Approximation Algorithms for Facial Cycles in Planar Embeddings

Giordano Da Lozzo
Computer Science Department, Roma Tre University, Italy
dalozzo@dia.uniroma3.it
https://orcid.org/0000-0003-2396-5174

Ignaz Rutter
Department of Computer Science and Mathematics, University of Passau, Germany
rutter@fim.uni-passau.de
https://orcid.org/0000-0002-3794-4406

Abstract

Consider the following combinatorial problem: Given a planar graph G and a set of simple cycles $\mathcal{C}$ in G, find a planar embedding $\mathcal{E}$ of G such that the number of cycles in $\mathcal{C}$ that bound a face in $\mathcal{E}$ is maximized. This problem, called Max Facial $\mathcal{C}$-Cycles, was first studied by Mutzel and Weiskircher [IPCO '99] and then proved NP-hard by Woeginger [Oper. Res. Lett., 2002].

We establish a tight border of tractability for Max Facial $\mathcal{C}$-Cycles in biconnected planar graphs by giving conditions under which the problem is NP-hard and showing that strengthening any of these conditions makes the problem polynomial-time solvable. Our main results are approximation algorithms for Max Facial $\mathcal{C}$-Cycles. Namely, we give a 2-approximation for series-parallel graphs and a $(4 + \varepsilon)$-approximation for biconnected planar graphs. Remarkably, this provides one of the first approximation algorithms for constrained embedding problems.

2012 ACM Subject Classification Mathematics of computing → Graph theory

Keywords and phrases Planar Embeddings, Facial Cycles, Complexity, Approximation Algorithms

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2018.41

Related Version A full version of the paper containing omitted or sketched proofs is available as [7], `https://arxiv.org/abs/1607.02347`.

Acknowledgements This research was partially supported by MIUR Project "MODE", by H2020-MSCA-RISE project "CONNECT", and by MIUR-DAAD JMP N° 34120.

1 Introduction

A planar graph is a graph that can be *embedded* into the plane, i.e., it can be drawn into the plane without crossings. Such an embedding partitions the plane into topologically connected regions, called *faces*. There is exactly one unbounded face, which is called *outer face*. While there exist infinitely many such embeddings, the embeddings for connected graphs can be grouped into finitely many equivalence classes of *combinatorial embeddings*, where two embeddings are *equivalent* if the clockwise cyclic order of the edges around each vertex is the same and their outer face is bounded by the same walk. Since a graph may admit exponentially many different such embeddings, several drawing algorithms for planar graphs simply assume that one embedding has been fixed beforehand and draw the graph with this fixed embedding. Often, however, the quality of the resulting drawing depends

29th International Symposium on Algorithms and Computation (ISAAC 2018).
Editors: Wen-Lian Hsu, Der-Tsai Lee, and Chung-Shou Liao; Article No. 41; pp. 41:1–41:13
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

strongly on this embedding; examples are the number of bends in orthogonal drawings [17] or the area requirement of planar straight-line drawings [8].

Consequently, there is a long line of research that seeks to optimize quality measures over all combinatorial embeddings. Not surprisingly, except for a few notable cases such as minimizing the radius of the dual graph [2, 15], many of these problems have turned out to be NP-complete. For example it is NP-complete to decide whether there exists a planar embedding that allows for a planar orthogonal drawing without bends or for an upward planar drawing [12]. While there has been quite a bit of work on solving these problems for special cases, e.g., for the orthogonal bend minimization problem [4, 5], to the best of our knowledge, approximation algorithms have rarely been considered.

Another way of describing a combinatorial embedding of a connected graph G is by describing its *facial walks*, i.e., by listing the walks of G that bound a face. In the case of biconnected planar graphs, the facial walks are simple, and we refer to them as *facial cycles*. In this paper, we consider the problem of optimizing the set of facial cycles. Given a list $\mathcal{C}$ of cycles in a biconnected graph G, the problem MAX FACIAL $\mathcal{C}$-CYCLES asks for an embedding $\mathcal{E}$ of G such that the number of cycles in $\mathcal{C}$ that are facial cycles of $\mathcal{E}$ is maximized. The practical motivation for this problem comes from the need to visualize graphs in such a way that particular substructures, in this case cycles, are clearly recognizable. These structures may be either provided manually by the user or be the result of an automated analysis. We note that, given a biconnected planar graph G and a set $\mathcal{C}$ of cycles of G, it can be efficiently decided whether there exists a planar embedding of G in which all cycles of $\mathcal{C}$ are facial cycles. For each cycle $C \in \mathcal{C}$, we subdivide each edge of C once and connect the subdivision vertex to a new vertex v_C. If the resulting graph is planar, then the desired embedding of G can be obtained by removing all vertices v_C and their incident edges. However, from a practical point of view, this approach is insufficient, as it does not produce a solution if it is not possible to simultaneously have all the cycles in $\mathcal{C}$ as facial cycles. Instead, we would like to compute an embedding that maximizes the number of cycles in $\mathcal{C}$ that are facial cycles.

The research on this problem was initiated by Mutzel and Weiskircher [16], who gave an integer linear program (ILP) for a weighted version. Woeginger [19] showed that the problem is NP-complete by showing that it is NP-complete to maximize the number of facial cycles that have size at most 4. Da Lozzo et al. [6] consider the problem of deciding whether there exists an embedding such that the maximum face size is k. They give polynomial-time algorithms for $k \leq 4$, show NP-hardness for $k \geq 5$, and give a factor-6 approximation for minimizing the size of the largest face. Dornheim [10] studies a decision problem subject to so-called *topological constraints*, which specify for certain cycles of a planar graph two subsets of edges of the graph that have to be embedded inside and outside the respective cycle; note that a cycle is a facial cycle if its interior is empty. He proved NP-completeness and reduced the connected case to the biconnected case. Another related problem, known as PARTIALLY EMBEDDED PLANARITY (for short PEP), has been studied by Angelini et al. [1]. Given a planar graph G and an embedding $\mathcal{E}_\mathcal{H}$ of a subgraph H of G, the PEP problem asks for the existence of an embedding $\mathcal{E}_\mathcal{G}$ of G that *extends* $\mathcal{E}_\mathcal{H}$, i.e., the restriction of $\mathcal{E}_\mathcal{G}$ to H coincides with $\mathcal{E}_\mathcal{H}$. As observed before, if H is biconnected, then its combinatorial embedding is fully specified by the set of cycles of H (of G) that are faces in $\mathcal{E}_\mathcal{H}$. Thus, the MAX FACIAL $\mathcal{C}$-CYCLES problem can be interpreted as an optimization counterpart of PEP in which one tries to minimize the set of faces of H that are not faces in the final embedding of G.

