γ of $C(G,T)$ with $\alpha \leq k$ can be constructed as follows. Initialize $\Gamma = \Gamma^*$. For each vertex v_i of G^*, draw a small disk d_i around it, and, for each edge e incident to v_i, place a vertex v'_i on the intersection between e and the boundary of d_i. Then, replace each edge $e = (v_i, v_j)$ in Γ with edge (v'_i, v'_j). Finally, remove each vertex v_i of G^* from Γ and represent each cluster μ_i in Γ as a region slightly surrounding disk d_i. Since each edge in Γ is represented as a Jordan curve that is a subset of the Jordan curve representing an edge in Γ^*, drawing Γ contains at most the same number of crossings as Γ^*.

Suppose that $\langle C(G,T), k \rangle$ admits a solution, that is, $C(G,T)$ has an $\langle \alpha, 0, 0 \rangle$-drawing Γ with $\alpha \leq k$. A drawing Γ^* of G^* with at most k^* crossings can be constructed as follows. Place vertex v_i on any interior point of region $R(\mu_i)$. For each intersection point p between the boundary of $R(\mu_i)$ and an edge incident to μ_i, draw a curve connecting v_i to p so that such curves do not cross each other. Remove each vertex v'_i and, for each edge e incident to v'_i, the part of e which lies inside $R(\mu_i)$. Also, remove all the regions representing clusters of T. The crossings in the resulting drawing Γ^* of G^* are a subset of the *ee*-crossings in Γ. Namely, the curves that exist in Γ^* and do not exist in Γ do not cross any edge of G^*, given that Γ has no *er*-crossing. This concludes the proof of the theorem. □

Theorem 17. *Problem β-CCN is NP-complete even for c-connected flat clustered graphs in which the underlying graph is a triconnected planar multigraph.*

Proof: The membership in NP follows from Lemma 7.

The NP-hardness is proved by means of a polynomial-time reduction from the NP-complete problem SteinerTreePlanarGraphs (STPG) [19], which is defined as follows: Given a planar graph $G(V,E)$ whose edges have weights $w : E \to \mathbb{N}$, given a set $S \subset V$ of *terminals*, and given an integer k, does a tree $T^*(V^*, E^*)$ exist such that (1) $V^* \subseteq V$, (2) $E^* \subseteq E$, (3) $S \subseteq V^*$, and (4) $\sum_{e \in E^*} w(e) \leq k$? The edge weights in w are bounded by a polynomial function $p(n)$ (see [19]). We are going to use the variant of STPG in which (A) G is a subdivision of a triconnected planar graph, where each subdivision vertex is not a terminal, and (B) all the edge weights are equal to 1.

In the following we sketch a reduction from STPG to STPG with the described properties. Let G be any edge-weighted planar graph. Augment G to any triconnected planar graph $G'(V', E')$ by adding dummy edges and by assigning weight $w(e) = n \cdot p(n)$ to each dummy edge e. Then, replace each edge e of G' with a path $P(e)$ with $w(e)$ edges, each with weight 1, hence obtaining a planar graph $G''(V'', E'')$. Let the terminals of G'' be the same terminals of G. Note that, by construction, G'' satisfies Properties (A) and (B). Also, since $|V''| \in O(n^2 \cdot p(n))$, the described reduction is polynomial.

We prove that $\langle G, S, k \rangle$ is a positive instance of STPG if and only if $\langle G'', S, k \rangle$ is a positive instance of the considered variant of STPG.

Suppose that $\langle G, S, k\rangle$ is a positive instance of STPG. Starting from the solution T^* of $\langle G, S, k\rangle$, we construct a solution $T^\diamond$ of $\langle G'', S, k\rangle$ by replacing each edge e with path $P(e)$. By construction, $\sum_{e\in T^\diamond} w(e) = \sum_{e\in T^*} w(e) \leq k$.

