Simultaneous Embedding: Edge Orderings, Relative Positions, Cutvertices[⋆]

Thomas Bläsius, Annette Karrer, and Ignaz Rutter[⋆⋆]

Karlsruhe Institute of Technology (KIT)
firstname.lastname@kit.edu

Abstract. A simultaneous embedding of two graphs $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ with common graph $G = G^{\textcircled{1}} \cap G^{\textcircled{2}}$ is a pair of planar drawings of $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ that coincide on G. It is an open question whether there is a polynomial-time algorithm that decides whether two graphs admit a simultaneous embedding (problem SEFE).

In this paper, we present two results. First, a set of three linear-time preprocessing algorithms that remove certain substructures from a given SEFE instance, producing a set of equivalent SEFE instances without such substructures. The structures we can remove are (1) cutvertices of the *union graph* $G^{\textcircled{1}} \cup G^{\textcircled{2}}$, (2) cutvertices that are simultaneously a cutvertex in $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ and that have degree at most 3 in G, and (3) connected components of G that are biconnected but not a cycle.

Second, we give an $O(n^2)$-time algorithm for SEFE where, for each pole u of a P-node μ (of a block) of the input graphs, at most three virtual edges of μ contain common edges incident to u. All algorithms extend to the sunflower case.

1 Introduction

A simultaneous embedding of two graphs $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ with common graph $G = G^{\textcircled{1}} \cap G^{\textcircled{2}}$ is a pair of planar drawings of $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$, that coincide on G. The problem to decide whether a simultaneous embedding exists is called SEFE (simultaneous embedding with fixed edges). This definition extends to more than two graphs. For three graphs SEFE is NP-complete [7]. In the *sunflower case* it is required that every pair of input graphs has the same intersection. See [2] for a survey on SEFE and related problems.

There are two fundamental approaches to solving SEFE in the literature. The first approach is based on the characterization of Jünger and Schulz [10] stating that finding a simultaneous embedding of two graphs $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ with common graph G is equivalent to finding planar embeddings of $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ that induce the same embedding on G. The second very recent approach by Schaefer [11] is based on Hanani-Tutte-style redrawing results. One tries to characterize the existence of a SEFE via the existence of drawings where no two independent edges of the same graph cross an odd number of times. The existence of such drawings can be expressed using a linear system of boolean equations.

When following the first approach, we need two things to describe the planar embedding of the common graph G. First, for each vertex v, a cyclic order of incident edges

[⋆] Partly done within GRADR – EUROGIGA project no. 10-EuroGIGA-OP-003.

[⋆⋆] Work was supported by a fellowship within the Postdoc-Program of the German Academic Exchange Service (DAAD).

S. Wismath and A. Wolff (Eds.): GD 2013, LNCS 8242, pp. 220–231, 2013.

around v. Second, for every pair of connected components H and H' of G, the face f of H containing H'. We call this relationship the *relative position* of H' with respect to H. To find a simultaneous embedding, one needs to find a pair of planar embeddings that induce the same cyclic edge orderings (*consistent edge orderings*) and the same relative positions (*consistent relative positions*) on the common graph G.

Most previous results use the first approach but none of them considers both consistent edge orderings and relative positions. Most of them assume the common graph to be connected or to contain no cycles. The strongest results of this type are the two linear-time algorithms for the case that G is biconnected by Haeupler et al. [9] and by Angelini et al. [1] and a quadratic-time algorithm for the case where $G^{①}$ and $G^{②}$ are biconnected and G is connected [4]. In the latter result, SEFE is modeled as an instance of the problem SIMULTANEOUS PQ-ORDERING. On the other hand, there is a linear-time algorithm for SEFE if the common graph consists of disjoint cycles [3], which requires to ensure consistent relative positions but makes edge orderings trivially consistent.

The advantage of the second approach is that it implicitly handles both, consistent edge orderings and consistent relative positions, at the same time. Thus, the results by Schaefer [11] are the first that handle SEFE instances where the common graph consists of several, non-trivial connected components. He gives a polynomial-time algorithm for the cases where each connected component of the common graph is biconnected or has maximum degree 3. Although this approach is conceptionally simple, very elegant, and combines several notions of planarity within a common framework, it has two disadvantages. The running time of the algorithms are quite high and the high level of abstraction makes it difficult to generalize the results, e.g., to the sunflower case.

Contribution & Outline. In this paper, we follow the first approach and show how to enforce consistent edge orderings and consistent relative positions at the same time, by combining different recent approaches, namely the algorithm by Angelini et al. [1] and result on SIMULTANEOUS PQ-ORDERING [4] for consistent edge orderings and the result on disjoint cycles [3] for consistent relative positions. To handle relative positions of connected components to each other without knowing their embedding, we show that these relative positions can be expressed in terms of relative positions with respect to a cycle basis. In addition to that, we are able to handle certain cutvertices of $G^{①}$ and $G^{②}$.

More precisely, we classify a vertex v to be a *union cutvertex*, a *simultaneous cutvertex*, and an *exclusive cutvertex* if v is a cutvertex of $G^{①} \cup G^{②}$, of $G^{①}$ and $G^{②}$ but not of $G^{①} \cup G^{②}$, and of $G^{①}$ but not $G^{②}$ or the other way around, respectively. We say that v has *common-degree* k if it is a common vertex with degree k in G. We present three preprocessing algorithms that simplify given instances of SEFE; see Section 3. They remove union cutvertices and simultaneous cutvertices with common-degree 3 (note that simultaneous cutvertices with common degree less than 3 cannot exist), and replace connected components of G that are biconnected by cycles. They run in linear time and can be applied independently. The latter algorithm together with the linear-time algorithm for disjoint cycles [3] improves the result by Schaefer [11] for instances where every connected component of G is biconnected to linear time and the sunflower case.

