

Characterizing Families of Cuts that can be Represented by Axis-Parallel Rectangles

Ulrik Brandes

Fachbereich Informatik & Informationswissenschaft,
Universität Konstanz
<http://www.inf.uni-konstanz.de/~brandes>
brandes@inf.uni-konstanz.de

Sabine Cornelsen

Fachbereich Informatik & Informationswissenschaft,
Universität Konstanz
<http://www.inf.uni-konstanz.de/~cornelse>
cornelse@inf.uni-konstanz.de

Dorothea Wagner

Institut für Logik, Komplexität und Deduktionssysteme,
Universität Karlsruhe
<http://i11www.informatik.uni-karlsruhe.de/algo/people/dwagner>
dwagner@ira.uka.de

Abstract

A drawing of a family of cuts of a graph is an augmented drawing of the graph such that every cut in the family is represented by a simple closed curve and vice versa.

We show that the families of cuts that admit a drawing in which every cut is represented by an axis-parallel rectangle are exactly those that have a cactus model that can be rooted such that edges of the graph that cross a cycle of the cactus point to the root. This includes the family of all minimum cuts of a graph. The proof also yields an efficient algorithm to construct a drawing with axis-parallel rectangles if it exists.

Article Type	Communicated by	Submitted	Revised
regular paper	G. Liotta	January 2004	July 2005

1 Introduction

A cut of a graph is a partition of its vertex set into two non-empty subsets. In a drawing of a graph, it is therefore natural to represent a cut by a closed curve partitioning the plane into two regions containing one subset each. See Fig. 1(a) for an example.

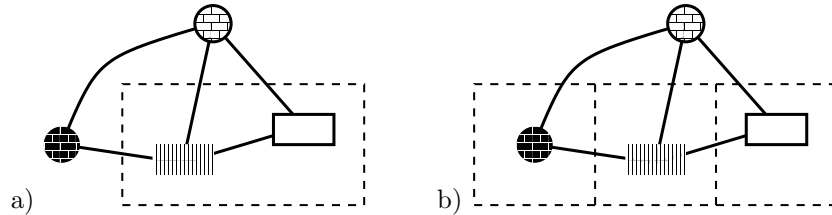


Figure 1: a) Example of a drawing of the cut separating rectangles from circles. b) Example of a simultaneous drawing of six cuts separating vertices according to shape, border, and filling style.

When a set of cuts is drawn in this way, the curves can intersect so that their union might contain closed curves that represent other cuts not contained in that set. We are interested in families of cuts that can be drawn without creating confusing non-members. In particular, we consider the problem of drawing families of cuts such that every member is represented by an axis-parallel rectangle and vice versa. For example, Fig. 1(b) shows a simultaneous drawing with axis-parallel rectangles of six cuts. The rectangles representing the cuts that separate vertices of different shape or vertices with different border, respectively, intersect. Drawings of this kind have been studied so far for the set of all minimum cuts of a weighted connected planar graph [1].

Our main result is a characterization of all families of cuts that can be represented by axis-parallel rectangles, namely those that can be modeled by a cactus containing a certain, rootable, node. These include the important family of all minimum cuts of a graph [4]. We give an algorithm to test this property. Provided a set of cuts of a graph with n vertices and m edges is represented by a cactus model, the test works in $\mathcal{O}(nm)$ time. We also show how to construct a drawing with axis-parallel rectangles if one exists. The construction works in two steps. First, we solve the problem for families of minimal cuts of a planar connected graph. Then, we use similar planarization techniques as in [3] to extend the result to general graphs.

The paper is organized as follows. In Sect. 2, we define drawings of families of cuts. In Sect. 3, we introduce the cactus model and show that its existence is necessary for the existence of a drawing with axis-parallel rectangles. Using hierarchically clustered graphs (briefly introduced in Sect. 4), we give an additional necessary condition for families of cuts that admit a drawing with axis-parallel rectangles in Sect. 5 and show how to test it. To show that the conditions are also sufficient, we construct a drawing in Sect. 6.

2 Drawing Families of Cuts

Throughout this paper, let $G = (V, E)$ denote a simple, connected, undirected graph with n vertices and m edges. A drawing \mathcal{D} of G maps vertices on distinct points in the plane and edges on simple curves connecting the drawings of their incident vertices. The interior of an edge must not contain the drawing of a vertex. A drawing is planar, if edges do not intersect but in common end points.

Let $S, T \subseteq V$. With $G(S)$ we denote the subgraph of G that is induced by S and with $E(S, T)$ we denote the set of edges that are incident to a vertex in S and a vertex in T . A *cut* of G is a partition $C = \{S, \bar{S}\}$ of the vertex set V into two non-empty subsets S and $\bar{S} := V \setminus S$. We say that the cut $\{S, \bar{S}\}$ is induced by S . The edges in $E(C) = E(S, \bar{S})$ are the *cut-edges* of C . A *minimum cut* is a cut with the minimum number of cut-edges among all cuts of G . A *minimal cut* is a cut $\{S, \bar{S}\}$ that is inclusion-minimal, i.e. there is no cut $\{T, \bar{T}\}$ with $E(T, \bar{T}) \subsetneq E(S, \bar{S})$. Note that both $G(S)$ and $G(\bar{S})$ are connected if and only if $\{S, \bar{S}\}$ is a minimal cut of the connected graph G .

A *drawing of a cut* $C = \{S, \bar{S}\}$ of G in a drawing $\mathcal{D}(G)$ is a simple closed curve γ , such that

- γ *separates* S and \bar{S} , i.e. the drawings of the edges and vertices in $G(S)$ and $G(\bar{S})$, respectively, are in different connected regions of $\mathbb{R}^2 \setminus \gamma$, and
- $|\mathcal{D}(e) \cap \gamma| = 1$ for $e \in E(C)$, i.e. the drawing of a cut-edge crosses the drawing of the cut exactly once.

Let \mathcal{C} be a set of cuts of a graph G . A mapping \mathcal{D} is a (*planar*) *drawing* of a graph G and a family of cuts \mathcal{C} of G , if

1. $\mathcal{D}(G)$ is a (planar) drawing of G and
2. $\mathcal{D}(C)$ is a drawing of C in $\mathcal{D}(G)$ for every cut $C \in \mathcal{C}$, and
3. every simple closed curve $\gamma \subseteq \bigcup_{C \in \mathcal{C}} \mathcal{D}(C)$ is a drawing of some cut in \mathcal{C} .