Contribution and outline. We thoroughly study the complexity of MAX FACIAL $\mathcal{C}$-CYCLES for biconnected planar graphs. We start with preliminaries concerning connectivity and the

SPQR-tree data structure in Section 2. In Section 3, we show that Max Facial $\mathcal{C}$-Cycles is NP-complete even if each cycle in $\mathcal{C}$ intersects any other cycle in $\mathcal{C}$ in at most two vertices and intersects at most three other cycles of $\mathcal{C}$. In Section 4, we complement these results with efficient algorithms for series-parallel and general planar graphs when the cycles intersect only few other cycles in more than one vertex. We note that, though these instances are fairly restricted, this establishes a tight border of tractability for the problem in the sense that dropping or strengthening any of the conditions for our algorithms yields an NP-hard problem. Moreover, the techniques for obtaining these results are the basis for our main result in Section 5, where we develop efficient approximation algorithms for the problem. More specifically, we give a 2-approximation for series-parallel graphs and a $(4 + \varepsilon)$-approximation for biconnected planar graphs, where $\varepsilon > 0$ is a constant. We remark that, to the best of our knowledge, this work and our contribution in [6] provide one of the very few known approximability results concerning constrained combinatorial embeddings.

A full version of the paper containing omitted or sketched proofs is available as [7].

2 Connectivity and SPQR-trees

A graph G is *connected* if there is a path between any two vertices. A *cutvertex* (*separating pair*) is a vertex (a pair of vertices) whose removal disconnects the graph. A connected (biconnected) graph is *biconnected* (*triconnected*) if it does not have a cutvertex (a separating pair).

We consider uv-graphs with two special *pole* vertices u and v, which can be recursively defined as follows. An edge uv is an uv-graph with poles u and v. Now let G_i be an uv-graph with poles u_i and v_i, for $i = 1, \dots, k$, and let H be a planar graph with two designated vertices u and v and $k+1$ edges $uv, e_1, \dots, e_k$. We call H the *skeleton* of the composition and its edges are called *virtual edges*; the edge uv is the *parent edge* and u and v are the poles of the skeleton H. To compose the G_i into an uv-graph with poles u and v, we remove the edge uv and replace each e_i by G_i, for $i = 1, \dots, k$, by removing e_i and identifying the poles of G_i with the endpoints of e_i. In fact, we only allow three types of compositions: in a *series composition* the skeleton H is a cycle of length $k + 1$, in a parallel composition H consists of two vertices connected by $k + 1$ parallel edge, and in a *rigid composition* H is triconnected.

It is known that for every biconnected graph G with an edge uv the graph $G - uv$ is an uv-graph with poles u and v. The uv-graph $G - uv$ gives rise to a (de-)composition tree $\mathcal{T}$ describing how it can be obtained from single edges. Refer to Fig. 1. The nodes of $\mathcal{T}$ corresponding to edges, series, parallel, and rigid compositions of the graph are *Q*-, *S*-, *P*-, and *R-nodes*, respectively. To obtain a composition tree for G, we add an additional root Q-node representing the edge uv. To fully describe the composition, we associate with each node μ its skeleton denoted by $\text{skel}(\mu)$. For a node μ of $\mathcal{T}$, the *pertinent graph* $\text{pert}(\mu)$ is the subgraph represented by the subtree with root μ. Similarly, for a virtual edge ε of a skeleton $\text{skel}(\mu)$, the *expansion graph* of ε, denoted by $\exp(\varepsilon)$, is the pertinent graph $\text{pert}(\mu')$ of the neighbour μ' of μ corresponding to ε when considering $\mathcal{T}$ rooted at μ.

The *SPQR-tree* of G with respect to the edge uv, originally introduced by Di Battista and Tamassia [9], is the (unique) smallest decomposition tree $\mathcal{T}$ for G. Using a different edge $u'v'$ of G and a composition of $G - u'v'$ corresponds to rerooting $\mathcal{T}$ at the node representing $u'v'$. It thus makes sense to say that $\mathcal{T}$ is the SPQR-tree of G. The SPQR-tree of G has size linear in G and can be computed in linear time [13]. Planar embeddings of G correspond bijectively to planar embeddings of all skeletons of $\mathcal{T}$; the choices are the orderings of the parallel edges in P-nodes and the embeddings of the R-node skeletons, which are unique up

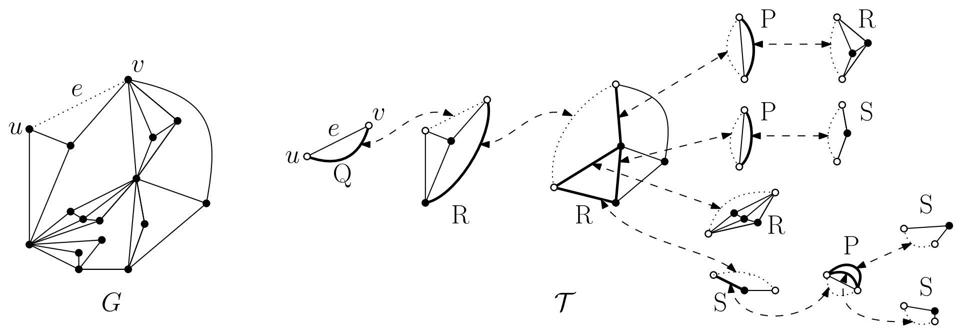

Figure 1 (left) A biconnected planar graph G and (right) the SPQR-tree $\mathcal{T}$ of G rooted at edge $e = uv$. The skeletons of all non-leaf nodes of $\mathcal{T}$ are depicted; virtual edges corresponding to edges of G are thin, whereas virtual edges corresponding to S-, P-, and R-nodes are thick. Dashed arrowed curves connect the (dotted) parent edge in the skeleton of a child node with the virtual edge representing the child node in the skeleton of its parent.

to a flip. When considering rooted SPQR-trees, we assume that the embedding of G is such that the root edge is incident to the outer face, which is equivalent to the parent edge being incident to the outer face in each skeleton. We remark that in a planar embedding of G, the poles of any node μ of $\mathcal{T}$ are incident to the outer face of pert(μ). Hence, in the following we only consider such embeddings.

Let μ be a node of $\mathcal{T}$ with poles u and v. We assume that edge uv is part of skel(μ) and pert(μ). Note that, due to this addition, pert(μ) may not be a subgraph of G anymore. The outer face of a embedding of pert(μ) is the one obtained from such an embedding after removing the edge (u, v) connecting its poles.

3 Complexity

In this section we study the computational complexity of the underlying decision problem Facial $\mathcal{C}$-Cycles of Max Facial $\mathcal{C}$-Cycles, which given a biconnected planar graph G, a set $\mathcal{C}$ of simple cycles of G, and a positive integer $k \le |\mathcal{C}|$, asks whether there exists a planar embedding $\mathcal{E}$ of G such that at least k cycles in $\mathcal{C}$ are facial cycles of $\mathcal{E}$. Facial $\mathcal{C}$-Cycles is in NP, as we can guess a set $\mathcal{C}' \subseteq \mathcal{C}$ of k cycles and then check in polynomial time (in $|G| + |\mathcal{C}|$) whether an embedding of G exists in which all cycles in $\mathcal{C}'$ are facial. We show NP-hardness for general graphs and for series-parallel graphs.