In Section 4 we show how to solve instances that have *common P-node degree 3* and contain neither union nor simultaneous cutvertices in quadratic time. An instance has common P-node degree k if, for each pole u of a P-node μ (of a block) of the input graphs, at most k virtual edges of μ contain common edges incident to u. Together with the preprocessing steps, this includes the case where every connected component of G is biconnected, has maximum degree 3, or is outerplanar with maximum degree 3 cutvertices. As before, this approach also applies to the sunflower case.

2 Preliminaries

Connectivity & SPQR-trees. A graph is *connected* if there exists a path between any pair of vertices. A *separating k-set* is a set of k vertices whose removal disconnects the graph. Separating 1-sets and 2-sets are *cutvertices* and *separation pairs*, respectively. A connected graph is *biconnected* if it has no cut vertex and *triconnected* if it has no separation pair. The maximal biconnected components of a graph are called *blocks*. The *split components* with respect to a separating k-set are the maximal subgraphs that are not disconnected by removing the separating k-set.

The SPQR-tree $\mathcal{T}$ of a biconnected graph G represents the decomposition of G along its *split pairs*, where a split pair is either a separating pair or a pair of adjacent vertices [6]. We consider the SPQR-tree to be unrooted, representing embeddings on the sphere, i.e., planar embeddings without a designated outer face.

Let $\{s, t\}$ be a split pair and let H_1 and H_2 be two subgraphs of G such that $H_1 \cup H_2 = G$ and $H_1 \cap H_2 = \{s, t\}$. Consider the tree consisting of two nodes μ_1 and μ_2 associated with the graphs $H_1 + \{s, t\}$ and $H_2 + \{s, t\}$, respectively. These graphs are called *skeletons* of the nodes μ_i, denoted by $\text{skel}(\mu_i)$, and the special edge $\{s, t\}$ is a *virtual edge*. The edge connecting the nodes μ_1 and μ_2 associates the virtual edges in $\text{skel}(\mu_1)$ and $\text{skel}(\mu_2)$ with each other. The *expansion graph* of a virtual edge $\{s, t\}$ is the subgraph of G it represents, that is in $\text{skel}(\mu_1)$ the expansion graph of $\{s, t\}$ is H_2 and the expansion graph of $\{s, t\}$ in $\text{skel}(\mu_2)$ is H_1. A combinatorial embedding of G uniquely induces a combinatorial embedding of $\text{skel}(\mu_1)$ and $\text{skel}(\mu_2)$ and vice versa.

Applying this kind of decomposition systematically yields the SPQR-tree $\mathcal{T}$. The skeletons of the internal nodes of $\mathcal{T}$ are either a cycle (S-node), a bunch of parallel edges (P-node) or a triconnected planar graph (R-node). The leaves are Q-nodes, and their skeleton consists of two vertices connected by a virtual and a normal edge. Thus, the only possible embedding choices are flipping skeletons of R-nodes and ordering the edges in skeletons of P-nodes. The SPQR-tree can be computed in linear time [8].

Let $\mathcal{T}^{①}$ by the SPQR-tree of a block of $G^{①}$ in an instance of SEFE and let G be the common graph. Let further μ be a P-node of $\mathcal{T}^{①}$. We say that μ has *common P-node degree k* if both vertices in $\text{skel}(\mu)$ are incident to common edges in the expansion graphs of at most k virtual edges (note that these can be different edges for the two vertices). We say that $G^{①}$ has common P-node degree k if each P-node in the SPQR-tree of each block of $G^{①}$ has common P-node degree k. If this is the case for $G^{①}$ and $G^{②}$, we say that the instance of SEFE has common P-node degree k.

PQ-trees. A PQ-tree, originally introduced by Booth and Lueker [5], is a tree, whose inner nodes are either P-nodes or Q-nodes (note that these P-nodes have nothing to do with the P-nodes of the SPQR-tree). The order of edges around a P-node can be ordered arbitrarily, the edges around a Q-node are fixed up to a flip. In this way, a PQ-tree represents a set of orders on its leaves. A rooted PQ-tree represents linear orders, an unrooted

PQ-tree represents cyclic orders (in most cases we consider unrooted PQ-trees). Given a PQ-tree T and a subset S of its leaves, there exists another PQ-tree T' representing exactly the orders represented by T where the elements in S are consecutive. The tree T' is the *reduction* of T with respect to S. The *projection* of T to S is a PQ-tree with leaves S representing exactly the orders on S that are represented by T.

The problem SIMULTANEOUS PQ-ORDERING has several PQ-trees sharing some leaves as input, that are related by identifying some of their leaves [4]. More precisely, every instance is a directed acyclic graph, where each node is a PQ-tree, and each arc (T, T') has the property that there is an injective map from the leaves of the *child* T' to the leaves of the *parent* T. For each PQ-tree in such an instance, one wants to find an order of its leaves such that for every arc (T, T') the order chosen for the parent T is an extension of the order chosen for the child T' (with respect to the injective map). We will later use instances of SIMULTANEOUS PQ-ORDERING to express relations between orderings of edges around vertices.