Note that the third condition is important if there are some cuts in \mathcal{C} whose drawings intersect. It eliminates any potential ambiguity regarding which cuts are in the family.

A drawing \mathcal{D} of a graph and a family \mathcal{C} of cuts is a *drawing with axis-parallel rectangles*, if every cut is drawn as an axis-parallel rectangle – more precisely, if every simple closed curve in $\bigcup_{C \in \mathcal{C}} \mathcal{D}(C)$ is an axis-parallel rectangle. See Fig. 2(a) for a planar drawing with axis-parallel rectangles of the set of all minimum and minimum+1 cuts of a graph.

3 Necessity of a Cactus Model

A *cactus* is a connected graph in which every edge belongs to at most one cycle. A *cactus model* for a set \mathcal{C} of cuts of a graph $G = (V, E)$ is a pair (\mathcal{G}, φ) that

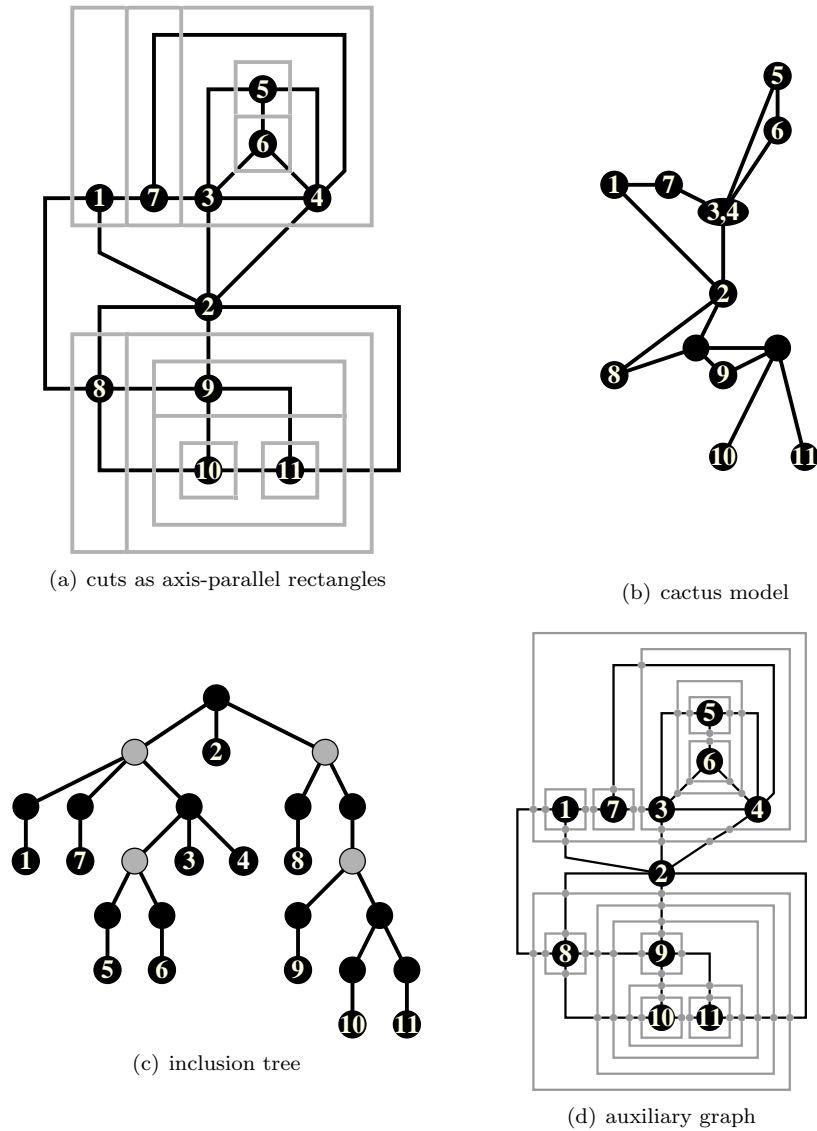


Figure 2: A drawing of the family of all minimum and minimum+1 cuts of a graph with axis-parallel rectangles. In the cactus model of the cuts, φ is indicated by node labels, and cycle-replacement nodes in the inclusion tree constructed from the cactus are shown in grey. The auxiliary graph is shown in a c-planar drawing of the hierarchical clustering represented by the inclusion tree.

consists of a cactus $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and a mapping $\varphi : V \rightarrow \mathcal{V}$ such that \mathcal{C} is *modeled* by the minimal cuts of \mathcal{G} , i.e.,

$$\mathcal{C} = \{\{\varphi^{-1}(S), \varphi^{-1}(\overline{S})\}; \{S, \overline{S}\} \text{ is a minimal cut of } \mathcal{G}\}.$$

To avoid confusion, the vertices of the cactus \mathcal{G} are called *nodes*. A node $\nu \in \mathcal{V}$ is called *empty* if $\varphi^{-1}(\nu) = \emptyset$.

The cactus model was introduced by Dinitz et al. [4] to model the set of all minimum cuts of a connected graph. A cactus model for the minimum and minimum+1 cuts of the graph in Fig. 2(a) is given in Fig. 2(b).

While not every family of cuts has a cactus model, we show that only those that do can have a drawing with axis-parallel rectangles. Two cuts $\{S, \overline{S}\}$ and $\{T, \overline{T}\}$ *cross*, if and only if the four *corners* $S \cap T$, $S \cap \overline{T}$, $T \cap \overline{S}$, and $\overline{S} \cap \overline{T}$ are non-empty. The four cuts induced by the four corners of two crossing cuts, respectively, are called *corner cuts*. The cut induced by $(S \setminus T) \cup (T \setminus S)$ is called the *diagonal cut*.

Theorem 1 ([6]) *A set \mathcal{C} of cuts of the graph G can be modeled by a cactus if and only if, for any two crossing cuts $\{S, \overline{S}\}$ and $\{T, \overline{T}\}$ in \mathcal{C} ,*

- *the four corner cuts are in \mathcal{C} , and*
- *the diagonal cut is not in \mathcal{C} .*

If a cactus model exists, there is always one with $O(n)$ nodes.

A constructive proof of Theorem 1 is given by Dinitz and Nutov in [5]. An algorithm for constructing a cactus model from another representation of a family of cuts can easily be deduced from their construction. However, the running time of such a construction is not discussed.

The properties of crossing cuts in the characterization in Theorem 1 are implied by overlapping axis-parallel rectangles.

