

Arc Diagrams, Flip Distances, and Hamiltonian Triangulations

Jean Cardinal^{*1}, Michael Hoffmann^{†2}, Vincent Kusters^{†2},
Csaba D. Tóth³, and Manuel Wettstein^{‡2}

1 Université libre de Bruxelles (ULB), Belgium

jcardin@ulb.ac.be

2 Department of Computer Science, ETH Zürich, Switzerland

{hoffmann,vincent.kusters,manuelwe}@inf.ethz.ch

3 California State University Northridge, Los Angeles, CA, USA

cdtoth@acm.org

Abstract

We show that every triangulation (maximal planar graph) on $n \geq 6$ vertices can be flipped into a Hamiltonian triangulation using a sequence of less than $n/2$ combinatorial edge flips. The previously best upper bound uses 4-connectivity as a means to establish Hamiltonicity. But in general about $3n/5$ flips are necessary to reach a 4-connected triangulation. Our result improves the upper bound on the diameter of the flip graph of combinatorial triangulations on n vertices from $5.2n - 33.6$ to $5n - 23$. We also show that for every triangulation on n vertices there is a simultaneous flip of less than $2n/3$ edges to a 4-connected triangulation. The bound on the number of edges is tight, up to an additive constant. As another application we show that every planar graph on n vertices admits an arc diagram with less than $n/2$ biarcs, that is, after subdividing less than $n/2$ (of potentially $3n - 6$) edges the resulting graph admits a 2-page book embedding.

1998 ACM Subject Classification G.2.2 Graph Theory

Keywords and phrases graph embeddings, edge flips, flip graph, separating triangles

Digital Object Identifier 10.4230/LIPIcs.STACS.2015.197

1 Introduction

An *arc diagram* (Figure 1) is a drawing of a graph in which vertices are represented by points on a horizontal line, called the *spine*, and edges are drawn either as one halfcircle (*proper arc*) or as a sequence of halfcircles centered on the line (forming a smooth Jordan arc). In a *proper arc diagram* all arcs are proper. Arc diagrams have been used and studied in many contexts since their first appearance in the mid-sixties [3, 24]. They constitute a well-studied geometric representation in graph drawing [14] that occurs, for instance, in the study of crossing numbers [1, 6] and universal point sets for circular arc drawings [4].

Bernhart and Kainen [5] proved that a planar graph admits a *plane* (i.e., crossing-free) proper arc diagram if and only if it can be augmented to a Hamiltonian planar graph by

* Partially supported by the ESF EUROCORES programme EuroGIGA, CRP ComPoSe.

† Partially supported by the ESF EUROCORES programme EuroGIGA, CRP GraDR and the Swiss National Science Foundation, SNF Project 20GG21-134306.

‡ Partially supported by the ESF EUROCORES programme EuroGIGA, CRP ComPoSe and the Swiss National Science Foundation, SNF Project 20GG21-134318/1.



© Jean Cardinal, Michael Hoffmann, Vincent Kusters,
Csaba D. Tóth, and Manuel Wettstein;
licensed under Creative Commons License CC-BY

32nd Symposium on Theoretical Aspects of Computer Science (STACS 2015).

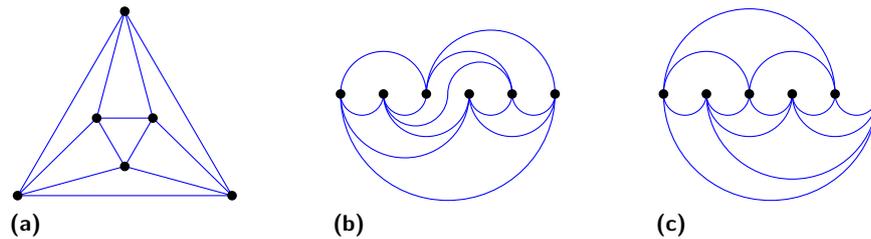
Editors: Ernst W. Mayr and Nicolas Ollinger; pp. 197–210