▶ **Theorem 1.** Facial $\mathcal{C}$-Cycles *is NP-complete, even if each cycle $C \in \mathcal{C}$*

- *intersects any other cycle in $\mathcal{C}$ in at most two vertices, and*
- *intersects at most three other cycles of $\mathcal{C}$ in more than one vertex.*

Proof sketch. We give a reduction from Maximum Independent Set in triconnected cubic planar graphs, which is NP-complete [14]. Let H be a triconnected cubic planar graph. Observe that H has a unique combinatorial embedding up to a flip [18]. We construct an instance $\langle G, \mathcal{C}, k\rangle$ of Facial $\mathcal{C}$-Cycles as follows; see Fig. 2. Take the planar dual $H^\star$ of H, which is a triangulation, and take $\mathcal{C}$ as the set of facial cycles of $H^\star$. The graph G is obtained from $H^\star$ by adding for each edge $e = uv \in E(H^\star)$ an *edge vertex* v_e with neighbors u and v. It is not hard to see that H admits an independent set of size k if and only if G admits a combinatorial embedding where k cycles in $\mathcal{C}$ are facial (see [7]). ◀

Figure 2 Illustrations for the proof of Theorem 1. (a) Graph H (black) and its planar dual H^* (red). Vertex v_1 is the only vertex in the MIS of H. (b) An embedding of graph G in which cycle C_1 (corresponding to the face of H^* that is dual to the vertex v_1 of H) bounds a face.

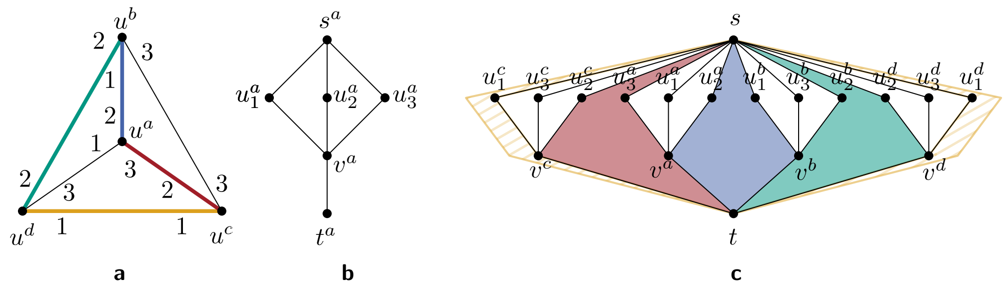

Figure 3 Illustrations for the proof of Theorem 2. (a) Cubic graph H with a Hamiltonian circuit Q (thick, colored edges). (b) Gadget G_a for a vertex $a \in V(H)$. (c) Combinatorial embedding of G corresponding to circuit Q (facial cycles have the same color as the corresponding edge in Q).

▶ **Theorem 2.** FACIAL $\mathcal{C}$-CYCLES *is NP-complete for series-parallel graphs, even if any two cycles in $\mathcal{C}$ share at most three vertices.*

Proof sketch. We reduce from HAMILTONIAN CIRCUIT, which is known to be NP-complete even for cubic graphs [11]. Let H be any such a graph.

Each vertex $a \in V(H)$ is represented by a gadget G_a consisting of

1. the graph $K_{2,3}$, where the vertices in the partition of size 2 are denoted s^a and v^a and the other vertices are denoted u_1^a, u_2^a, u_3^a, and of
2. an additional vertex t^a adjacent to v^a; see Fig. 3b.

To define graph G, we merge the vertices s_a into a single vertex s and the vertices t_a into a single vertex t. To define $\mathcal{C}$, we number the incident edges of each vertex of H from 1 to 3. If ab is the i-th edge for a and the j-th edge for b, we define $C_{ab} \in \mathcal{C}$ as the cycle $(s, u_i^a, v^a, t, v^b, u_j^b, s)$; see Fig. 3a and 3c. We claim that G admits a combinatorial embedding with $|V(H)|$ facial cycles in $\mathcal{C}$ if and only if H is Hamiltonian.

If Q is a Hamiltonian circuit of H, we embed G such that the order of the gadgets G_a is the same as the order of the vertices in Q. We then choose embeddings of the gadgets such that, for each edge ab of Q, the cycle C_{ab} bounds the face between G_a and G_b; this yields the claimed number of facial cycles in $\mathcal{C}$. Conversely, observe that if C_{ab} is a facial cycle of an embedding of G, then G_a and G_b, where ab is an edge of H, must be consecutive in the circular order around s. If G has $|V(H)|$ facial cycles in $\mathcal{C}$, it follows that the vertices corresponding to the gadgets form a Hamiltonian circuit in this order. ◀

4 Polynomial-time Solvable Cases

In this section we discuss special cases of MAX FACIAL $\mathcal{C}$-CYCLES that admit a polynomial-time solution. In particular, we show that strengthening any of the conditions in Theorem 1

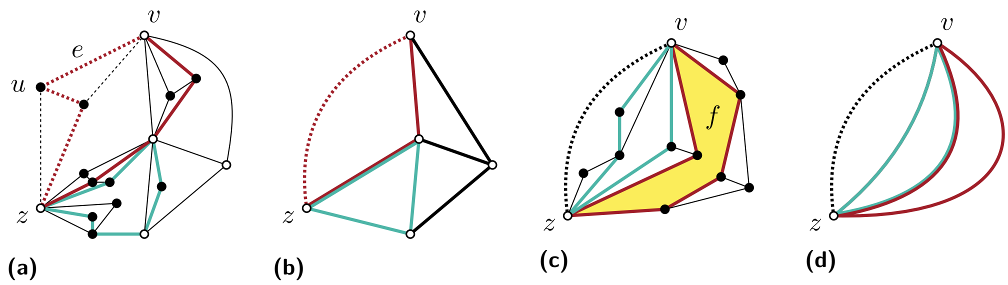

Figure 4 (b) Skeleton of the R-node μ of the SPQR-tree $\mathcal{T}$ rooted at edge $e = uv$ of the graph depicted in (a); see Fig. 1 for an illustration of tree $\mathcal{T}$. The green and red cycles in (a) are relevant for μ as they project to the green and red cycle in (b), respectively. The red cycle is an interface cycle. Dotted edges and dotted virtual edges are associated with the parent of μ. (c) Pertinent graph of the P-node depicted in (d) with three virtual edges corresponding to children from left to right realizing none, the green and the red, and only the red cycle, respectively. The red cycle bounds face f in (c) since, in addition, the second and third child are adjacent in the embedding of the skeleton.

or Theorem 2 makes the problem tractable.