3 Preprocessing Algorithms

In this section, we present several algorithms that can be used as a preprocessing of a given SEFE instance. The result is usually a set of SEFE instances that admit a solution if and only if the original instance admits one. The running time of the preprocessing algorithms is linear, and so is the total size of the equivalent set of SEFE instances. Each of the preprocessing algorithms removes certain types of structures form the instance, in particular from the common graph. Namely, we show that we can eliminate union cutvertices, simultaneous cutvertices with common-degree 3, and connected components of G that are biconnected but not a cycle. None of these algorithms introduces new cutvertices in G or increases the degree of a vertex. Thus, the preprocessing algorithms can be successively applied to a given instance, removing all the claimed structures.

Let $(G^{①}, G^{②})$ be a SEFE instance with common graph $G = G^{①} \cap G^{②}$. We can equivalently encode such an instance in terms of its *union graph* $G_{\cup} = G^{①} \cup G^{②}$, whose edges are labeled $\{1\}$, $\{2\}$, or $\{1, 2\}$, depending on whether they are contained exclusively in $G^{①}$, exclusively in $G^{②}$, or in G, respectively. Any graph with such an edge coloring can be considered as a SEFE instance. Since sometimes the coloring version is more convenient, we use these notions interchangeably throughout this section.

3.1 Union Cutvertices and Simultaneous Cutvertices

It is not hard to see that union cutvertices of a SEFE instance can be used to split it into independent instances. A simultaneous cutvertex with common-degree 3 can be modified as in Fig. 1, yielding an equivalent instance. Exhaustively applying these ideas, yields the following results; proofs are omitted due to space constraints.

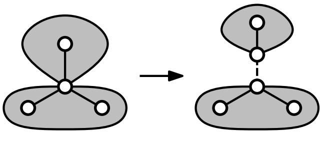

Fig. 1. A simultaneous cutvertex with common-degree 3. The gray regions are the split components of $G^{①}$, the new dashed edge belongs to $G^{②}$.

Theorem 1. *There is a linear-time algorithm that decomposes a* SEFE *instance into an equivalent set of* SEFE *instances that do not contain union cutvertices.*

Theorem 2. *Let $(G^{\textcircled{1}}, G^{\textcircled{2}})$ be an instance of* SEFE *such that every simultaneous cutvertex has common-degree 3. An equivalent instance without simultaneous cutvertices can be computed in linear time.*

3.2 Connected Components that are Biconnected

Let $(G^{\textcircled{1}}, G^{\textcircled{2}})$ be a SEFE instance. Throughout this section, we assume without loss of generality that $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ are connected [3] and that the common graph G is an induced subgraph of $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$. The latter can be achieved by subdividing each exclusive edge once, which clearly does not alter the existence of a SEFE.

Let C be a connected component of G that is a cycle. A *union bridge* of $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ with respect to C is a connected component of $G_{\cup} - C$ together with all its attachment vertices on C. Similarly, there are $\textcircled{1}$*-bridges* and $\textcircled{2}$*-bridges*, which are connected components of $G^{\textcircled{1}} - C$ and $G^{\textcircled{2}} - C$ together with their attachment vertices on C, respectively. We say that two bridges B_1 and B_2 *alternate* if there are attachments a_1, b_1 of B_1 and attachments a_2, b_2 of B_2, such that the order along C is $a_1 a_2 b_1 b_2$. We have the following lemma.

Lemma 1. *Let $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ be two planar graphs and let C be a connected component of the common graph that is a cycle. Then the graphs $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ admit a* SEFE *where C is the boundary of the outer face if and only if (i) each union bridge admits a* SEFE *together with C and (ii) no two $\textcircled{i}$-bridges of C alternate for $i = 1, 2$.*

Proof. Clearly the conditions are necessary; we prove sufficiency. Let $B_1, \dots, B_k$ be the union bridges with respect to C, and let $(\mathcal{E}_1^{\textcircled{1}}, \mathcal{E}_1^{\textcircled{2}}), \dots, (\mathcal{E}_k^{\textcircled{1}}, \mathcal{E}_k^{\textcircled{2}})$ be the corresponding simultaneous embeddings of B_i together with C, which exist by condition (i). Note that each union bridge is connected, and hence all its edges and vertices are embedded on the same side of C. After possibly flipping some of the embeddings, we may assume that each of them has C with the same clockwise orientation as the outer face.

We now glue $\mathcal{E}_1^{\textcircled{1}}, \dots, \mathcal{E}_k^{\textcircled{1}}$ to an embedding $\mathcal{E}^{\textcircled{1}}$ of $G^{\textcircled{1}}$, which is possible by condition (ii). In the same way, we find an embedding $\mathcal{E}^{\textcircled{2}}$ of $G^{\textcircled{2}}$ from $\mathcal{E}_1^{\textcircled{2}}, \dots, \mathcal{E}_k^{\textcircled{2}}$. We claim that $(\mathcal{E}_1, \mathcal{E}_2)$ is a SEFE of $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$. For the consistent edge orderings, observe that any common vertex v with common-degree at least 3 is contained, together with all neighbors, in some union bridge B_i. The compatibility of the edge ordering follows since $(\mathcal{E}_i^{\textcircled{1}}, \mathcal{E}_i^{\textcircled{2}})$ is a SEFE. Concerning the relative position of a vertex v and some common cycle C', we note that the relative positions clearly coincide in $\mathcal{E}^{\textcircled{1}}$ and $\mathcal{E}^{\textcircled{2}}$ for $C = C'$. Otherwise C' is contained in some union bridge. If v is embedded in the interior of C' in one of the two embeddings, then it is contained in the same union bridge as C', and the compatibility follows. If this case does not apply, it is embedded outside of C' in both embeddings, which is compatible as well. □