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany





■ **Figure 1** A plane straight-line drawing (a), an arc diagram (b) and a proper arc-diagram (c) of the same graph.

adding new edges. Such planar graphs are also called *subhamiltonian*, and they are NP-hard to recognize [26]. A Hamiltonian cycle in the augmented graph directly yields a feasible order for the vertices on the spine. Every planar graph can be subdivided into a subhamiltonian graph with at most one subdivision vertex per edge [22]. Consequently, every planar graph admits a plane *biarc diagram* in which each edge is either a proper arc or the union of two halfcircles (a *biarc*); one above and one below the spine. Di Giacomo et al. [15] showed that every planar graph even admits a *monotone* biarc diagram in which every biarc is *x-monotone*—such an embedding is also called a *2-page topological book embedding*. See [14] for various other applications of subhamiltonian subdivisions of planar graphs.

Eppstein [13] said: “Arc diagrams (with one arc per edge) are very usable and practical but can only handle a subset of planar graphs.” Using biarcs allows to represent all planar graphs, but adds to the complexity of the drawing. Hence it is a natural question to ask: How close can we get to a proper arc diagram, while still being able to represent all planar graphs? A natural measure of complexity is the number of biarcs used.

Previous methods for subdividing an n -vertex planar graph into a subhamiltonian graph use at most one subdivision per edge [14, 15, 19, 22], consequently the number of biarcs in an arc diagram is bounded by the number of edges. Our main goal in this paper is to tighten the upper and lower bound on the minimum number of biarcs in an arc diagram (or, alternatively, the number of subdivision vertices in a subhamiltonian subdivision) of a planar graph with n vertices. Minimizing the number of biarcs is clearly NP-hard, since the number of biarcs is zero if and only if the graph is subhamiltonian.

Our results. We show that the number of biarcs can be bounded by n , even when they are restricted to be monotone. Although previous methods can be shown to yield less than the trivial $3n - 6$ biarcs [19], or ensure monotonicity [15], we give the first proof that both properties can be guaranteed simultaneously. The algorithm is similar to the canonical ordering-based method of Di Giacomo et al. [15].

► **Theorem 1.** *Every planar graph on $n \geq 4$ vertices admits an arc diagram using at most $n - 4$ biarcs, all of which are monotone.*

For arbitrary (not necessarily monotone) biarcs we achieve better bounds. Our main tool is relating subhamiltonian planar graphs to edge flips in triangulations. A *flip* in a triangulation involves switching the diagonal of a quadrilateral made of two adjacent facial triangles. We consider *combinatorial* flips, which can be regarded as an operation on an abstract graph. The *flip graph* induced by flips on the set of all triangulations on n vertices, and the corresponding *flip distance* between two triangulations, have been the topic of extensive research [9, 11]. For instance, the flip diameter restricted to the interior of a convex polygon is equivalent to rotation distance of binary trees [25, 23].

We prove that in every triangulation there exists a set of less than $2n/3$ edges that can be flipped *simultaneously* so that the resulting triangulation is 4-connected, and that this bound is tight up to an additive constant (Section 4). Since by Tutte's Theorem every 4-connected planar graph is Hamiltonian, we can transform every planar graph into a subhamiltonian graph by subdividing at most $2n/3$ edges. The fact that a single simultaneous flip can make a triangulation 4-connected has already been established by Bose et al. [8]. However, they do not give any bound on the number of flipped edges.

► **Theorem 2.** *Every maximal planar graph on $n \geq 6$ vertices can be transformed into a 4-connected maximal planar graph using a simultaneous flip of at most $\lfloor (2n - 7)/3 \rfloor$ edges.*

► **Theorem 3.** *For every $i \in \mathbb{N}$, there is a maximal planar graph G_i on $n_i = 3i + 4$ vertices such that no simultaneous flip of less than $(2n_i - 8)/3 = 2i$ edges results in a 4-connected graph.*

Finally, we prove an upper bound on the flip distance of a planar triangulation to Hamiltonicity, that is, on the worst-case number of successive flips required to reach a Hamiltonian triangulation (Section 5). Given the hardness of determining whether a given planar graph is Hamiltonian, we should not expect a nice characterization of (non-)Hamiltonicity. Hence, in the context of planar graphs, 4-connectivity is often used as a substitute because by Tutte's Theorem it is a sufficient condition for Hamiltonicity. Bose et al. [10] gave a tight bound (up to an additive constant) of $3n/5$ flips to transform a given triangulation on n vertices into a 4-connected triangulation. We show that fewer flips are sufficient to guarantee Hamiltonicity. Obviously, the target triangulation is not 4-connected in general, which means it possibly contains separating triangles.