Lemma 1 *If a set of cuts has a drawing with axis-parallel rectangles, it has a cactus model.*

Proof: Let \mathcal{D} be a drawing of a set \mathcal{C} of cuts with axis-parallel rectangles and suppose that \mathcal{C} contains crossing cuts $\{S, \overline{S}\}$ and $\{T, \overline{T}\}$. There are essentially the seven cases indicated in Fig. 3 for the drawings of two crossing cuts by axis-parallel rectangles. Since the cases in Fig. 3(a) contain simple closed curves that are not axis-parallel rectangles, only the case in Fig. 3(b) needs to be considered.

Let $D_S, D_T \subseteq \mathbb{R}^2$ be the rectangular regions bounded by $\mathcal{D}(\{S, \overline{S}\})$ and $\mathcal{D}(\{T, \overline{T}\})$, respectively. Then regions $D_S \cap D_T$, $D_S \setminus D_T$, $D_T \setminus D_S$, and $D_S \cup D_T$ are bounded by axis-parallel rectangles in $\mathcal{D}(\{S, \overline{S}\}) \cup \mathcal{D}(\{T, \overline{T}\})$. These are drawings of the four corner cuts of $\{S, \overline{S}\}$ and $\{T, \overline{T}\}$. Hence, they are in \mathcal{C} .

On the other hand, suppose the diagonal cut C induced by $(S \setminus T) \cup (T \setminus S)$ is in \mathcal{C} and let D_C be the rectangular region bounded by its drawing $\mathcal{D}(C)$. Without loss of generality, suppose that $\mathcal{D}(G(S)) \subset D_S$ and $\mathcal{D}(G(T)) \subset D_T$.

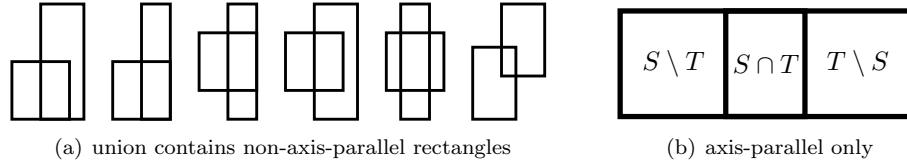


Figure 3: Drawings of two crossing cuts $\{S, \bar{S}\}$ and $\{T, \bar{T}\}$ with axis-parallel rectangles.

Either D_C contains $\mathcal{D}(G(S \setminus T))$ and $\mathcal{D}(G(T \setminus S))$, but not $\mathcal{D}(G(S \cap T))$, or it contains $\mathcal{D}(G((S \setminus T) \cup (T \setminus S)))$, but not $\mathcal{D}(G((S \setminus T) \cup (T \setminus S)))$. In the first case, the drawing contains a simple closed curve bounding the region $D_C \cap D_S \cap D_T$ (thus inducing an empty cut), and in the second case the union of the three rectangles contains a simple closed curve that is not an axis-parallel rectangle. \square

4 Cactus-Induced Hierarchical Clusterings

Since a cactus model is necessary for a drawing to exist, we can make use of a transformation originally developed for drawing the particular family of all minimum cuts of a planar connected graph [1]. Given a cactus model (\mathcal{G}, φ) of a family \mathcal{C} of cuts of G , we use the tree $T = \mathcal{T}(\mathcal{G}, \varphi)$ that is constructed as follows.

1. Replace each cycle c of \mathcal{G} by a star, i.e. delete every edge of c , add a new empty node ν_c – called *cycle-replacement node* of c – to \mathcal{G} , and for every node ν of c , add an edge $\{\nu_c, \nu\}$ to \mathcal{G} .
2. For each $v \in V$, add v to the vertex set of \mathcal{G} and add an edge $\{v, \varphi(v)\}$ to \mathcal{G} .

The construction is illustrated in Fig. 2(c). We will consider T as a rooted tree, i.e., we will fix an inner vertex r of T as the root. This yields a triple (G, T, r) of

- a graph $G = (V, E)$,
- a tree T , and
- an inner vertex r of T
- where the set of leaves of T is exactly V .

This is the hierarchically clustered graph model introduced by Feng et al. [7]. G is called the *underlying graph* and T the *inclusion tree* of (G, T, r) . Inner vertices of T are called *nodes*. We denote the tree T rooted at r by (T, r) . Each

node ν of T represents the *cluster* $V_r(\nu)$ of leaves in the subtree of (T, r) rooted at ν . Note that $\{\{V_r(\nu), \overline{V_r(\nu)}\}; \nu \neq r \text{ node of } \mathcal{T}(\mathcal{G}, \varphi)\}$ equals the subset \mathcal{C}_{\parallel} of *non-crossing cuts* of \mathcal{C} , i.e. the set of cuts in \mathcal{C} that do not cross any other cut in \mathcal{C} . A common representation of a rooted tree (T, r) is the *inclusion representation*: Each node or leaf of T is represented by a simple closed region bounded by a simple closed curve. Leaves may also be represented as points. The drawing of a node or leaf ν of T is contained in the interior of the region representing a node μ of T if and only if μ is contained in the path from ν to r in T . The drawings of two nodes μ and ν are disjoint if neither μ is contained in the path from ν to r nor ν is contained in the path from μ to r in T . A *c-planar drawing* [7] of a hierarchically clustered graph (G, T, r) consists of

- a planar drawing of the underlying graph G and
- an inclusion representation of the rooted tree (T, r) such that
- each edge crosses the boundary of the drawing of a node of T at most once.

Note that the vertices of G are the leaves of T and thus have the same drawing. Also note, that a c-planar drawing of $(G, \mathcal{T}(\mathcal{G}, \varphi), r)$ contains a planar drawing of the set of all non-crossing cuts \mathcal{C}_{\parallel} . A hierarchically clustered graph (G, T, r) is *c-planar* if it has a c-planar drawing. It is *completely connected*, if each cluster and the complement of each cluster induces a connected subgraph of G . A hierarchically clustered graph with planar underlying graph does not have to be c-planar. However, a completely connected hierarchically clustered graph is c-planar if and only if the underlying graph is planar [2].