Now consider a connected component C of the common graph G of a SEFE instance such that C is biconnected. Such a component is called *2-component*. If C is a cycle, it is a *trivial* 2-component. We define the union bridges, and the $\textcircled{1}$- and $\textcircled{2}$- bridges of $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ with respect to C as above. We call an embedding $\mathcal{E}$ of C together with an assignment of the union bridges to its faces *admissible* if and only if, (i) for each union bridge, all attachments are incident to the face to which it is assigned, and (ii) no two ①- or ②-bridges that are assigned to the same face alternate. For a union bridge B, let C_B denote the cycle consisting of the attachments of B in the ordering of an arbitrary cycle of G containing all the attachments. It can be shown that this cycle is uniquely determined. Let G_B denote the graph consisting of B and C_B. We call these graphs the *union bridge graphs*.

Lemma 2. *Let $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ be two connected planar graphs and let C be 2-component of the common graph G. Then the graphs $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ admit a* SEFE *if and only if (i) C admits an admissible embedding, and (ii) each union bridge graph admits a* SEFE. *If a* SEFE *exists, the embedding of C can be chosen as an arbitrary admissible embedding.*

Proof. Clearly, a SEFE of $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ defines an embedding of C and a bridge assignment that is admissible. Moreover, it induces a SEFE of each union bridge graph.

Conversely, assume that C admits an admissible embedding and each union bridge graph admits a SEFE. We obtain a SEFE of $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ as follows. Embed C with the admissible embedding and consider a face f of this embedding with facial cycle C_f. Let $B_1, \dots, B_k$ denote the union bridges that are assigned to this face, and let $(\mathcal{E}_1^{\textcircled{1}}, \mathcal{E}_1^{\textcircled{2}}), \dots, (\mathcal{E}_k^{\textcircled{1}}, \mathcal{E}_k^{\textcircled{2}})$ be simultaneous embeddings of the bridge graphs G_B. By subdividing the cycle C_B, in each of the embeddings, we may assume that the outer face of each B_i in the embedding $(\mathcal{E}_i^{\textcircled{1}}, \mathcal{E}_i^{\textcircled{2}})$ is the facial cycle C_f with the same orientation in each of them. By Lemma 1, we can hence combine them to a single SEFE of all union bridges whose outer face is the cycle C_f. We embed this SEFE into the face f of C. Since we can treat the different faces of C independently, applying this step for each face yields a SEFE of $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ with the claimed embedding of C. □

Lemma 2 suggests a simple strategy for reducing SEFE instances containing non-trivial 2-components. Namely, take such a component, construct the corresponding union bridge graphs, where C occurs only as a cycle, and find an admissible embedding of C. Finding an admissible embedding for C can be done as follows. To enforce the non-overlapping attachment property, replace each ①-bridge of C by a *dummy ①-bridge* that consists of a single vertex that is connected to the attachments of that bridge via edges in $E^{\textcircled{1}}$. Similarly, we replace ②-bridges, which are connected to attachments via exclusive edges in $E^{\textcircled{2}}$. We seek a SEFE of the resulting instance (where the common graph is biconnected), additionally requiring that dummy bridges belonging to the same union bridge are embedded in the same face. We also refer to such an instance as SEFE *with union bridge constraints*. A slight modification of the algorithm by Angelini et al. [1] can decide the existence of such an embedding in polynomial time. It then remains to treat the bridge graphs. Exhaustively applying Lemma 2 results in a set of SEFE instances where each 2-component is trivial.

Linear-Time Decomposition. We now show that the set of instances resulting from exhaustively applying Lemma 2 can be computed in linear time.

Theorem 3. *Given a* SEFE *instance, an equivalent set of instances of total linear size such that each 2-component of these instances is trivial can be computed in linear time.*

Let G be a planar graph and let $C_1, \dots, C_k$ be connected components of G. We are interested in simultaneously determining for each component C_i the number of incident bridges, and for each such bridge its attachment vertices. For this, we introduce the notion of *subbridges*. A subbridge of G with respect to $C_1, \ldots, C_k$ is a maximal connected subgraph of G that does not become disconnected by removing all vertices of one component C_i. It is readily seen that each C_i-bridge B contains a unique subbridge S incident to C_i and that the attachments of S at C_i are exactly the attachments of B at C_i. We will thus rather work with the subbridges than the actual bridges as they represent the same information but in a more compact way. Our reduction now works in three phases.

1. Compute for each 2-component of G the number of ①-, ②-bridges, for each such bridge its attachments, and the grouping of these bridges into union bridges.
2. Find for each 2-component an admissible embedding with respect to its bridges.
3. Compute for each subbridge of G with respect to its 2-components a corresponding instance where each 2-component has been replaced by a suitable cycle.

The correctness of this approach descends from Lemma 2, since the set of instances computed by the procedure is exactly the one that can be obtained by exhaustively applying this lemma. The details of the implementation of this procedure are deferred to the full version of this paper. Here we only sketch the main ideas. For step 1, we exploit the fact that, after contracting each connected component of G that is biconnected to a single vertex, (almost) every such component is a cutvertex, and the union subbridges are essentially the blocks of the resulting graph. We can then traverse for each cutvertex its incident edges and label them by the block (subbridge) containing them. This allows us to construct the dummy-bridges and union bridges that are solved in step 2. For step 2, we modify the algorithm due to Angelini et al.[1]. Augmenting it such that it computes admissible embedding in polynomial time is straightforward. Achieving linear running time is quite technical and, like the linear version of the original algorithm, requires some intricate data structures. Step 3 is finally implemented by taking the admissible embeddings from step 2. We then traverse each such face exactly one, and construct, during this traversal, the corresponding cycles in all incident subbridges that are embedded in this face.