► **Theorem 4.** *Every maximal planar graph on $n \geq 6$ vertices can be transformed into a Hamiltonian maximal planar graph using a sequence of at most $\lfloor (n - 3)/2 \rfloor$ edge flips.*

In this case we do not have a matching lower bound. The best lower bound we know can be obtained using *Kleetopes* [16]. These are convex polytopes that are generated from another convex polytope by replacing every face by a small pyramid. In the language of planar graphs, we start from a 3-connected planar graph and for every face add a new vertex that is connected to all vertices on the boundary of the face. If the graph we start from has enough faces, then the added vertices form a large independent set so that the resulting graph is not Hamiltonian. Aichholzer et al. [2] describe such a construction explicitly in the context of flipping a triangulation to a Hamiltonian triangulation, but state the asymptotics only. A precise counting reveals the following figures.

► **Theorem 5.** *For every $i \in \mathbb{N}$, there is a maximal planar graph G_i on $n_i = 3i + 8$ vertices such that no sequence of less than $(n_i - 8)/3 = i$ edge flips produces a Hamiltonian graph and no set of less than $(n_i - 8)/3 = i$ subdivision vertices produces a subhamiltonian graph.*

Our proof for Theorem 4 is constructive, and each flip in the sequence involves an edge of the initial graph G and is incident to a separating triangle of G . Some of these edges may be incident to a common facial triangle, so they cannot always be flipped simultaneously. However, we show that if we *subdivide* each of these edges (instead of successively flipping them), we obtain a subhamiltonian graph. Combined with the characterization of Bernhart and Kainen [5], this yields a new bound on the number of biarcs.

► **Corollary 6.** *Every planar graph on $n \geq 6$ vertices admits a biarc diagram with at most $\lfloor (n - 3)/2 \rfloor$ biarcs.*

As another corollary of Theorem 4, we establish a new upper bound on the diameter of the flip graph of all triangulations on n vertices, improving on the previous best bound of $5.2n - 33.6$ by Bose et al. [10]. Mori et al. [21] showed that any two Hamiltonian triangulations on n vertices can be transformed into each other by at most $\max\{4n - 20, 0\}$ flips. Combined with Theorem 4, this implies the following.

► **Corollary 7.** *Every two triangulations on $n \geq 6$ vertices can be transformed into each other using at most $5n - 23$ edge flips.*

Due to space constraints, many proofs have to be omitted.

2 Notation

A *drawing* of a graph G in \mathbb{R}^2 maps the vertices into distinct points in the plane and maps each edge to a Jordan arc between (the images of) the two vertices that is disjoint from (the image of) any other vertices. To avoid notational clutter it is common to identify vertices and edges with their geometric representation. A drawing is called *plane* (or an *embedding*) if no two edges intersect except at a possible common endpoint. Only planar graphs admit plane drawings, but not every drawing of a planar graph is plane. A *maximal planar* graph on n vertices is a planar graph with $3n - 6$ edges. In this paper the term *triangulation* is used as a synonym for maximal planar graph.¹

In a plane drawing of a triangulation G , every face (including the outer face) is bounded by three edges. Hence, every triangulation with $n \geq 4$ vertices is 3-connected [12][Lemma 4.4.5]. Every 3-connected planar graph has a topologically unique plane drawing, apart from the choice of the outer face. Specifically, the facial triangles are precisely the nonseparating chordless cycles of G in every plane drawing [12][Proposition 4.2.7]. Consequently, G has a well-defined dual graph G^* (independent of the drawing): the vertices of G^* correspond to the faces of G , and two vertices of G^* are adjacent if and only if the corresponding faces share an edge. A triangle of G that is not facial is called a *separating* triangle and its removal disconnects the graph.