According to the construction in [1, 10], we associate an auxiliary graph $G_{\mathcal{D}}$ with a c-planar drawing \mathcal{D} of a hierarchically clustered graph. Let V' be the set of points, in which drawings of edges and boundaries of drawings of clusters intersect. Then the vertex set of $G_{\mathcal{D}}$ is $V \cup V'$. The edge set of $G_{\mathcal{D}}$ contains two types of edges. For an edge $e = \{v, w\}$ of G , let v_1, \dots, v_k be the points in $\mathcal{D}(e) \cap V'$ in the order they occur in the drawing of e from v to w . Then $G_{\mathcal{D}}$ contains the edges $\{v, v_1\}, \{v_1, v_2\}, \dots, \{v_k, w\}$. Let $\nu \neq r$ be a node of T . Let v_1, \dots, v_k be the points in $\partial\mathcal{D}(\nu) \cap V'$ in the order they occur in the boundary $\partial\mathcal{D}(\nu)$ of the drawing of ν . Then $G_{\mathcal{D}}$ contains the edges $\{v_1, v_2\}, \dots, \{v_{k-1}, v_k\}, \{v_k, v_1\}$. The cycle v_1, \dots, v_k of $G_{\mathcal{D}}$ is called the *boundary cycle* of ν . (To avoid loops and parallel edges, additional vertices of degree two may be inserted into boundary cycles). See Fig. 2(d).

5 Towards a Characterization

Another necessary condition for families of cuts that have a drawing with axis-parallel rectangles depends also on the edges in the graph. Let (\mathcal{G}, φ) be the cactus model of a set \mathcal{C} of cuts of a graph G . For a cycle $c : \nu_1, \dots, \nu_k$ in \mathcal{G} we consider the tree $T = \mathcal{T}(\mathcal{G}, \varphi)$ rooted at the cycle-replacement node ν_c . Let

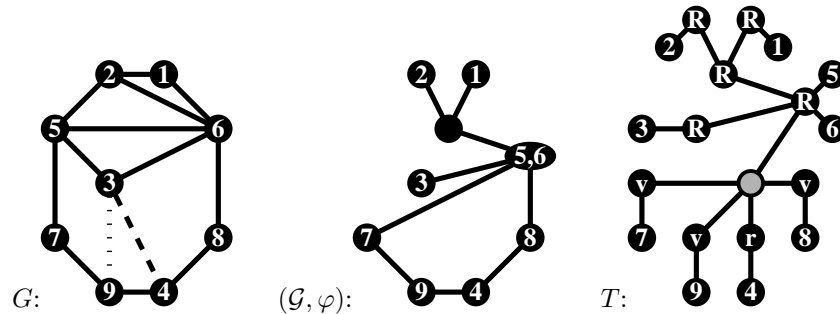


Figure 4: The dashed edge in G crosses the only cycle in the cactus model \mathcal{G} towards a root labeled r or R of the inclusion tree $T = \mathcal{T}(\mathcal{G}, \varphi)$ but not to a root labeled v . The nodes labeled R are exactly the rootable nodes.

$V_i := V_{\nu_c}(\nu_i), i = 1, \dots, k$. An edge $\{v, w\}$ of G crosses the cycle c if there are $1 \leq i, j \leq k, i \neq j$ such that

$$v \in V_i, w \in V_j, \text{ and } i - j \not\equiv \pm 1 \pmod k.$$

For example, in the cactus model \mathcal{G} in Fig. 4 there is one cycle $c : \nu_1 = 7, \nu_2 = 9, \nu_3 = 4, \nu_4 = 8, \nu_5 = 5, 6$. The corresponding subsets of the set of vertices of G are $V_1 = \{7\}$, $V_2 = \{9\}$, $V_3 = \{4\}$, $V_4 = \{8\}$, and $V_5 = \{1, 2, 3, 5, 6\}$. The dashed edge of G is incident to V_5 and V_3 . Hence it crosses c . The only other edge that crosses c is the dotted edge $\{3, 9\}$.

If (\mathcal{G}, φ) is the cactus of all minimum cuts of G , then no edge of G crosses a cycle of \mathcal{G} (see, e.g., [1, Lemma 7]). In general, it depends on the edges that cross a cycle of the cactus model, whether a set of cuts has a drawing with axis-parallel rectangles. More precisely, if \mathcal{C} has a drawing \mathcal{D} with axis-parallel rectangles and e crosses the cycle c then there exists an $i \in \{1, \dots, k\}$ such that e is incident to a vertex in V_i and the drawing of \bar{V}_i is contained in the simply connected region bounded by $\mathcal{D}(V_i, \bar{V}_i)$.

To formalize this statement, we say that an edge e of G that crosses a cycle c of the cactus model \mathcal{G} crosses c towards a root r of the inclusion tree $\mathcal{T}(\mathcal{G}, \varphi)$ if e is incident to a vertex in $V \setminus V_r(\nu_c)$. See Fig. 4 for an example.

Lemma 2 *A family of cuts that has a cactus model (\mathcal{G}, φ) has a drawing with axis-parallel rectangles only if the root r of $\mathcal{T}(\mathcal{G}, \varphi)$ can be chosen such that*

(R) *each edge of G that crosses a cycle of \mathcal{G} crosses it towards r .*

Proof: Suppose that the set of cuts modeled by (\mathcal{G}, φ) has a drawing \mathcal{D} with axis-parallel rectangles. Let r be a node of $T = \mathcal{T}(\mathcal{G}, \varphi)$ such that for every node $\nu \neq r$ of T the set $V_r(\nu)$ is contained in the simple region bounded by $\mathcal{D}(\{V_r(\nu), \bar{V}_r(\nu)\})$. Using the fact that $\bar{V}_r(\mu) = V_\mu(r)$ for two adjacent nodes r and μ of T , it can be shown that such a node r exists: Let r be such that for the

maximum number of nodes $\nu \neq r$ of T the drawing of $V_r(\nu)$ is contained in the simple region bounded by $\mathcal{D}(\{V_r(\nu), \overline{V_r(\nu)}\})$. It suffices to show that all adjacent nodes μ of r have the above mentioned property. So suppose there is an adjacent node μ of r such that the drawing of $V_r(\mu)$ is not contained in the simple region bounded by $\mathcal{D}(\{V_r(\mu), \overline{V_r(\mu)}\})$. Hence, on one hand, $\overline{V_r(\mu)} = V_\mu(r)$ is contained in the simple region bounded by $\mathcal{D}(\{V_r(\mu), \overline{V_r(\mu)}\})$. On the other hand, for every node $\nu \notin \{\mu, r\}$ it holds that the drawing of $V_r(\nu)$ is contained in the simple region bounded by $\mathcal{D}(\{V_r(\nu), \overline{V_r(\nu)}\})$ if and only if the drawing of $V_\mu(\nu)$ is contained in the simple region bounded by $\mathcal{D}(\{V_\mu(\nu), \overline{V_\mu(\nu)}\})$ – contradicting the choice of r .