4.1 General Planar Graphs

We study Max Facial $\mathcal{C}$-Cycles when each cycle in $\mathcal{C}$ intersects at most two other cycles in $\mathcal{C}$ in more than one vertex. In this setting, we give in Theorem 9 a quadratic-time algorithm for biconnected planar graphs. For series-parallel graphs we present in Theorem 10 an FPT algorithm with respect to the maximum number of cycles in $\mathcal{C}$ sharing two or more vertices with any cycle in $\mathcal{C}$. We remark that our algorithms imply that strengthening *any* of the two conditions of Theorem 1 results in a polynomial-time solvable problem. In particular, Max Facial $\mathcal{C}$-Cycles is polynomial-time solvable if any two cycles in $\mathcal{C}$ share at most one vertex.

We compute the optimal solution in these cases by a dynamic program that works bottom-up in the SPQR-tree $\mathcal{T}$ of G. Let μ be a node of $\mathcal{T}$. We call a cycle $C \in \mathcal{C}$ *relevant* for μ (or for skel(μ)) if it *projects* to a cycle in skel(μ), that is, the vertices of C in skel(μ) and the edges of skel(μ) that contain vertices or edges in C form a cycle C' in skel(μ) with at least two edges. The cycle C' is the *projection* of the cycle C in skel(μ). Similarly, we also define the projection of a cycle $C \in \mathcal{C}$ to pert(μ). The cycle C is an *interface cycle* of μ if its projection C' contains the parent edge of skel(μ). Refer to Figs. 4a and 4b. We denote the set of relevant cycles and of interface cycles of a node μ by $\mathcal{R}(\mu)$ and by $\mathcal{I}(\mu)$, respectively. Clearly, $\mathcal{I}(\mu) \subseteq \mathcal{R}(\mu)$. We denote $I(\mu) = \{X \subseteq \mathcal{I}(\mu) \mid |X| \leq 2\}$ as the set of possible interfaces. Let μ be a node of $\mathcal{T}$. We have the following two important observations.

▶ **Observation 3.** *If each cycle in $\mathcal{C}$ intersects at most two other cycles in $\mathcal{C}$ in more than one vertex, then $|\mathcal{I}(\mu)| \leq 3$.*

▶ **Observation 4.** *In any combinatorial embedding $\mathcal{E}$ of G at most two interface cycles of μ can simultaneously bound a face in $\mathcal{E}$.*

Observation 3 holds since all interface cycles of a node μ share at least the poles of μ. Observation 4 holds since each interface cycle can only bound one of the two faces incident to the virtual edge representing the parent of μ in skel(μ).

Thus, to the rest of G, the only relevant information about a combinatorial embedding $\mathcal{E}_\mu$ of pert(μ) is

(a) the number of facial cycles in $\mathcal{C}$ that bound a face of $\mathcal{E}_\mu$ and

(b) the set of cycles in $\mathcal{C}$ that project to the facial cycles incident to the parent edge of $\mathrm{pert}(\mu)$.

The reason for (a) is that cycles of $\mathcal{C}$ that are facial cycles of $\mathcal{E}_\mu$ not incident to the parent edge will be facial cycles of any embedding of G where the embedding of $\mathrm{pert}(\mu)$ is $\mathcal{E}_\mu$. So it suffices to track their number rather than which cycles are facial. For (b) observe that only those cycles that project to a face incident to $\mathcal{E}_\mu$ can potentially be realized by any embedding of G where the embedding of $\mathrm{pert}(\mu)$ is $\mathcal{E}_\mu$. We thus have to keep track of them. However, by Observation 4, at most two of them can eventually become facial cycles, and hence it suffices to consider any combination of at most two cycles for this interface.

If $\mathcal{E}$ is a combinatorial embedding of $\mathrm{pert}(\mu)$ and the elements of $I \in I(\mu)$ project to distinct faces incident to the parent edge in $\mathrm{pert}(\mu)$, we say that $\mathcal{E}$ *realizes* I; see Figs. 4c and 4d.

For any node μ and any set $I \in I(\mu)$, we denote by $T[\mu, I]$ the maximum number k such that there exists a combinatorial embedding $\mathcal{E}$ of $\mathrm{pert}(\mu)$ that realizes I and such that k cycles in $\mathcal{C}$ bound a face of $\mathcal{E}$ that is not incident to the parent edge of $\mathrm{pert}(\mu)$. If no such embedding exists, we set $T[\mu, I] = -\infty$. Due to Observation 4, for convenience we extend the definition of T to the case in which the size of I is larger than 2; in this case, we define $T[\mu, I] = -\infty$.

We show how to compute the entries of T in a bottom-up fashion in the SPQR-tree $\mathcal{T}$ of G. It is not hard to modify the dynamic program to additionally output a corresponding combinatorial embedding of G. We root $\mathcal{T}$ at an arbitrary Q-node ρ. Let ϕ be the unique child of ρ. Note that the maximum number of facial cycles in $\mathcal{C}$ for any combinatorial embedding of G is $\max_{I \in I(\phi)} |I| + T[\phi, I]$. For any leaf Q-node μ, we have that $T[\mu, I] = 0$ for each $I \in I(\mu)$. The following lemmata deal with the different types of inner nodes in an SPQR-tree.

▶ **Lemma 5.** *Let μ be an S-node with children μ_i, $i = 1, \dots, k$. Then, $T[\mu, I] = \sum_{i=1}^{k} T[\mu_i, I]$, for $I \in I(\mu)$. Also, each entry $T[\mu, I]$ can be computed in $O(k)$ time.*

Proof. The lemma follows easily from the observation that a combinatorial embedding of $\mathrm{pert}(\mu)$ realizes I if and only if each of its children realizes I. ◀

▶ **Lemma 6.** *Let μ be a P-node with children $\mu_1, \dots, \mu_k$. Then*

$$T[\mu, I] = \max_{I \subseteq C \subseteq \mathcal{R}(\mu)} \left(\sum_{i=1}^{k} T[\mu_i, C_{\mu_i}] + f(C)\right),$$

where (i) $C_{\mu_i} = C \cap \mathcal{I}(\mu_i)$ and (ii) $f(C) = |C \setminus I|$ if $\mathrm{skel}(\mu)$ admits a planar embedding $\mathcal{E}$ such that (a) each two virtual edges e_i and e_j corresponding to children μ_i and μ_j of μ, respectively, such that $|C_{\mu_i} \cap C_{\mu_j}| = 1$ are adjacent in $\mathcal{E}$, and where (b) the virtual edges e' and e'' corresponding to the children μ' and μ'' of μ such that $C'_\mu \cap I \neq \emptyset$ and $C''_\mu \cap I \neq \emptyset$, respectively, are incident to the outer face of $\mathcal{E}$, and $f(C) = -\infty$ otherwise.