4 Instances with Common P-Node Degree 3

We consider instances of SEFE that have common P-node degree 3. Recall that a simultaneous embedding must induce consistent edge orderings and consistent relative positions on the common graph. We show how to address both requirements separately, by formulating necessary and sufficient constraints using linear equations over $\mathbb{F}_2$. Both resulting systems of equations share all variables representing embedding choices. Satisfying both sets of linear equations at the same time then solves SEFE.

Before we can follow this strategy, we need to address one problem. The relative position of a component H' of G with respect to another connected component H, denoted by $\mathrm{pos}_H(H')$, is the face of H containing H'. However, the set of faces of H depends on the embedding of H. To be able to handle relative positions independently from edge orderings, we need to express the relative positions independently from faces.

4.1 Relative Positions with Respect to a Cycle Basis

A *generalized cycle* C in a graph H is a subset of its edges such that every vertex of H has an even number of incident edges in C. The *sum* $C + C'$ of two generalized cycles is the symmetric difference between the edge sets, i.e., an edge e is contained in $C + C'$ if and only if it is contained in C or in C' but not in both. The resulting edge set $C + C'$ is again a generalized cycle. The set of all generalized cycles in H is a vector space over $\mathbb{F}_2$. A basis of this vector space is called *cycle basis* of H.

Instead of considering the relative position $\mathrm{pos}_H(H')$ of a connected component H' with respect to another component H, we choose a cycle basis $\mathcal{C}$ of H and show that the relative positions of H' with respect to the cycles in $\mathcal{C}$ suffice to uniquely define $\mathrm{pos}_H(H')$, independent from the embedding of H. We assume H to be biconnected. All results can be extended to connected graphs by using a cycle basis for each block.

Let $C_0, \dots, C_k$ be the set of facial cycles with respect to an arbitrary planar embedding of H. The set $\mathcal{C} = \{C_1, \dots, C_k\}$ obtained by removing one of the facial cycles is a cycle basis of G. A cycle basis that can be obtained in this way is called *planar cycle basis*. In the following we assume all cycle bases to be planar cycle bases. Moreover, we consider all cycles to have an arbitrary but fixed orientation, which has the effect, that $\mathrm{pos}_C(p)$ for any cycle C and any point p can have either the value LEFT or RIGHT.

Theorem 4. *Let H be a planar graph embedded on the sphere, let p be a point on the sphere, and let $\mathcal{C} = \{C_1, \dots, C_k\}$ be an arbitrary planar cycle basis of H. Then the face containing p is determined by the relative positions $\mathrm{pos}_{C_i}(p)$ for $1 \leq i \leq k$.*

Proof (sketch). Clearly, the point p and the face f containing p have to lie on the same side of each of the cycles in $\mathcal{C}$. It remains to show that the face with this property is unique. Let C be the facial cycle of f and let $C = C_1 + \cdots + C_\ell$ be the linear combination of basis cycles of C. The *position vector* of a point p with respect to the facial cycle C is $\mathrm{pos}(p) = (\mathrm{pos}_{C_1}(p), \dots, \mathrm{pos}_{C_\ell}(p))$. It can be seen that inside f, the vector $\mathrm{pos}(p)$ has a different parity of values LEFT than outside, which shows, that no other face can have the same relative positions with respect to all cycles in $\mathcal{C}$. □

4.2 Consistent Edge Orderings

We first assume that the graphs $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ are biconnected. There exists an instance of SIMULTANEOUS PQ-ORDERING that has a solution if and only if $G^{\textcircled{1}}$ and $G^{\textcircled{2}}$ admit embeddings with consistent edge ordering [4]. This solution is based on the *PQ-embedding representation*, an instance of SIMULTANEOUS PQ-ORDERING representing all embeddings of a biconnected planar graph. We describe this embedding representation and show how to simplify it for instances that have common P-node degree 3.

For each vertex $v^{\textcircled{1}}$ of $G^{\textcircled{1}}$, the PQ-embedding representation, denoted by $D(G^{\textcircled{1}})$, contains the *embedding tree* $T(v^{\textcircled{1}})$ having a leaf for each edge incident to $v^{\textcircled{1}}$, representing all possible orders of edges around $v^{\textcircled{1}}$. For every P-node $\mu^{\textcircled{1}}$ in the SPQR-tree $\mathcal{T}^{\textcircled{1}}$ of $G^{\textcircled{1}}$ that contains $v^{\textcircled{1}}$ in $\mathrm{skel}(\mu^{\textcircled{1}})$ there is a P-node in $T(v^{\textcircled{1}})$ representing the choice to reorder the virtual edges in $\mathrm{skel}(\mu^{\textcircled{1}})$. Similarly, for every R-node $\mu^{\textcircled{1}}$ in $\mathcal{T}^{\textcircled{1}}$ containing $v^{\textcircled{1}}$ there is a Q-node in $T(v^{\textcircled{1}})$ whose flip corresponds to the flip of $\mathrm{skel}(\mu^{\textcircled{1}})$. As the orders of edges around different vertices of $G^{\textcircled{1}}$ cannot be chosen independently

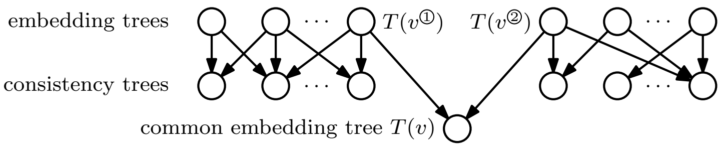

Fig. 2. The Q-embedding representations of $G^{①}$ and $G^{②}$ and one common embedding tree

from each other, so called *consistency trees* are added as common children to enforce Q-nodes stemming from the same R-node in $\mathcal{T}^{①}$ to have the same flip and P-nodes stemming from the same P-node to have consistent (i.e., opposite) orders. Every solution of the resulting instance corresponds to a planar embedding of $G^{①}$ and vice versa [4].