Suppose that $\langle G'', S, k\rangle$ is a positive instance of the variant of STPG. Let $T^\diamond(V^\diamond, E^\diamond)$ be the solution of $\langle G'', S, k\rangle$. Assume that $T^\diamond$ is the *optimal* solution to $\langle G'', S, k\rangle$, i.e., there exists no tree $T^\sharp(V^\sharp, E^\sharp)$ such that $T^\sharp(V^\sharp, E^\sharp)$ is a solution to $\langle G'', S, k\rangle$ and $\sum_{e\in E^\sharp} w(e) < \sum_{e\in E^\diamond} w(e)$. Observe that, if an edge of a path $P(e)$ belongs to $E^\diamond$, then all the edges of $P(e)$ belong to $E^\diamond$. Moreover, no edge of a path $P(e)$ such that e is a dummy edge belongs to $E^\diamond$, since the total weight of the edges of each path $P(e)$ such that e is a dummy edge is $n \cdot p(n)$. Starting from $T^\diamond$, we construct a solution T^* of $\langle G, S, k\rangle$ by replacing all the edges of each path $P(e)$ with an edge e. By construction, $\sum_{e\in T^*} w(e) = \sum_{e\in T^\diamond} w(e) \leq k$.

Next we show a polynomial-time reduction from the variant of STPG in which all the instances satisfy Properties (A) and (B) to β-CCN. Refer to Fig. 25. Let $\langle G, S, k\rangle$ be an instance of the variant of STPG. Since G is a subdivision of a triconnected planar graph, it admits a unique planar embedding, up to a flip and to the choice of the outer face. Construct a planar embedding Γ_G of G. Construct the dual graph H of Γ_G. Note that, since G is a subdivision of a triconnected planar graph, its dual H is a planar triconnected multigraph. For each terminal $s \in S$, consider the set $E_G(s)$ of the edges incident to s in G and consider the face f_s of H composed of the edges that are dual to the edges in $E_G(s)$; add s to the vertex set of H, embed it inside f_s, and connect it to the vertices incident to f_s. Define the inclusion tree T as follows. For each vertex $s_i \in S$, with $1 \leq i \leq |S|$, T has a cluster $\mu_i = \{s_i\}$; all the other vertices in the vertex set of H belong to the same cluster ν. Then, the instance of β-CCN is $\langle C(H, T), k\rangle$.

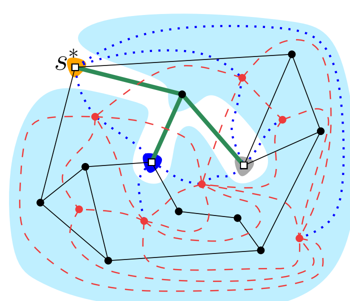

Figure 25: Illustration for the proof of Theorem 17. Solid (black) lines are edges of G; dashed (red) and dotted (blue) lines are edges of H; green edges are the edges of T^*; black circles and white squares are non-terminal vertices and terminals in G, respectively; finally, red circles and white squares are vertices in H.

We show that $\langle C(H, T), k\rangle$ admits a solution if and only if $\langle G, S, k\rangle$ does.

Suppose that $\langle G, S, k\rangle$ admits a solution T^*. Consider a terminal vertex $s^* \in S$ and construct a planar embedding of H such that s^* is incident to the outer face. Construct a drawing of cluster ν as a simple region $R(\nu)$ that entirely encloses H, except for a small region surrounding T^* (observe that such a simple region $R(\nu)$ exists since s^* is incident to the outer face and $s^* \in T^*$). Draw each cluster μ_i as a region $R(\mu_i)$ surrounding s_i sufficiently small so that it does not intersect $R(\nu)$. Observe that the resulting drawing of $C(H,T)$ is a $\langle 0, \infty, 0 \rangle$-drawing. Moreover, $R(\nu)$ intersects all and only the edges dual to edges in T^*, hence there are at most k edge-region crossings, that is, $C(H,T)$ is a $\langle 0, \beta, 0 \rangle$-drawing with $\beta \leq k$.