Proof. Consider an embedding of $\mathrm{pert}(\mu)$ that embeds $T[\mu, I]$ cycles of $\mathcal{C}$ as facial cycles and the corresponding embedding $\mathcal{E}$ of $\mathrm{skel}(\mu)$. Let $C \subseteq \mathcal{R}(\mu)$ denote the set of cycles in $\mathcal{C}$ that are facial cycles in $\mathcal{E}$ or that are in I. Obviously, to make a cycle $c \in C \setminus I$ a facial cycle, each of the two children of μ that contain c in their interface (i) must be adjacent in $\mathcal{E}$ and (ii) must both realize cycle c. Also, in order for the cycles in I to bound the outer face of the embedding of $\mathrm{pert}(\mu)$, the two children of μ containing such interface cycles (i) must be incident to the outer face of $\mathcal{E}$ and (ii) must each realize one of these cycles in their interface. Hence $T[\mu, C]$ is a lower bound on the number of facial cycles in $\mathcal{C}$ in the embedding of $\mathrm{pert}(\mu)$. On the other hand, it is not hard to see that by picking the maximum over all subsets $C \subseteq \mathcal{R}(\mu)$ this bound is attained. ◀

We note that the existence of a corresponding embedding for a P-node μ with k children can be tested in $O(k)$ time for any set $C \subseteq \mathcal{R}(\mu)$, thus allowing us to evaluate $f(C)$ efficiently as follows. Consider the auxiliary multigraph O that contains a vertex for each virtual edge of $\text{skel}(\mu)$, except for the edge representing the parent of μ, and two such edges are adjacent if and only if there is a cycle in $C \setminus I$ that contains edges from both expansion graphs. Also, if there exist two virtual edges in $\text{skel}(\mu)$ containing edges from cycles in I, multigraph O contains an edge between them. A corresponding embedding exists if and only if O is either a simple cycle or it is a collection of paths. In latter case, O can be augmented to a simple cycle and the order of the virtual edges along this cycle defines a suitable embedding of $\text{skel}(\mu)$.

Generally, the size of $\mathcal{R}(\mu)$ can be large. However, if every cycle $C \in \mathcal{C}$ shares two or more vertices with at most r other cycles in $\mathcal{C}$, the running time can be bounded as follows.

▶ **Lemma 7.** *Let μ be a P-node with children $\mu_1, \dots, \mu_k$ such that any cycle of $\mathcal{R}(\mu)$ shares two or more vertices with at most r other cycles in $\mathcal{R}(\mu)$. For each set $I \in I(\mu)$, table $T[\mu, I]$ can be computed in $O(r^2 2^r \cdot k)$ time from $T[\mu_i, \cdot]$ with $i = 1, \dots, k$.*

Proof. We employ Lemma 6. It is $|\mathcal{R}(\mu)| \le r+1$, and $|I(\mu)| = O(r^2)$. For each $I \in I(\mu)$ we need to consider all the sets $C \subseteq \mathcal{R}(\mu)$ such that $I \subseteq C$. There are $O(2^r)$ such sets C and for each of them we evaluate $f(C)$ in $O(k)$ time. ◀

We now deal with *R*-nodes. Let μ be an R-node. Note that the instance in the hardness of Theorem 1 is an R-node whose children are a parallel of an edge and a path of length 2. If, however, any cycle in $\mathcal{C}$ shares two or more vertices with at most two other cycles from $\mathcal{C}$, then the subgraph of the dual of $\text{skel}(\mu)$ induced by the faces that are projections of cycles in $\mathcal{R}(\mu)$ consists of paths and cycles. We exploit the fact that these graphs have maximum degree 2 to give an efficient algorithm via dynamic programming.

▶ **Lemma 8.** *Let μ be an R-node with children $\mu_1, \dots, \mu_k$. There is an $O(k^2)$-time algorithm for computing $T[\mu, \cdot]$ from $T[\mu_i, \cdot]$ for $i = 1, \dots, k$, provided that cycles in $\mathcal{C}$ share two or more vertices with at most two other cycles from $\mathcal{C}$.*

Altogether, Lemmas 5, 7, and 8 imply the following theorem.

▶ **Theorem 9.** MAX FACIAL $\mathcal{C}$-CYCLES *can be solved in $O(n^2)$ time if every cycle in $\mathcal{C}$ intersects at most two other cycles in more than one vertex.*

4.2 Series-Parallel Graphs

In this section we consider MAX FACIAL $\mathcal{C}$-CYCLES on series-parallel graphs. Combining the results from Lemma 5 and Lemma 7 yields the following.

▶ **Theorem 10.** MAX FACIAL $\mathcal{C}$-CYCLES *is solvable in $O(r^2 2^r \cdot n)$ time for series-parallel graphs if any cycle in $\mathcal{C}$ intersects at most r other cycles in two or more vertices.*

▶ **Corollary 11.** MAX FACIAL $\mathcal{C}$-CYCLES *is solvable in $O(n)$ time for series-parallel graphs if any cycle in $\mathcal{C}$ intersects at most two other cycles.*

In the following we show that MAX FACIAL $\mathcal{C}$-CYCLES can be solved in polynomial time for series-parallel graphs if any two cycles in $\mathcal{C}$ share at most two vertices. The next lemma shows the special structure of relevant cycles in P-nodes of the SPQR-tree in this case.

▶ **Lemma 12.** *Let G be a series-parallel graph and $\mathcal{C}$ be a set of cycles in G such that any two cycles share at most two vertices. For each P-node μ any two cycles in $\mathcal{C}$ that are relevant for μ are either edge-disjoint in* $\text{skel}(\mu)$ *or they share the unique virtual edge of* $\text{skel}(\mu)$ *that corresponds to a Q-node child of μ, if any.*

Proof. Let C and C' be two relevant cycles for some P-node μ with poles u and v. Clearly C and C' share the two poles u and v. Now assume that C and C' additionally share a virtual edge e of skel(μ). Consider the expansion graph G_e of e and observe that $\{u, v\}$ cannot be a separation pair of G_e, since μ is a P-node. Thus the corresponding child ν of μ must be either a Q- or an S-node. If it is an S-node, however, then G_e contains a cutvertex c, which is contained in both C and C', a contradiction. Further observe that a P-node may have at most one child that is a Q-node. This concludes the proof. ◀

We again use a bottom-up traversal of the SPQR-tree of a series-parallel graph to obtain the following theorem. The S-nodes are handled using Lemma 5 and the structural properties guaranteed by Lemma 12 allow for a simple handling of the P-nodes.

▶ **Theorem 13.** Max Facial $\mathcal{C}$-Cycles *is solvable in $O(n)$ for series-parallel graphs if any two cycles in $\mathcal{C}$ share at most two vertices.*

5 Approximation Algorithms

In this section we derive constant-factor approximations for Max Facial $\mathcal{C}$-Cycles in series-parallel graphs and in biconnected planar graphs. Again, we use dynamic programming on the SPQR-tree. This time, however, instead of computing $T[\mu, I]$, we compute an approximate version $\tilde{T}[\mu, I]$ of it. A table $\tilde{T}[\mu, \cdot]$ is a *c-approximation* of $T[\mu, \cdot]$ if $1/c \cdot T[\mu, I] \leq \tilde{T}[\mu, I] \leq T[\mu, I]$, for all $I \in I(\mu)$. For P-nodes, we give an algorithm that approximates each entry within a factor of 2, for R-nodes, we achieve an approximation ratio of $(4 + \varepsilon)$ for any $\varepsilon > 0$. In the following lemmas we deal separately with S-, P-, and R-nodes.