As we are only interested in the order of common edges, we modify $D(G^{①})$ by projecting each PQ-tree to the leaves representing common edges. As $G^{①}$ and $G^{②}$ have common P-node degree 3, all P-nodes of the resulting PQ-trees have degree 3 and can be assumed to be Q-nodes representing a binary decision. We call the resulting instance *Q-embedding representation* and denote it by $D(G^{①})$. Let $\mu^{①}$ be an R-node of the SPQR-tree $\mathcal{T}^{①}$ whose embedding has influence on the ordering of common edges around a vertex. It is not hard to see that the Q-embedding representation contains a consistency tree consisting of a single Q-node representing the flip of $\mathrm{skel}(\mu^{①})$. We associate the binary variable $\mathrm{ord}(\mu^{①})$ with this decision. For a P-node $\mu^{①}$ we get a similar result. Let $u^{①}$ and $v^{①}$ be the nodes in $\mathrm{skel}(\mu^{①})$. If the consistency tree enforcing a consistent decision in the embedding trees $T(u^{①})$ and $T(v^{①})$ has degree 3, its flip represents the embedding decision for $\mathrm{skel}(\mu^{①})$ and we again get a binary variable $\mathrm{ord}(\mu^{①})$. Otherwise, this consistency tree contains two or less leaves and can be ignored. Then the choices for the Q-nodes corresponding to $\mu^{①}$ in $T(u^{①})$ and $T(v^{①})$ are independent and we get one binary variable for each of these Q-nodes. We denote these variables by $\mathrm{ord}(\mu_u^{①})$ and $\mathrm{ord}(\mu_v^{①})$. We call these variables *PR-ordering variables.*

For a common vertex v occurring as $v^{①}$ and $v^{②}$ in $G^{①}$ and $G^{②}$, respectively, we can ensure a consistent edge ordering by adding a so called *common embedding tree* $T(v)$ as child of the embedding trees $T(v^{①})$ and $T(v^{②})$ in the Q-embedding representations of $G^{①}$ and $G^{②}$; see Fig. 2. We get the following lemma.

Lemma 3. *Let $G^{①}$ and $G^{②}$ be two biconnected graphs with common P-node degree 3. Requiring the common edges to be ordered consistently is equivalent to satisfying a system of linear equations M_{ord} over $\mathbb{F}_2$ with the following properties.*

(i) All equations in M_{ord} are of the type $x + y = c$ for $c \in \mathbb{F}_2$.
(ii) M_{ord} contains all PR-ordering variables.
(iii) M_{ord} has linear size and can be computed in linear time.

In the following, we extend this result to the case where we allow exclusive cutvertices. Let $B_1^{①}, \dots, B_k^{①}$ be the blocks of $G^{①}$ and let $B_1^{②}, \dots, B_\ell^{②}$ be the blocks of $G^{②}$. We say that embeddings of these blocks have *blockwise consistent edge orderings* if for every pair of blocks $B_i^{①}$ and $B_j^{②}$ sharing a vertex v the edges incident to v they share are ordered consistently. To have consistent edge orderings, it is obviously necessary to have blockwise consistent edge orderings.

When composing the embeddings of two blocks that share a cutvertex, the edges of each of the two blocks have to appear consecutively (note that this is no longer true for three or more blocks), which leads to another necessary condition. Let v be an exclusive cutvertex of $G^{\textcircled{1}}$. Then v is contained in a single block of $G^{\textcircled{2}}$ whose embedding induces an order $O^{\textcircled{2}}$ on all common edges incident to v. For every pair of blocks $B_i^{\textcircled{1}}$ and $B_j^{\textcircled{2}}$ containing v, the edges in $B_i^{\textcircled{1}}$ must appear consecutively in the order of the edges incident to v in $B_i^{\textcircled{1}}$ and $B_j^{\textcircled{2}}$ that is induced by $O^{\textcircled{2}}$. If this is true for every exclusive cutvertex, we say that the embeddings have *pairwise consecutive blocks*.

Lemma 4. *Two graphs without simultaneous cutvertices admit embeddings with consistent edge orderings if and only if their blocks admit embeddings that have blockwise consistent edge orderings and pairwise consecutive blocks.*

To extend Lemma 3 to the case where we allow exclusive cutvertices, we enforce blockwise consistent edge orderings and pairwise consecutive blocks by adding additional PQ-trees to the above instance of SIMULTANEOUS PQ-ORDERING. As before, we get direct access to the embedding chosen for each block, via the PR-ordering variables. We want to get access to the ordering of common edges around a cutvertex v of $G^{\textcircled{1}}$ in a similar way. Let $B^{\textcircled{1}}$ be a block that contains the common edge e incident to v and let e_1 and e_2 be two common edges incident to v that are contained in a different block. We use the *cutvertex-ordering variable* $\mathrm{ord}(e_1, e_2, B^{\textcircled{1}})$ to denote the order of e_1, e_2, and e. Note that this is independent from the choice of the edge e of $B^{\textcircled{1}}$. To decrease the number of variables, we only consider those variables that are *required by a cycle basis $\mathcal{C}$*, where $\mathrm{ord}(e_1, e_2, B^{\textcircled{1}})$ is required by $\mathcal{C}$ if e_1 and e_2 share a cycle in $\mathcal{C}$.