Let c be a cycle of \mathcal{G} and let e be an edge of G that crosses c . Let ν_1, \dots, ν_k be the children of ν_c such that ν_1, \dots, ν_k is a path in c . Suppose that e is not incident to a vertex in $V \setminus V_r(\nu_c)$. Then $e \in E(V_r(\nu_i), V_r(\nu_j))$ for some $1 < i+1 < j \leq k$. Let $S = \bigcup_{\ell=1}^{i+1} V_r(\nu_\ell)$ and $T = \bigcup_{\ell=i+1}^k V_r(\nu_\ell)$. Then $\{S, \overline{S}\}$ and $\{T, \overline{T}\}$ are both modeled by (\mathcal{G}, φ) , more precisely by the two minimal cuts $\{S', \overline{S'}\}$ and $\{T', \overline{T'}\}$ of \mathcal{G} such that $E(S', \overline{S'}) = \{\{\nu_0, \nu_1\}, \{\nu_{i+1}, \nu_{i+2}\}\}$ and $E(T', \overline{T'}) = \{\{\nu_i, \nu_{i+1}\}, \{\nu_k, \nu_0\}\}$, where ν_0 is the predecessor of ν_c in T .

Now, on one hand, $e \in E(S \setminus T, T \setminus S)$. But, on the other hand, $\{S, \overline{S}\}$ and $\{T, \overline{T}\}$ are two crossing cuts that are modeled by (\mathcal{G}, φ) such that S and T are contained in the simple closed region D_S and D_T bounded by $\mathcal{D}(\{S, \overline{S}\})$ and $\mathcal{D}(\{T, \overline{T}\})$, respectively. Analogously to the argumentation in Lemma 1, we have again the situation as indicated in Fig. 3(b). By the definition of drawings of cuts, the drawing of the edges and vertices of $G(S \cup T)$ has to be contained inside $D_S \cup D_T$ and the drawing of the edges and vertices of $\mathcal{D}(G(\overline{S} \cap \overline{T}))$ may not intersect $D_S \cap D_T$. This implies that $E(S \setminus T, T \setminus S) = \emptyset$. \square

In the following, a node of the inclusion tree $T = \mathcal{T}(\mathcal{G}, \varphi)$ is called *rootable*, if it fulfills Condition (R) of the previous lemma. We give an algorithm for finding all rootable nodes of the inclusion tree $T = \mathcal{T}(\mathcal{G}, \varphi)$. We assume that the size of \mathcal{G} is in $\mathcal{O}(n)$.

We use a node and edge array `CROSSED` on T to store the information about edges of G that cross a cycle of the cactus. Let c be a cycle of \mathcal{G} . `CROSSED`(ν_c) is true if and only if there is an edge of G that crosses c . `CROSSED`($\{\nu_c, \nu\}$) is true for an edge $\{\nu_c, \nu\}$ of T if and only if `CROSSED`(ν_c) is true and each edge that crosses c is incident to a vertex in $V_{\nu_c}(\nu)$. For all other edges and nodes of T `CROSSED` is false. In the example tree T in Fig. 4, `CROSSED`(ν_c) is true for the only cycle c of \mathcal{G} and `CROSSED`($\{\nu_c, R\}$) is true for the adjacent node of ν_c labeled R . For all other edges and nodes of T , `CROSSED` is false. Clearly, it can be tested in linear time, which cycles are crossed by an edge. Hence, `CROSSED` can be computed in $\mathcal{O}(mn)$ time. The next remark is a reformulation of Condition (R).

Remark 1 *A node r of T is rootable if and only if for each node ν of T with `CROSSED`(ν) = TRUE there exists a node μ adjacent to ν with `CROSSED`{ ν, μ } = TRUE such that μ is in the path from ν to r .*

Hence, we can obtain the possible roots for constructing a drawing with axis-parallel rectangles of \mathcal{C} – possibly including some cycle-replacement nodes – by deleting for all marked nodes ν the subtrees of (T, ν) rooted at those adjacent nodes μ of ν for which $\{\nu, \mu\}$ is not marked. Proceeding first from the leaves to an arbitrary root and then from this root to the leaves of T , this can be done in time $\mathcal{O}(n)$.

By constructing a drawing, we show in the next section that the conditions in Lemma 2 are also sufficient.

6 The Drawing

Let (\mathcal{G}, φ) continue to be a cactus model of a set \mathcal{C} of cuts of the graph G and let $T = \mathcal{T}(\mathcal{G}, \varphi)$ be the inclusion tree constructed in Sect. 4. We assume that the root r of T is a rootable node. Without loss of generality, we may assume that r is not a cycle-replacement node. We show how to construct a drawing with axis-parallel rectangles for \mathcal{C} . In a first step, we consider the case that \mathcal{C} is a set of minimal cuts of a connected planar graph. Using planarization, we generalize the result to general sets of cuts of not necessarily planar graphs.

6.1 Planar Graphs

In this section, we assume that \mathcal{C} is a set of minimal cuts of a planar connected graph G . We show how the ideas for drawing the set of all minimum cuts of a planar graph [1] yield a drawing with axis-parallel rectangles of \mathcal{C} . We briefly sketch the general construction and explain in more details the parts that differ from the case of drawings for the set of minimum cuts. The construction starts with a c-planar drawing \mathcal{D} of the hierarchically clustered graph (G, T, r) . Such a c-planar drawing always exists: Since \mathcal{C} is a set of minimal cuts of a connected graph it follows that (G, T, r) is completely connected and hence c-planar.