A graph is *Hamiltonian* if it contains a cycle through all vertices. By a famous theorem of Tutte, all 4-connected planar graphs are Hamiltonian. For triangulations, 4-connectivity is equivalent to the absence of separating triangles. A vertex or an edge is *incident* to a triangle T in a graph if it is a vertex or edge of T .

A triangulation G can be partitioned into a *4-block tree* \mathcal{B} . Each vertex of \mathcal{B} is either a maximal 4-connected component of G or a subgraph of G that is isomorphic to K_4 . Two vertices of \mathcal{B} are adjacent if they share a separating triangle of G . The 4-block tree is similar to the standard block-tree for 2-connected components, but the generalization of the notion “component” to higher connectivity is not straightforward in general. For a triangulation, however, the 4-block tree is well-defined and can be computed in linear time and space [18].

Flips. Consider an edge ab of a triangulation G and let abc and adb denote the two incident facial triangles. The *flip* of ab replaces the edge ab by the edge cd . If this operation produces a triangulation (i.e., if the edge cd is not already present in G), we call ab *flippable*².

¹ In contrast, a maximal plane *straight-line* drawing may have fewer edges, depending on the number of points on the convex hull.

² We consider *combinatorial flips*, as opposed to *geometric flips* defined for straight-line plane drawings, where an edge is flippable if and only if the quadrilateral formed by the two incident facial triangles is convex.

A closely related concept is the *simultaneous flip* of a set F of flippable edges in a triangulation $G = (V, E)$, which is defined as follows. For $e \in F$ denote by $c(e)$ the edge created by flipping e in G , and let $C(F) = \bigcup_{e \in F} c(e)$. Then the simultaneous flip of F in G results in the graph $G' = (V, (E \setminus F) \cup C(F))$. Bose et al. [8] introduced this notion and showed that the result of a simultaneous flip is a triangulation if every facial triangle of G is incident to at most one edge from F and the edges $c(e)$, for $e \in F$, are all distinct and not present in E .

3 Biarc Diagrams

► **Lemma 8.** *If a planar graph G has a simultaneous flip of k edges such that the resulting graph is Hamiltonian, then G admits an arc diagram with at most k biarcs.*

Proof. Let H be a Hamiltonian graph obtained from G by simultaneously flipping an edge set E_1 to E_2 with $|E_1| = k$. Without loss of generality, assume that E_1 is a minimal set of edges that must be flipped in order to obtain a Hamiltonian graph. Consequently, every Hamiltonian cycle C in H passes through all k edges in E_2 . If we subdivide each edge in E_2 , we obtain a Hamiltonian graph H' . Now consider the graph G' obtained from G by subdividing each edge in E_1 , and identify the subdivision vertices of the corresponding edges in G' and H' . Notice that the union of G' and H' is a plane graph that contains H' , hence it is Hamiltonian. Consequently G' is subhamiltonian. By the characterization of Bernhart and Kainen [5], G admits an arc diagram with k biarcs, as claimed. ◀

In order to obtain a general statement about arc diagrams from Lemma 8, we need a bound on the number of edges to simultaneously flip in a given graph in order to make it Hamiltonian. Even the existence of such a simultaneous flip—regardless of the number of edges involved—is not obvious to begin with. For instance, consider a vertex that has linear degree in a triangulation T_1 and constant degree in a triangulation T_2 . As a single simultaneous flip can only change about half of the edges incident to a vertex, at least a logarithmic number of simultaneous flips is required to transform T_1 into T_2 [8].

Bose et al. [8] showed that every triangulation on $n \geq 6$ vertices can be transformed to a 4-connected (hence Hamiltonian) triangulation by a single simultaneous flip. However, no bound is known on the number of flipped edges, which leaves us with the trivial bound of $(2n - 4)/2 = n - 2$. Note that the resulting bound on the number of biarcs is similar to the one from Theorem 1, but there we could guarantee that all biarcs are monotone. Using Lemma 8 we do not have any control over the type of biarcs used.