Lemma 5. *Given two graphs without union or simultaneous cutvertices with common P-node degree 3, requiring the common edges to be ordered consistently is equivalent to satisfying a system of linear equations M_{ord} with the following properties.*

(i) All equations in M_{ord} are of the type $x + y = 0$ or $x + y = 1$.
(ii) M_{ord} contains all PR-ordering variables and all cutvertex-ordering variables required by a cycle basis of the common graph.
(iii) M_{ord} has size $O(\min\{n^2, n\Delta^2\})$ (where Δ is the maximum degree in the common graph) and can be computed in linear time in its size.

4.3 Consistent Relative Positions

In this section, we present a system of linear equations M_{pos} containing the PR-ordering and cutvertex-ordering variables such that satisfying M_{pos} is equivalent to requiring consistent relative positions for an instance of SEFE. Let H and H' be two connected components of the common graph G. To represent the relative position $\mathrm{pos}_{H'}(H)$ of H with respect to H', we use the relative positions $\mathrm{pos}_C(H)$ of H with respect to cycles C in the cycle basis of H' (Theorem 4). To get binary variables, we use $\mathrm{pos}_C(H) = 0$ if H lies to the right of C and $\mathrm{pos}_C(H) = 1$ if H lies to the left of C. When we consider the graph $G^{\textcircled{1}}$ containing G, it is known that the value of $\mathrm{pos}_C(H)$ is determined by a single, very local embedding decision of $G^{\textcircled{1}}$ [3]. In the following we consider the three possible cases that $\mathrm{pos}_C(H)$ is determined by an R-node, by a P-node or by a cutvertex.

R-Node. If $\mathrm{pos}_C(H)$ is determined by an R-node μ, then C induces a cycle κ in $\mathrm{skel}(\mu)$ and parts of H are contained in a virtual edge ε not contained in κ. The relative position of H with respect to C is the same as the position of ε with respect to κ [3]. As the value of $\mathrm{pos}_\kappa(\varepsilon)$ changes, when the embedding of $\mathrm{skel}(\mu)$ changes, we can simply set $\mathrm{pos}_C(H) + \mathrm{ord}(\mu) = c$ (where $c \in \mathbb{F}_2$ depends on the reference embedding of $\mathrm{skel}(\mu)$ and the orientation of C). Note that this implicitly ensures the consistency of all relative positions that are determined by thee embedding of $\mathrm{skel}(\mu)$.

P-Node. If $\mathrm{pos}_C(H)$ is determined by the embedding of $\mathrm{skel}(\mu)$ of a P-node μ, then C induces a cycle κ (of length 2) in $\mathrm{skel}(\mu)$ and H is completely contained in a single edge ε of $\mathrm{skel}(\mu)$ not belonging to κ. Again $\mathrm{pos}_C(H)$ in G is the same as $\mathrm{pos}_\kappa(\varepsilon)$ in $\mathrm{skel}(\mu)$. However, this time the embedding choices of $\mathrm{skel}(\mu)$ are more complicated than to flip or not to flip. Thus, we have to consider all relative positions decided by the embedding of $\mathrm{skel}(\mu)$ at once, to get the dependencies between them.

We only consider the case where the common graph induces paths between the vertices of $\mathrm{skel}(\mu)$ in the expansion graphs of three edges ε_1, ε_2, and ε_3 (all other cases are simpler as μ has common P-node degree 3). Cycles in the cycle basis $\mathcal{C}$ can induce three different cycles, namely $\kappa_{1,2}$, $\kappa_{2,3}$, and $\kappa_{1,3}$ consisting of the virtual edges $(\varepsilon_1, \varepsilon_2)$, $(\varepsilon_2, \varepsilon_3)$, and $(\varepsilon_1, \varepsilon_3)$, respectively. For every virtual edge $\varepsilon \neq \varepsilon_i$, we get the three variables $\mathrm{pos}_{\kappa_{1,2}}(\varepsilon)$, $\mathrm{pos}_{\kappa_{2,3}}(\varepsilon)$, and $\mathrm{pos}_{\kappa_{1,3}}(\varepsilon)$ determining the position of ε with respect to these three cycles. Recall that the variable $\mathrm{ord}(\mu)$ determines the ordering of the three edges ε_1, ε_2, and ε_3. Consider the case that $\mathrm{ord}(\mu) = 0$ and assume that the reference order of ε_1, ε_2, and ε_3 as well as the orientation of the cycles $\kappa_{1,2}$, $\kappa_{2,3}$, and $\kappa_{1,3}$ is as shown in Fig. 3. Then either all three relative positions have the value 0 (which corresponds to RIGHT), or exactly two relative position has the value 1. Thus, a combination of values for the positions $\mathrm{pos}_{\kappa_{1,2}}(\varepsilon)$, $\mathrm{pos}_{\kappa_{2,3}}(\varepsilon)$, and $\mathrm{pos}_{\kappa_{1,3}}(\varepsilon)$ is possible if and only if there is an even number of 1s. When setting $\mathrm{ord}(\mu) = 1$, there need to be an odd number of 1s. This yields the constraint $\mathrm{ord}(\mu) + \mathrm{pos}_{\kappa_{1,2}}(\varepsilon) + \mathrm{pos}_{\kappa_{2,3}}(\varepsilon) + \mathrm{pos}_{\kappa_{1,3}}(\varepsilon) = 0$ (a different reference embedding or different orientations of the cycles may lead to a 1 on the right-hand side of the equation). As for R-nodes we set $\mathrm{pos}_C(H) + \mathrm{pos}_{\kappa_{i,j}}(\varepsilon) = c$ ($c \in \mathbb{F}_2$), if C induces $\kappa_{i,j}$ in $\mathrm{skel}(\mu)$ and H is contained in the expansion graph of ε.