▶ **Lemma 14.** *Let μ be an S-node with children $\mu_1, \dots, \mu_k$ and let $\tilde{T}[\mu_i, I]$ be a c-approximation of $T[\mu_i, I]$ for $i = 1, \dots, k$. Then, $\tilde{T}[\mu, I] = \sum_{i=1}^{k} \tilde{T}[\mu_i, I]$ is a c-approximation of $T[\mu, I]$.*

Proof. To see this, observe that by Lemma 5, it is $1/c \cdot T[\mu, I] = 1/c \cdot \sum_{i=1}^{k} T[\mu_i, I] \leq \sum_{i=1}^{k} \tilde{T}[\mu_i, I]$ and $\sum_{i=1}^{k} \tilde{T}[\mu_i, I] \leq \sum_{i=1}^{k} T[\mu_i, I] = T[\mu, I]$. ◀

Next we deal with a P-node μ with children $\mu_1, \dots, \mu_k$. The algorithm works as follows. Fix a set $I \in I(\mu)$. We construct an auxiliary graph H as follows. The vertices of H are the children $\mu_1, \dots, \mu_k$ of μ. Two vertices μ_i and μ_j are adjacent in H if and only if there exists a cycle $C \in \mathcal{C}$ that intersects μ_i and μ_j such that $\tilde{T}[\mu_x, (I \cup \{C\}) \cap \mathcal{I}(\mu_x)] = \tilde{T}[\mu_x, I \cap \mathcal{I}(\mu_x)]$ for $x \in \{i, j\}$, i.e., according to the approximate table $\tilde{T}$ additionally realizing C in the interface of the children μ_i and μ_j does not decrease the number of facial cycles of pert(μ_i) in $\mathcal{C}$. If $|I| = 2$, assume that μ_1 and μ_2 are the two children intersected by the cycles in I. Unless μ_1 and μ_2 are the only children of μ, we remove the edge $\mu_1\mu_2$ from H if it is there. This reflects the fact that, due to the restrictions imposed by I, it is not possible to realize a corresponding cycle. Now compute a maximum matching M in H. The matching M corresponds to a set $C_M \subseteq \mathcal{R}(\mu)$ of relevant cycles of μ that are pairwise edge-disjoint. We set $\tilde{T}[\mu, I] = \sum_{i=1}^{k} \tilde{T}[\mu_i, (I \cup C_M) \cap \mathcal{I}(\mu_i)] + |M| = \sum_{i=1}^{k} \tilde{T}[\mu_i, I \cap \mathcal{I}(\mu_i)] + |M|$. We claim that this gives a $\max\{2, c\}$-approximation of $T[\mu, \cdot]$ if the $\tilde{T}[\mu_i, \cdot]$ are c-approximations of $T[\mu_i, \cdot]$.

▶ **Lemma 15.** *Let μ be a P-node an let $\tilde{T}[\mu, \cdot]$ be the table computed in the above fashion. Then, $\tilde{T}[\mu, \cdot]$ is a $\max\{2, c\}$-approximation of $T[\mu, \cdot]$ if $\tilde{T}[\mu_i, \cdot]$ is a c-approximation of $T[\mu_i, \cdot]$.*

Proof. We first show that $\tilde{T}[\mu, \cdot] \leq T[\mu, \cdot]$. To this end, it suffices to show that, for any $I \in I(\mu)$, there exists an embedding of pert(μ) that realizes I and has $\tilde{T}[\mu, I]$ facial cycles in $\mathcal{C}$. Consider the multigraph with vertex set $\{\mu_1, \dots, \mu_k\}$ and edge set $C_M \cup I$. This graph has maximum degree 2 and, due to the special treatment of the edge $\mu_1\mu_2$, unless $k = 2$, none of its connected components is a cycle. We can thus always complete this graph into a cycle containing all μ_i, which defines a circular order of $\mu_1, \dots, \mu_k$, and hence an embedding of skel(μ). In this embedding, all the cycles in $C_M \cup I$ project to facial cycles. Realizing all these cycles yields $\tilde{T}[\mu, I] = \sum_{i=1}^{k} \tilde{T}[\mu_i, (I \cup C_M) \cap \mathcal{I}(\mu_i)] + |M| \leq \sum_{i=1}^{k} T[\mu_i, (I \cup C_M) \cap \mathcal{I}(\mu_i)] + |M|$ realized cycles. By the definition of the $T[\mu_i, \cdot]$ we get embeddings for the pert(μ_i) with a corresponding number of cycles in $\mathcal{C}$. Combining them according to the embedding of skel(μ) from above, yields an embedding of pert(μ) that realizes I and has at least $\tilde{T}[\mu, I]$ facial cycles in $\mathcal{C}$. Hence $\tilde{T}[\mu, I] \leq T[\mu, I]$.

Conversely, consider $T[\mu, I]$ and a corresponding embedding of skel(μ). Denote by C_{opt} the set of facial cycles in $\mathcal{C}$ in an optimal solution that project to facial cycles of skel(μ). We consider two cycles in C_{opt} as adjacent if they intersect the same child of μ. Clearly, each child μ_i is intersected by at most two cycles in C_{opt} and, moreover, the two faces of skel(μ) incident to the parent edge are not realized. Hence the corresponding graph is a collection of paths, and it can be edge-colored with two colors. Let C'_{opt} be the cycles in the larger color class. We have $|C'_{\mathrm{opt}}| \geq |C_{\mathrm{opt}}|/2$ and no two distinct cycles in C'_{opt} intersect the same child μ_i of μ, i.e., interpreting the cycles in C'_{opt} as edges on the vertex set $\{\mu_1, \dots, \mu_k\}$ yields a matching M'. We would like to argue that our matching M in the auxiliary graph H is larger than M', and hence we realize at least half of the cycles of the optimum. However, this argument is not valid, since M' may contain edges that are not present in H due to approximation errors in the $\tilde{T}[\mu_i, \cdot]$. We will show that the contribution of these edges is irrelevant and hence the intuition about comparing the matching sizes indeed applies.

Let $M'_1 = M' \setminus E(H)$ and $M'_2 = M' \cap E(H)$. Let $J = \{1, \dots, k\}$ and let $J_1 = \{i \in J \mid \exists C \in M'_1$ that intersects $\mu_i\}$ be the indices of children that are intersected by a cycle in M'_1. The set $J_2 = J \setminus J_1$ contains the remaining indices.