The obvious open question is: Can we give a better bound on the number of edges needed in a simultaneous flip to a Hamiltonian triangulation?

4 Simultaneous Flip Distance to 4-connectivity

In this section we determine the maximum number of edges needed to transform an n -vertex triangulation into a Hamiltonian triangulation using a single simultaneous flip. Consider a triangulation $G = (V, E)$. As there is no 4-connected triangulation on fewer than six vertices, suppose that G has at least six vertices. We would like to transform G into a 4-connected triangulation by simultaneously flipping a set $F \subset E$ of edges such that all separating triangles are destroyed and none created. We use the following criterion by Bose et al. [8] to check whether a simultaneous flip produces a 4-connected triangulation.

► **Lemma 9** (Bose et al. [8]). *Let F be a set of edges in a triangulation G such that no two edges in F are incident to a common triangle, every edge in F is incident to a separating triangle, and for every separating triangle T there is at least one edge in F that is incident to T . Then F is simultaneously flippable in G and the resulting triangulation is 4-connected.*

For a simultaneously flippable set F of edges no face of the triangulation is incident to more than one edge. Recall that the edges of a triangulation G and its dual G^* are in one-to-one correspondence. Consequently, the dual edges of F form a matching in G^* . As all faces of a triangulation are triangles, G^* is cubic (3-regular). Moreover, every triangulation on $n \geq 4$ vertices is 3-connected and so its dual is 2-edge-connected (bridgeless). By a famous theorem of Tait the following statement is equivalent to the Four-Color Theorem:

► **Theorem 10** (Tait [7] Chapter 11). *Every bridgeless cubic planar graph admits a partition of the edge set into three perfect matchings.*

In particular, this applies to the dual of a triangulation. Call a set $F \subseteq E$ of edges of a triangulation $G = (V, E)$ a (perfect) *dual matching* if the corresponding set of edges in the dual graph G^* forms a (perfect) matching of G^* . While it is clear that a perfect dual matching contains exactly one edge of each facial triangle, this is not obvious for separating triangles. But it follows from a simple parity argument, as the following lemma shows.³

► **Lemma 11.** *Every perfect dual matching of a triangulation G contains an edge of every triangle of G .*

The last missing bit to prove Theorem 2 is an upper bound on the number of edges in a triangulation that can be incident to separating triangles.

► **Lemma 12.** *At most $2n - 7$ edges of a maximal planar graph on $n \geq 4$ vertices are incident to separating triangles. This bound is the best possible.*

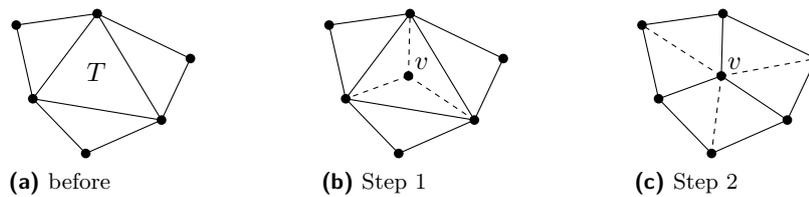
► **Theorem 2.** *Every maximal planar graph on $n \geq 6$ vertices can be transformed into a 4-connected maximal planar graph using a simultaneous flip of at most $\lfloor (2n - 7)/3 \rfloor$ edges.*

Proof. Consider a maximal planar graph G on n vertices. By Theorem 10 the $3n - 6$ edges of G can be partitioned into three perfect dual matchings $M_1, M_2,$ and M_3 , of $n - 2$ edges each. Let M'_i , for $i \in \{1, 2, 3\}$, denote the dual matching that results from removing all edges from M_i that are not incident to any separating triangle. By Lemma 12 at most $2n - 7$ edges of G are incident to separating triangles. Therefore, one of $M'_1, M'_2,$ and M'_3 contains at most $\lfloor (2n - 7)/3 \rfloor$ edges. By Lemma 9 these edges are simultaneously flippable and the resulting graph is 4-connected. ◀