Let $c : \nu_1, \dots, \nu_\ell$ be a cycle of \mathcal{G} . Suppose that ν_c is a child of ν_ℓ in the tree T rooted at r . Let $V_i = V_{\nu_c}(\nu_i)$, $i = 1, \dots, \ell$ and let $V_0 = V_\ell$. Let $i \in \{1, \dots, \ell - 1\}$ and let e_1, \dots, e_k be the sequence of edges incident to any vertex in V_i in the cyclic order around the boundary of V_i . Recall that by hypothesis $E(V_i, V_j)$ is empty for $j \neq i - 1, i + 1, \ell$. We will now show that $E(V_i, V_{i-1})$ and $E(V_i, V_{i+1})$ are non-empty subsequences of e_1, \dots, e_k , i.e. suppose $e_1 \in E(V_i, V_{i-1})$ and $e_k \notin E(V_i, V_{i-1})$, then there are indices $1 < k_1 \leq k_2 < k_3 \leq k$ such that

$$\underbrace{e_1, \dots, e_{k_1}}_{E(V_i, V_{i-1})} \underbrace{e_{k_1+1}, \dots, e_{k_2}}_{E(V_i, V_\ell)} \underbrace{e_{k_2+1}, \dots, e_{k_3}}_{E(V_i, V_{i+1})} \underbrace{e_{k_3+1}, \dots, e_k}_{E(V_i, V_\ell)}. \quad (1)$$

Further, $E(V_\ell, V_1)$ and $E(V_\ell, V_{\ell-1})$ are non-empty subsequence of the edges around $V \setminus V_\ell$. We prove the observation in two steps.

1. First, we show that the sets $E(V_{i-1}, V_i)$, $i = 1, \dots, \ell$ are non-empty. Else suppose without loss of generality that $E(V_\ell, V_1) = \emptyset$. Let $C = \{V_\ell \cup V_1, V_2 \cup \dots \cup V_{\ell-1}\}$. Then it follows that $E(C)$ is the disjoint union of

$E(V_1, \overline{V_1})$ and $E(V_\ell, \overline{V_\ell})$. Since G is connected it follows that $\emptyset \neq E(V_\ell, \overline{V_\ell})$. Hence, on one hand, $E(V_1, \overline{V_1}) \subsetneq E(C)$. On the other hand, C is modeled by \mathcal{G} and hence it is contained in the set \mathcal{C} of minimal cuts – a contradiction.

2. Second, we show that $E(V_i, V_{i-1}), E(V_i, V_{i+1}), i = 1, \dots, \ell - 1$ are subsequences of the cyclic sequence of edges around V_i . Recall that by the connectedness of G and the minimality of the cuts in \mathcal{C} all sets $V_i, i = 1, \dots, \ell$ are connected. Let $W = V_{i-1}$ or $W = V_{i+1}$, respectively. Now suppose that in the cyclic order around a cluster $V_i, i = 1, \dots, \ell - 1$ there is a subsequence e_1, e_2, e_3, e_4 such that $e_1, e_3 \in E(V_i, W)$ and $e_2, e_4 \notin E(V_i, W)$. Let $e_j = \{v_j, w_j\}, j = 1, \dots, 4$ with $v_j \in V_i$. Let p be a path in V_i from v_1 to v_3 and let q be a path in W from w_3 to w_1 . Then the cycle $\{w_1, v_1\}p\{v_3, w_3\}q$ divides the plane into two connected regions such that w_2 and w_4 are in different regions. By Step 1 and the connectivity of the sets $V_i, i = 1, \dots, \ell$, there is a path from w_2 to w_4 in G that does not intersect V_i or W . But this contradicts planarity. Analogously, it can be shown that $E(V_\ell, V_1)$ and $E(V_\ell, V_{\ell-1})$ are subsequences of the cyclic sequence of the edges around $V \setminus V_\ell$.

Besides, let $e \in E(V_i, V_j)$ for some $i, j \in \{1, \dots, \ell\}, i \neq j$. Then c-planarity implies that the two vertices in the auxiliary graph $G_{\mathcal{D}}$ that represent the intersection of e with the boundary of ν_i and ν_j (or the boundary of ν_c , if $i = \ell$ or $j = \ell$), respectively, are adjacent. Hence the situation in the auxiliary graph $G_{\mathcal{D}}$ is as indicated in Fig. 5(a). For $i \in \{1, \dots, \ell - 1\}$, let p_i^\pm be the path in $G_{\mathcal{D}}$ that is induced by the intersection of the edges in $E(V_i, V_{i\pm 1})$ with the boundary of ν_i . Similarly, let p_0^+ (p_ℓ^-) be the path in $G_{\mathcal{D}}$ that is induced by the intersection of the edges in $E(V_\ell, V_1)$ ($E(V_\ell, V_{\ell-1})$) with the boundary of ν_c .

A planar drawing of \mathcal{C} with axis-parallel rectangles can now be obtained as follows. First, for each i , paths p_i^+ and p_{i+1}^- are united to one path. The parts of the boundary-cycle of ν_c that are not in p_0^+ or p_ℓ^- are removed. The result can be seen in Fig. 5(b). Finally, at each end of each thus united path p_i^+/p_{i+1}^- , $i = 1, \dots, \ell - 2$, an additional vertex is inserted as indicated in Fig. 5(c). In the following, we refer to these vertices as *cycle-path end-vertices*. Let G' be the graph that results from $G_{\mathcal{D}}$ by doing this for every cycle of \mathcal{G} . Applying a min-cost flow approach for quasi-orthogonal drawings [11, 9, 8, 12] with appropriate restrictions on the flow to G' yields a drawing of \mathcal{C} with axis-parallel rectangles. There are two restrictions on the flow necessary:

1. the flow over a boundary edge from outside a boundary cycle into the inside of a boundary cycle is zero;
2. the flow from a cycle-path end-vertex into the inside of a boundary cycle is minimum, i.e. there is a rectangle.

As demonstrated in [1, Lemma 19], it can be shown that there is always a feasible flow for the thus restricted flow network. Note that now the corners of the drawings of the cuts of \mathcal{C} are either at the cycle-path end-vertex or at bends.

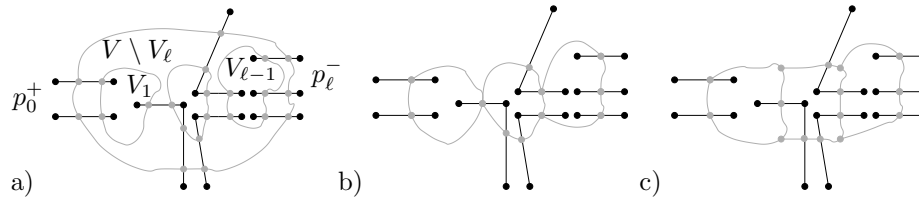


Figure 5: Constructing a drawing for a set of cuts from a c-planar drawing.

Finally, let us briefly mention the case where a combinatorial embedding \mathcal{E} for G is given. By [2], the c-planar drawing \mathcal{D} of (G, T, r) can be chosen such that the underlying drawing of G is according to \mathcal{E} . Note, however, that the outer face has to be chosen in a suitable way. Since the rest of the construction does not effect the embedding, it follows that a drawing of a set of minimal cuts of a planar graph with a fixed combinatorial embedding can be constructed such that the embedding (but not necessarily the outer face) is respected.