Clearly, we have $T[\mu, I] = \sum_{i=1}^{k} T[\mu_i, (I \cup C_{\mathrm{opt}}) \cap \mathcal{I}(\mu_i)] + |C_{\mathrm{opt}}|$ according to Lemma 6. Realizing instead of C_{opt} just the set of cycles $C_{M'} = C'_{\mathrm{opt}}$ corresponding to M' drops at most $|C_{\mathrm{opt}}|/2$ facial cycles in $\mathcal{C}$, while imposing weaker interface constraints on the children. We therefore have

$$T[\mu, I] = \sum_{i=1}^{k} T[\mu_i, (I \cup C_{\mathrm{opt}}) \cap \mathcal{I}(\mu_i)] + |C_{\mathrm{opt}}| \leq \sum_{i=1}^{k} T[\mu_i, (I \cup C_{M'}) \cap \mathcal{I}(\mu_i)] + 2|M'|$$

We now use the fact that the $\tilde{T}[\mu_i, \cdot]$ are c-approximations of the $T[\mu_i, \cdot]$, and hence also c'-approximations for $c' = \max\{c, 2\}$, and we also separate the sum by the index set J_1 and J_2 and consider the two matchings M'_1 and M'_2 separately.

$$\begin{aligned} \sum_{i=1}^{k} T[\mu_i, (I \cup C_{M'}) \cap \mathcal{I}(\mu_i)] + 2|M'| \leq c' \cdot \sum_{i \in J_1} \tilde{T}[\mu_i, (C_{M'_1} \cup I) \cap \mathcal{I}(\mu_i)] + 2|M'_1| \\ + c' \cdot \sum_{i \in J_2} \tilde{T}[\mu_i, (C_{M'_2} \cup I) \cap \mathcal{I}(\mu_i)] + 2|M'_2|. \end{aligned} \tag{1}$$

Observe that the indices of the children intersected by cycles that form a matching M_2 in H are all contained in J_2. By the definition of H, we have $\tilde{T}[\mu_i, (C_{M'_2} \cup I) \cap \mathcal{I}(\mu_i)] = \tilde{T}[\mu_i, I \cap \mathcal{I}(\mu_i)]$, for $i \in J_2$.

For the first term, observe that, for each edge $\mu_i\mu_j \in M'_1$, we have $\tilde{T}[\mu_x, (M'_1 \cup I) \cap \mathcal{I}(\mu_x)] \leq \tilde{T}[\mu_x, I \cap \mathcal{I}(\mu_x)] - 1$ for at least one $x \in \{i, j\}$. Otherwise the edge would be in H, and hence in M'_2. Let $J'_1 \subseteq J_1$ denote the set of indices where this happens and let $J''_1 = J_1 \setminus J'_1$.

Observe that $|J'_1| \geq |M'_1|$. We thus have

$$\begin{aligned}
&c' \cdot \sum_{i \in J_1} \tilde{T}[\mu_i, (C_{M'_1} \cup I) \cap \mathcal{I}(\mu_i)] + 2|M'_1| \\
&= c' \cdot \sum_{i \in J'_1} \tilde{T}[\mu_i, (C_{M'_1} \cup I) \cap \mathcal{I}(\mu_i)] + c' \cdot \sum_{i \in J''_1} \tilde{T}[\mu_i, (C_{M'_1} \cup I) \cap \mathcal{I}(\mu_i)] + 2|M'_1| \\
&\leq c' \cdot \sum_{i \in J'_1} (\tilde{T}[\mu_i, I \cap \mathcal{I}(\mu_i)] - 1) + c' \cdot \sum_{i \in J''_1} \tilde{T}[\mu_i, I \cap \mathcal{I}(\mu_i)] + 2|M'_1| \\
&\leq c' \cdot \left(\sum_{i \in J_1} \tilde{T}[\mu_i, I \cap \mathcal{I}(\mu_i)] - |J'_1| \right) + 2|M'_1| \leq c' \cdot \sum_{i \in J_1} \tilde{T}[\mu_i, I \cap \mathcal{I}(\mu_i)],
\end{aligned}$$

where the last step uses the fact that $c' \geq 2$. Plugging this information into Eq. 1, yields

$$\begin{aligned}
&c' \cdot \sum_{i \in J_1} \tilde{T}[\mu_i, (C_{M'_1} \cup I) \cap \mathcal{I}(\mu_i)] + 2|M'_1| + c' \cdot \sum_{i \in J_2} \tilde{T}[\mu_i, (C_{M'_2} \cup I) \cap \mathcal{I}(\mu_i)] + 2|M'_2| \\
&\leq c' \cdot \sum_{i=1}^{k} \tilde{T}[\mu_i, I \cap \mathcal{I}(\mu_i)] + 2|M'_2| \leq c' \cdot \sum_{i=1}^{k} \tilde{T}[\mu_i, I \cap \mathcal{I}(\mu_i)] + 2|M| \\
&\leq c' \cdot \left(\sum_{i=1}^{k} \tilde{T}[\mu_i, I \cap \mathcal{I}(\mu_i)] + |M| \right) = c' \left(\sum_{i=1}^{k} \tilde{T}[\mu_i, (I \cup C_M) \cap \mathcal{I}(\mu_i)] + |M| \right)
\end{aligned}$$

The last three steps use the facts that $M \subseteq E(H)$ is a maximum matching, and hence larger than M'_2, that $c' \geq 2$, and that $C_M \subseteq E(H)$, respectively. ◀

We note that the bottleneck for computing $T[\mu, I]$ is finding a maximum matching in a graph with $O(|\operatorname{skel}(\mu)|)$ vertices and $O(|\mathcal{C}|)$ edges. Hence the running time for one step is $O(|\operatorname{skel}(\mu)| + \sqrt{|\operatorname{skel}(\mu)|} \cdot |\mathcal{C}|)$. Since $|I(\mu)| \leq |\mathcal{C}|^2$, the running time for processing a single P-node μ is $O(|\operatorname{skel}(\mu)||\mathcal{C}|^2 + \sqrt{|\operatorname{skel}(\mu)|} \cdot |\mathcal{C}|^3)$. The total time for processing all P-nodes is $O(n|\mathcal{C}|^2 + \sqrt{n}|\mathcal{C}|^3)$.

▶ **Theorem 16.** *There is a* 2*-approximation algorithm with running time* $O(n|\mathcal{C}|^2 + \sqrt{n}|\mathcal{C}|^3)$ *for* Max Facial $\mathcal{C}$-Cycles *in series-parallel graphs.*

Next we deal with R-nodes. Let μ be an R-node with children $\mu_1, \dots, \mu_k$. For each face f of $\operatorname{skel}(\mu)$ let J_f denote the indices of the children μ_i whose corresponding virtual edge in $\operatorname{skel}(\mu)$ is incident to f.

Fix $I \in \mathcal{I}(\mu)$. We propose the following algorithm for computing $\tilde{T}[\mu, I]$. Consider the subgraph H of the dual of $\operatorname{skel}(\mu)$ induced by those vertices v corresponding to a face f not incident to the parent edge of $\operatorname{skel}(\mu)$ and such that there exists a cycle $C_v \in \mathcal{C}$ that projects to the boundary of f and such that $\tilde{T}[\mu_i, (\{C_v\} \cup I) \cap \mathcal{I}(\mu)] = \tilde{T}[\mu_i, I \cap \mathcal{I}(\mu)]$, i.e., requiring that C_v is realized in μ_i does not change the approximate number of faces of $\operatorname{pert}(\mu_i)$ in $\mathcal{C}$.