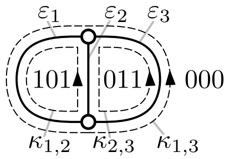

Fig. 3. A P-node with the three cycles $\kappa_{1,2}$, $\kappa_{1,2}$, and $\kappa_{1,2}$ (dashed)

Cutvertex. The relative position $\mathrm{pos}_C(H)$ is determined by a cutvertex v of $G^{\textcircled{1}}$ if C contains v and H lies in a split component S (with respect to v) different from the split component containing C. We can either embed the whole split component S to the left or to the right of C. For this decision, we introduce the variable $\mathrm{pos}_C(S)$. Clearly, we get $\mathrm{pos}_C(H) = \mathrm{pos}_C(S)$ for every connected component H of G in S. Moreover, there are no further constraints on the relative positions determined by the embedding at the cutvertex v [4]. In case the split component S contains a common edge incident to v, fixing $\mathrm{pos}_C(S)$ is equivalent fixing the cutvertex-ordering variable $\mathrm{ord}(e_1, e_2, B)$, where e_1 and e_2 are the edges in C incident to v and B is the block in S containing v. Thus, we get $\mathrm{ord}(e_1, e_2, B) + \mathrm{pos}_C(S) = c$ ($c \in \mathbb{F}_2$). Together with the constraints for the relative positions determined by the P- and R-nodes, we get the following lemma.

Lemma 6. *Let $G^{①}$ and $G^{②}$ be two graphs without union or simultaneous cutvertices with common P-node degree 3. Requiring the relative positions to be consistent is equivalent to satisfying a system of linear equations M_{pos} with the following properties.*

(i) All equations in M_{pos} are of the type $x + y = c$ (with $c \in \mathbb{F}_2$) except for a linear number of equations of size 4.

(ii) M_{pos} contains all PR-ordering variables and all cutvertex-ordering variables required by a cycle basis of G.

(iii) M_{pos} has quadratic size and can be computed in quadratic time.

Lemma 5 and Lemma 6 yield the following theorem. We obtain the quadratic running time by first eliminating all equations of size 1, and then solving the remaining system of $O(n)$ linear equations of size 4 with the algorithm by Wiedemann [12].

Theorem 5. SEFE *can be solved in quadratic time for two graphs without union or simultaneous cutvertex with common P-node degree 3.*

References

1. Angelini, P., Di Battista, G., Frati, F., Patrignani, M., Rutter, I.: Testing the simultaneous embeddability of two graphs whose intersection is a biconnected or a connected graph. Journal of Discrete Algorithms 14, 150–172 (2012)
2. Bläsius, T., Kobourov, S.G., Rutter, I.: Simultaneous Embedding of Planar Graphs. In: Handbook of Graph Drawing and Visualization, pp. 349–381. Chapman and Hall/CRC (2013)
3. Bläsius, T., Rutter, I.: Disconnectivity and relative positions in simultaneous embeddings. In: Didimo, W., Patrignani, M. (eds.) GD 2012. LNCS, vol. 7704, pp. 31–42. Springer, Heidelberg (2013)
4. Bläsius, T., Rutter, I.: Simultaneous PQ-ordering with applications to constrained embedding problems. In: Proc. 24th ACM-SIAM Sympos. Discrete Algorithm, SODA 2013, pp. 1030–1043. ACM (2013)
5. Booth, K.S., Lueker, G.S.: Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. Journal of Computer and System Sciences 13, 335–379 (1976)
6. Di Battista, G., Tamassia, R.: On-line maintenance of triconnected components with SPQR-trees. Algorithmica 15(4), 302–318 (1996)
7. Gassner, E., Jünger, M., Percan, M., Schaefer, M., Schulz, M.: Simultaneous graph embeddings with fixed edges. In: Fomin, F.V. (ed.) WG 2006. LNCS, vol. 4271, pp. 325–335. Springer, Heidelberg (2006)
8. Gutwenger, C., Mutzel, P.: A linear time implementation of SPQR-trees. In: Marks, J. (ed.) GD 2000. LNCS, vol. 1984, pp. 77–90. Springer, Heidelberg (2001)
9. Haeupler, B., Jampani, K., Lubiw, A.: Testing simultaneous planarity when the common graph is 2-connected. J. Graph Algorithms Appl. 17(3), 147–171 (2013)
10. Jünger, M., Schulz, M.: Intersection graphs in simultaneous embedding with fixed edges. Journal of Graph Algorithms and Applications 13(2), 205–218 (2009)
11. Schaefer, M.: Toward a theory of planarity: Hanani-tutte and planarity variants. In: Didimo, W., Patrignani, M. (eds.) GD 2012. LNCS, vol. 7704, pp. 162–173. Springer, Heidelberg (2013)
12. Wiedemann, D.: Solving sparse linear equations over finite fields. IEEE Transactions on Information Theory 32(1), 54–62 (1986)