6.2 General Graphs

In this section, we extend the results on planar drawings for families of minimal cuts to not necessarily planar drawings. Similar to the method described in [3] for hierarchically clustered graphs, the idea for the construction uses planarization techniques. Recall that (\mathcal{G}, φ) is a cactus model for a family of cuts of G . We assume again that the root r of $T = T(\mathcal{G}, \varphi)$ is rootable and not a cycle-replacement node.

1. Let $G_P = (V, E_P)$ be a planar connected graph on the same set of vertices as G , such that \mathcal{G} models a set of minimal cuts and each edge of G_P that crosses a cycle \mathcal{G} crosses it towards r .
2. Construct the auxiliary graph G'_P for $(G_P, \mathcal{G}, \varphi)$ as it is described in Sect. 6.1.
3. Construct a planar graph G' , by inserting the remaining edges of G into G'_P such that each edge crosses the boundary cycle of a cut at most once.
4. Apply an algorithm for quasi-orthogonal drawings with the restrictions indicated in Sect. 6.1 to G' .
5. Delete the edges that do not correspond to original edges in G .

Remarks to Step 1. Note that even if \mathcal{G} models a set of minimal cuts of G , it is not always possible to choose the graph G_P with the properties required in Step 1 as a subgraph of G . For example consider that G is constructed from a $K_{3,3}$ by inserting three vertices into each edge and that (\mathcal{G}, φ) is the cactus of all minimum cuts of G . The following lemma shows, however, that at least some graph with the properties required in Step 1 exists.

Lemma 3 *Let (\mathcal{G}, φ) be a cactus model for a set of cuts of a graph with vertex set V . Then there exists a planar connected graph H with vertex set V such that:*

1. \mathcal{G} models a set of minimal cuts of H and
2. no edge of H crosses a cycle of \mathcal{G} .

Proof: Let μ be a node of T that is adjacent to a leaf v . In a first step, we start with an empty graph. Proceeding from bottom to top of the tree T rooted at μ , we can construct a graph $H_0 = (V, E_0)$ such that:

1. for any node $\nu \neq \mu$ the subgraph of H_0 induced by $V_\mu(\nu)$ is a path,
2. no edge of H_0 crosses a cycle of \mathcal{G} , and
3. $E_0(V_\mu(\nu), \overline{V_\mu(\nu)}) = \emptyset$ if ν is adjacent to μ .

Let $\nu \neq \mu$ be a node of T and let ν_1, \dots, ν_k be the children of ν that are adjacent to ν . If ν is a cycle-replacement node, suppose that we ordered the ν_i such that ν_1, \dots, ν_k is a path in \mathcal{G} . If ν_i is a leaf for some i , let $V_\mu(\nu_i) = \{\nu_i\}$. We assume that $V_\mu(\nu_1), \dots, V_\mu(\nu_k)$ already induce paths in G_0 . Now, for each $i = 1, \dots, k-1$, let $v_1 \in V_\mu(\nu_i)$ be a vertex of degree one or zero in $G_0(\bigcup_{j=1}^i V_\mu(\nu_j))$ and let v_2 be a vertex of degree one or zero in $G_0(V_\mu(\nu_{i+1}))$. Add $\{v_1, v_2\}$ to G_0 .

In a second step, for each node ν of T that is adjacent to μ let v_1 and v_2 be the two vertices of degree one or let $v_1 = v_2$ be the vertex of degree zero in the subgraph of H_0 induced by $V_\mu(\nu)$, respectively. Add edges $\{v, v_1\}$ and $\{v, v_2\}$ to H_0 . Finally, for all leaves $w \neq v$ of T that are adjacent to μ , add an edge $\{v, w\}$ to H_0 .

It is now easy to see that no edge of the thus constructed graph H crosses a cycle of \mathcal{G} and that \mathcal{G} models a set of minimal cuts. H is a cactus and hence it is planar and connected. \square

Let $H = (V, E_H)$ be the graph from the previous lemma and let G_P be a maximal planar subgraph of $(V, E \cup E_H)$ such that H is a subgraph of G_P . Then G_P is planar and connected, \mathcal{G} models a set of minimal cuts of G_P and r is rootable. Hence, G_P fulfills the properties required in Step 1.

Remarks to Step 3. We use an extension of planarization-techniques that is similar to the method introduced by Di Battista et al. [3] for drawing non-planar hierarchically clustered graphs. Edges in G that are not represented in G_P are routed iteratively through the dual graph of G'_P , replacing every crossing of an edge and a dual in the route by a vertex of degree four. Suppose, we want to insert an edge $\{v, w\}$. Let F be the set of faces of G'_P . Let $E_{\{v,w\}}^*$ be the following set of arcs. For every edge e of G'_P , let $f_1, f_2 \in F$ be the faces that are incident to e . Then $E_{\{v,w\}}^*$ contains the *dual arcs* $(f_1, f_2)_e$ and $(f_2, f_1)_e$. Further, $E_{\{v,w\}}^*$ contains the arcs $(f, x), (x, f)$, $x \in \{v, w\}, f \in F$, f incident to x . There are no other arcs contained in $E_{\{v,w\}}^*$.

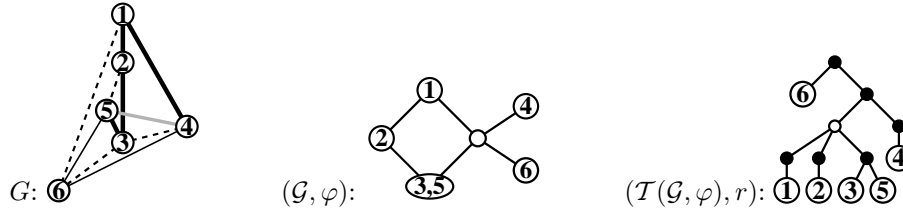


Figure 6: A non-planar graph G and a cactus model. Deleting the grey edge yields a maximal planar subgraph. Solid thick black edges induce a spanning path in every cluster $S \neq V$ of $(G, \mathcal{T}(\mathcal{G}, \varphi), r)$. All solid black edges indicate the graph H constructed in the proof of Lemma 3.