Suppose that $C(H,T)$ admits a $\langle 0, \beta, 0 \rangle$-drawing Γ with $\beta \leq k$ edge-region crossings and assume that Γ is optimal (that is, there is no $\langle 0, \infty, 0 \rangle$-drawing with fewer er-crossings). Consider the graph T^* composed of the edges that are dual to the edges of H participating in some edge-region crossing. We claim that T^* has at least one edge incident to each terminal in S and that T^* is connected. The claim implies that T^* is a solution to the instance $\langle G, S, k \rangle$ of STPG, since T^* has at most k edges and since Γ is optimal. Consider any terminal $s \in S$. If none of the edges incident to s in G belongs to T^*, it follows that none of the edges of H incident to face f_s has a crossing with the region $R(\nu)$ representing ν in Γ. Observe that, since Γ is optimal, there exists a terminal s^* incident to the outer face of Γ. If $s \neq s^*$, then since all the vertices incident to f_s have to lie inside $R(\nu)$, either $R(\nu)$ is not a simple region or it contains s, in both cases contradicting the assumption that Γ is a $\langle 0, \infty, 0 \rangle$-drawing. Also, if $s = s^*$, then either $R(\nu)$ is not a simple region or it contains all the vertices of H, and hence also vertices not in ν, in both cases contradicting the assumption that Γ is a $\langle 0, \infty, 0 \rangle$-drawing. Suppose that T^* contains (at least) two connected components T^*_1 and T^*_2. At least one of them, say T^*_2, does not contain any edge that is dual to an edge incident to the outer face of Γ. Therefore, $R(\nu)$ is not a simple region, a contradiction. This concludes the proof of the theorem. □

7. Open Problems

Given a clustered graph whose underlying graph is planar we defined and studied its $\langle \alpha, \beta, \gamma \rangle$-drawings, where the number of ee-, er-, and rr-crossings is equal to α, β, and γ, respectively.

This paper opens several problems. First, some of them are identified by non-tight bounds in the tables of the Introduction. Second, in order to study how allowing different types of crossings impacts the features of the drawings, we concentrated most of the attention on $\langle \alpha, \beta, \gamma \rangle$-drawings where two out of α, β, and γ are equal to zero. It would be interesting to study classes of clustered graphs that have drawings where the values of α, β, and γ are balanced in some way. Third, we have seen that not all clustered graphs whose underlying graph is planar admit $\langle 0, 0, \infty \rangle$-drawings. It would be interesting to characterize the class of clustered graphs that admit one and to extend our testing algorithm to simply-connected clustered graphs.

References

[1] P. Angelini, G. Di Battista, F. Frati, M. Patrignani, and I. Rutter. Testing the simultaneous embeddability of two graphs whose intersection is a biconnected or a connected graph. *J. Discrete Algorithms*, 14:150–172, 2012.

[2] P. Angelini, F. Frati, and M. Patrignani. Splitting clusters to get c-planarity. In *GD*, volume 5849 of *LNCS*, pages 57–68, 2010.

[3] K. S. Booth and G. S. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using pq-tree algorithms. *J. Comput. Syst. Sci.*, 13(3):335–379, 1976.

[4] P. Braß, E. Cenek, C. A. Duncan, A. Efrat, C. Erten, D. Ismailescu, S. G. Kobourov, A. Lubiw, and J. S. B. Mitchell. On simultaneous planar graph embeddings. *Comput. Geom.*, 36(2):117–130, 2007.

[5] S. Cornelsen and D. Wagner. Completely connected clustered graphs. *Journal of Discrete Algorithms*, 4(2):313–323, 2006.

[6] P. F. Cortese and G. Di Battista. Clustered planarity (invited lecture). In *Twenty-first annual symposium on Computational Geometry (proc. SoCG 05)*, ACM, pages 30–32, 2005.

[7] P. F. Cortese, G. Di Battista, F. Frati, M. Patrignani, and M. Pizzonia. C-planarity of c-connected clustered graphs. *J. Graph Alg. Appl.*, 12(2):225–262, 2008.

[8] P. F. Cortese, G. Di Battista, M. Patrignani, and M. Pizzonia. Clustering cycles into cycles of clusters. *J. Graph Alg. Appl.*, 9(3):391–413, 2005.