Now we compute a $(1 + \varepsilon/4)$-approximation of a maximum independent set of H, which can be done in time polynomial in $|\operatorname{skel}(\mu)|$ (and exponential in $(1/\varepsilon)$) [3]. Let X denote this independent set, and let $C_X = \{C_v \mid v \in X\}$ be a set of corresponding cycles in $\mathcal{C}$. We set $\tilde{T}[\mu, I] = \sum_{i=1}^{k} \tilde{T}[\mu_i, (I \cup X) \cap \mathcal{I}(\mu_i)] + |X| = \sum_{i=1}^{k} \tilde{T}[\mu_i, I \cap \mathcal{I}(\mu_i)] + |X|$, and claim that in this fashion $\tilde{T}[\mu, \cdot]$ is a $\max\{c, (4 + \varepsilon)\}$-approximation provided that $\tilde{T}[\mu_i, \cdot]$ is a c-approximation of $T[\mu_i, \cdot]$. The proof is similar to that of Lemma 15. It 4-colors the facial cycles $C_{\text{opt}} \subset \mathcal{C}$ of an optimal solution and considers the largest color class, which is an independent set of faces that has size at least $|C_{\text{opt}}|/4$.

▶ **Lemma 17.** *Let $\tilde{T}[\mu, \cdot]$ be the table computed in the above fashion. Then, $\tilde{T}[\mu, \cdot]$ is a $\max\{c, (4+\varepsilon)\}$-approximation of $T[\mu, \cdot]$ provided that $\tilde{T}[\mu_i, \dots]$ is a c-approximation of $T[\mu_i, \cdot]$.*

Overall, we obtain the following theorem.

▶ **Theorem 18.** Max Facial $\mathcal{C}$-Cycles *for biconnected planar graphs admits an efficient $(4+\varepsilon)$-approximation algorithm for any $\varepsilon > 0$.*

6 Conclusions

In this paper, we explored the boundaries of the computational complexity of Max Facial $\mathcal{C}$-Cycles. In particular, we proved the problem NP-hard under restrictive conditions, showed that slightly stronger conditions make the problem tractable, and gave constant-factor approximations for series-parallel and biconnected planar graphs with approximation guarantees of 2 and $4+\varepsilon$ for any $\varepsilon > 0$, respectively. Our main open question is whether these approximation guarantees may be improved.

References

1 Patrizio Angelini, Giuseppe Di Battista, Fabrizio Frati, Vít Jelínek, Jan Kratochvíl, Maurizio Patrignani, and Ignaz Rutter. Testing Planarity of Partially Embedded Graphs. *ACM Trans. Alg.*, 11(4):32:1–32:42, 2015. doi:10.1145/2629341.

2 Patrizio Angelini, Giuseppe Di Battista, and Maurizio Patrignani. Finding a Minimum-Depth Embedding of a Planar Graph in $O(n^4)$ Time. *Algorithmica*, 60:890–937, 2011. doi:10.1007/s00453-009-9380-6.

3 Brenda S. Baker. Approximation Algorithms for NP-complete Problems on Planar Graphs. *J. ACM*, 41(1):153–180, January 1994. doi:10.1145/174644.174650.

4 Thomas Bläsius, Sebastian Lehmann, and Ignaz Rutter. Orthogonal graph drawing with inflexible edges. *Comput. Geom.*, 55:26–40, 2016. doi:10.1016/j.comgeo.2016.03.001.

5 Thomas Bläsius, Ignaz Rutter, and Dorothea Wagner. Optimal Orthogonal Graph Drawing with Convex Bend Costs. *ACM Trans. Algorithms*, 12(3):33:1–33:32, 2016.

6 Giordano Da Lozzo, Vít Jelínek, Jan Kratochvíl, and Ignaz Rutter. Planar Embeddings with Small and Uniform Faces. In *ISAAC'14*, volume 8889 of *LNCS*, pages 633–645. Springer, 2014. doi:10.1007/978-3-319-13075-0.

7 Giordano Da Lozzo and Ignaz Rutter. On the complexity of realizing facial cycles. *CoRR*, abs/1607.02347, 2016. arXiv:1607.02347.

8 Giuseppe Di Battista and Fabrizio Frati. Drawing Trees, Outerplanar Graphs, Series-Parallel Graphs, and Planar Graphs in a Small Area. In *Thirty Essays on Geometric Graph Theory*, pages 121–165. Springer New York, 2013. doi:10.1007/978-1-4614-0110-0_9.

9 Giuseppe Di Battista and Roberto Tamassia. On-Line Graph Algorithms with SPQR-Trees. In Michael S. Paterson, editor, *ICALP'90*, volume 443 of *LNCS*, pages 598–611. Springer, 1990. doi:10.1007/BFb0032061.

10 Christoph Dornheim. Planar Graphs with Topological Constraints. *JGAA*, 6(1):27–66, 2002. doi:10.7155/jgaa.00044.

11 M. R. Garey, David S. Johnson, and Robert Endre Tarjan. The Planar Hamiltonian Circuit Problem is NP-Complete. *SIAM J. Comput.*, 5(4):704–714, 1976. doi:10.1137/0205049.

12 Ashim Garg and Roberto Tamassia. On the Computational Complexity of Upward and Rectilinear Planarity Testing. *SIAM J. on Comput.*, 31(2):601–625, 2001. doi:10.1137/S0097539794277123.

13 Carsten Gutwenger and Petra Mutzel. A Linear Time Implementation of SPQR-trees. In Joe Marks, editor, *GD'00*, volume 1984 of *LNCS*, pages 77–90. Springer, 2001. doi:10.1007/3-540-44541-2_8.

14 Joseph Douglas Horton and Kyriakos Kilakos. Minimum Edge Dominating Sets. *SIAM J. Discrete Math.*, 6(3):375–387, 1993. doi:10.1137/0406030.

15 Frank Kammer. Determining the Smallest k Such That G Is k-Outerplanar. In *ESA'07*, volume 4698 of *LNCS*, pages 359–370. Springer, 2007. doi:10.1007/978-3-540-75520-3_33.

16 Petra Mutzel and René Weiskircher. Optimizing over All Combinatorial Embeddings of a Planar Graph. In *IPCO'99*, volume 1610 of *LNCS*, pages 361–376. Springer, 1999. doi:10.1007/3-540-48777-8_27.

17 Roberto Tamassia. On Embedding a Graph in the Grid with the Minimum Number of Bends. *SIAM J. Comput.*, 16(3):421–444, 1987. doi:10.1137/0216030.

18 Hassler Whitney. Congruent Graphs and the Connectivity of Graphs. *American Journal of Mathematics*, 54(1):150–168, 1932. URL: http://www.jstor.org/stable/2371086.

19 Gerhard J. Woeginger. Embeddings of planar graphs that minimize the number of long-face cycles. *Oper. Res. Lett.*, 30(3):167–168, 2002. doi:10.1016/S0167-6377(02)00119-0.