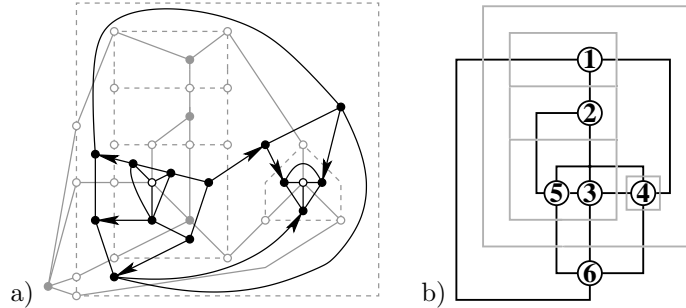


Figure 7: a) Black edges show a connected component of the restricted extended dual for adding the edge $\{5, 4\}$ from 5 to 4 of the graph in Fig. 6. Bidirected edges are indicated by simple curves without arrows. b) A drawing with axis-parallel rectangles for the family of cuts modeled by (\mathcal{G}, φ) in Fig. 6.

In general, any simple path in the *extended dual* $G_{\{v,w\}}^* = (\{v, w\} \cup F, E_{\{v,w\}}^*)$ from v to w can be used as a route for the edge $\{v, w\}$. To achieve that the drawing of $\{v, w\}$ does not cross a boundary cycle twice, we use a restricted version of the extended dual. Let $v = \nu_1, \nu_2, \dots, \nu_\ell = w$ be the path in T between v and w . Let $k \in \{2, \dots, \ell - 1\}$ be such that ν_{k-1} and ν_{k+1} are children of ν_k . For each boundary edge e of G'_P :

- if e is contained in the drawing γ of $\{V_r(\nu_i), \overline{V_r(\nu_i)}\}, i = 2, \dots, k-1$, delete the dual arc of e that is directed from the outside of γ to the inside of γ from $E_{\{v,w\}}^*$;
- if e is contained in the drawing γ of $\{V_r(\nu_i), \overline{V_r(\nu_i)}\}, i = k+1, \dots, \ell-1$, delete the dual arc of e that is directed from the inside of γ to the outside of γ from $E_{\{v,w\}}^*$;
- if e is contained in the drawing of a cut $\{V_r(\nu), \overline{V_r(\nu)}\}$ for a node $\nu \neq \nu_i, i = 2, \dots, k-1, k+1, \dots, \ell-1$ of T , delete both dual arcs of e from $E_{\{v,w\}}^*$.

Thus the restricted dual (see Fig. 7 for an illustration) guarantees that any path from v to w crosses any boundary cycle at most once. On the other hand, there is a path from v to w in the restricted dual if and only if either ν_k is not a cycle-replacement node or ν_k is a cycle-replacement node and ν_{k-1}, ν_{k+1} are adjacent in \mathcal{G} . Hence, by Condition (R), there is a path from v to w . Summarizing, we have shown the following characterization

Theorem 2 *A family \mathcal{C} of cuts of a graph G has a drawing with axis-parallel rectangles if and only if*

1. \mathcal{C} has a cactus model (\mathcal{G}, φ) and
2. the root r of $\mathcal{T}(\mathcal{G}, \varphi)$ can be chosen such that each edge that crosses a cycle of \mathcal{G} crosses it towards r .

Provided a cactus model is given, it can be tested in time $\mathcal{O}(nm)$ whether a drawing with axis-parallel rectangles exists.

Moreover, if a set of minimal cuts of a planar graph has a drawing with axis-parallel rectangles, then it has a planar drawing with axis-parallel rectangles.

References

- [1] U. Brandes, S. Cornelsen, and D. Wagner. How to draw the minimum cuts of a planar graph. *Computational Geometry: Theory and Applications*, 29:117–133, 2004.
- [2] S. Cornelsen and D. Wagner. Completely connected clustered graphs. In H. L. Bodlaender, editor, *Graph Theoretic Concepts in Computer Science, 29th International Workshop, (WG 2003)*, volume 2880 of *Lecture Notes in Computer Science*, pages 168–179. Springer, 2003.
- [3] G. Di Battista, W. Didimo, and A. Marcandalli. Planarization of clustered graphs. In M. Jünger and P. Mutzel, editors, *Proceedings of the 9th International Symposium on Graph Drawing (GD 2001)*, volume 2265 of *Lecture Notes in Computer Science*, pages 60–74. Springer, 2002.
- [4] Y. Dinitz, A. V. Karzanov, and M. Lomonosov. On the structure of a family of minimal weighted cuts in a graph. In A. Fridman, editor, *Studies in Discrete Optimization*, pages 290–306. Nauka, 1976. (in Russian).
- [5] Y. Dinitz and Z. Nutov. Cactus tree type models for families of bisections of a set. Manuscript.
- [6] Y. Dinitz and Z. Nutov. A 2-level cactus model for the system of minimum and minimum+1 edge-cuts in a graph and its incremental maintenance. In *Proceedings of the 27th Annual ACM Symposium on the Theory of Computing (STOC '95)*, pages 509–518. ACM, The Association for Computing Machinery, 1995.
- [7] Q. Feng, R. F. Cohen, and P. Eades. Planarity for clustered graphs. In P. Spirakis, editor, *Proceedings of the 3rd European Symposium on Algorithms (ESA '95)*, volume 979 of *Lecture Notes in Computer Science*, pages 213–226. Springer, 1995.
- [8] U. Fößmeier and M. Kaufmann. Drawing high degree graphs with low bend numbers. In F. J. Brandenburg, editor, *Proceedings of the 3rd International Symposium on Graph Drawing (GD '95)*, volume 1027 of *Lecture Notes in Computer Science*, pages 254–266. Springer, 1996.
- [9] G. W. Klau and P. Mutzel. Quasi orthogonal drawing of planar graphs. Technical Report MPI-I-98-1-013, Max-Planck-Institut für Informatik, Saarbrücken, Germany, 1998. Available at <http://data.mpi-sb.mpg.de/internet/reports.nsf>.
- [10] D. Lütke-Hüttmann. Knickminimales Zeichnen 4-planarer Clustergraphen. Master's thesis, Universität des Saarlandes, 1999. (Diplomarbeit).
- [11] R. Tamassia. On embedding a graph in the grid with the minimum number of bends. *SIAM Journal on Computing*, 16:421–444, 1987.

- [12] R. Tamassia, G. Di Battista, and C. Batini. Automatic graph drawing and readability of diagrams. *IEEE Transactions on Systems, Man and Cybernetics*, 18(1):61–79, 1988.