[9] P. F. Cortese, G. Di Battista, M. Patrignani, and M. Pizzonia. On embedding a cycle in a plane graph. *Discrete Mathematics*, 309(7):1856–1869, 2009.

[10] E. Dahlhaus. A linear time algorithm to recognize clustered graphs and its parallelization. In *Proc. Latin American Theoretical Informatics*, volume 1380 of *LNCS*, pages 239–248, 1998.

[11] G. Di Battista, W. Didimo, and A. Marcandalli. Planarization of clustered graphs. In P. Mutzel, M. Juenger, and S. Leipert, editors, *Graph Drawing (Proc. GD '01)*, volume 2265 of *Lecture Notes in Computer Science*, pages 60–74, 2002.

[12] G. Di Battista and F. Frati. Efficient c-planarity testing for embedded flat clustered graphs with small faces. *J. Graph Alg. Appl.*, 13(3):349–378, 2009.

[13] G. Di Battista and R. Tamassia. On-line maintenance of triconnected components with SPQR-trees. *Algorithmica*, 15(4):302–318, 1996.

[14] G. Di Battista and R. Tamassia. On-line planarity testing. *SIAM J. Comput.*, 25:956–997, 1996.

[15] C. Erten and S. G. Kobourov. Simultaneous embedding of planar graphs with few bends. *J. Graph Algorithms Appl.*, 9(3):347–364, 2005.

[16] Q. Feng, R. F. Cohen, and P. Eades. Planarity clustered graphs. In *Proc. European Symposium on Algorithms*, volume 979 of *LNCS*, pages 213–226, 1995.

[17] Q. W. Feng, R. F. Cohen, and P. Eades. How to draw a planar clustered graph. In *COCOON'95*, volume 959 of *LNCS*, pages 21–30, 1995.

[18] M. Forster. *Crossings in clustered level graphs.* PhD thesis, University of Passau, 2005.

[19] M. R. Garey and D. S. Johnson. The Rectilinear Steiner Tree Problem is NP-Complete. *SIAM J. Appl. Math.*, 32:826–834, 1977.

[20] M. R. Garey and D. S. Johnson. Crossing Number is NP-Complete. *SIAM Journal on Algebraic and Discrete Methods*, 4(3):312–316, 1983.

[21] M. T. Goodrich, G. S. Lueker, and J. Z. Sun. C-planarity of extrovert clustered graphs. In *GD*, volume 3843 of *LNCS*, pages 211–222, 2006.

[22] C. Gutwenger, M. Jünger, S. Leipert, P. Mutzel, M. Percan, and R. Weiskircher. Advances in c-planarity testing of clustered graphs. In *GD'02*, LNCS, pages 220–235, 2002.

[23] C. Gutwenger and P. Mutzel. A linear time implementation of SPQR-trees. In J. Marks, editor, *Graph Drawing (GD '00)*, volume 1984 of *LNCS*, pages 77–90, 2001.

[24] V. Jelínek, E. Jelínková, J. Kratochvíl, and B. Lidicky. Clustered planarity: Embedded clustered graphs with two-component clusters. In *GD '08*, volume 5417 of *LNCS*, pages 121–132, 2008.

[25] V. Jelínek, O. Suchý, M. Tesar, and T. Vyskocil. Clustered planarity: Clusters with few outgoing edges. In *GD '08*, pages 102–113, 2009.

[26] E. Jelínková, J. Kara, J. Kratochvíl, M. Pergel, O. Suchý, and T. Vyskocil. Clustered planarity: Small clusters in eulerian graphs. In *GD '07*, volume 4875 of *LNCS*, pages 303–314, 2007.

[27] F. Kammer. Simultaneous embedding with two bends per edge in polynomial area. In *SWAT*, volume 4059 of *LNCS*, pages 255–267, 2006.

[28] J. Pach and G. Tóth. Which crossing number is it anyway? *J. Comb. Theory, Ser. B*, 80(2):225–246, 2000.