5 Flip Distance to Hamiltonicity

With regard to arc diagrams, there is actually no reason to insist that the triangulation be 4-connected. In order to apply Lemma 8 we only need the triangulation to be Hamiltonian. Hence the obvious question: Can we always find a simultaneous flip of fewer than $2n/3$ edges to obtain a Hamiltonian triangulation? In this section we go one step further and in addition lift the restriction that the flip be simultaneous. Instead, an arbitrary sequence

³ Bose et al. [8] derive this property from the explicit Tait coloring. The statement here is slightly more general because it holds for every perfect dual matching.



■ **Figure 2** Example of a dummy flip.

of edge flips is allowed. In this case tight bounds are known if the goal is to obtain a 4-connected triangulation. Bose et al. [10] showed that $\lfloor (3n - 9)/5 \rfloor$ flips are always sufficient and sometimes $(3n - 10)/5$ flips are necessary to transform a given triangulation on n vertices into a 4-connected triangulation.

In general, a non-simultaneous flip sequence has no direct implication for arc diagrams. But if only edges of the original triangulation are flipped, then we can subdivide those edges rather than flipping them. In the resulting arc diagram only the subdivided edges may appear as biarcs. But a bound on the flip distance to a Hamiltonian triangulation is of independent interest. For instance, it is directly related to the current best upper bound on the diameter of the flip graph of combinatorial triangulations [10, 20, 21]. The argument uses a single so-called canonical triangulation and shows that every triangulation can be transformed into this canonical triangulation in two steps: First at most $\lfloor (3n - 9)/5 \rfloor$ flips are needed to obtain a 4-connected triangulation and then an additional at most $2n - 15$ flips are needed to transform any 4-connected triangulation into the canonical one. Combining two such flip sequences yields an upper bound of $5.2n - 33.6$ on the diameter of the flip graph [10]. The bound of $2n - 15$ flips for the second step is actually tight [20]. The corresponding bound for a triangulation that is Hamiltonian (but not necessarily 4-connected) is slightly worse only: It can be transformed into the canonical triangulation using at most $2n - 10$ flips [21]. Hence our focus is to improve the first step by showing that fewer flips are needed to guarantee a Hamiltonian triangulation than a 4-connected one.

► **Theorem 4.** *Every maximal planar graph on $n \geq 6$ vertices can be transformed into a Hamiltonian maximal planar graph using a sequence of at most $\lfloor (n - 3)/2 \rfloor$ edge flips.*

Proof outline. The proof is constructive and consists of two steps. In a first step we apply a sequence of elementary operations that transform a triangulation G into a 4-connected triangulation G' . An elementary operation is either a usual edge flip or a *dummy flip*, where a facial triangle T is subdivided into three triangles by inserting a new (dummy) vertex and then all three edges of T are flipped. All this will be done in such a way that G' becomes 4-connected and, therefore, contains a Hamiltonian cycle H' . We then remove all dummy vertices and construct a Hamiltonian cycle H'' resembling H' in the resulting triangulation G'' . Finally, we argue that G'' can be obtained from G with at most $n/2$ (usual) edge flips. Specifically, we show that each dummy flip can be implemented using at most two edge flips.

Dummy flips. Given a triangulation G on $n \geq 4$ vertices and a facial triangle T of G , a dummy flip of T transforms G as follows (Figure 2): First, insert a new (dummy) vertex v in the interior of face T and connect it to all three vertices of T . Note that T becomes a separating triangle in the resulting graph. Second, flip all three edges of T in an arbitrary order. Similarly to the usual flip operation, a dummy flip may create multiple edges. But we

will use this operation in specific situations only—as specified in the lemma below—where we can show that it produces a triangulation (that is, no multiple edges).

► **Lemma 13.** *Let G be a maximal planar graph and let T be a facial triangle of G such that every edge of T is incident to a separating triangle of G . Then the dummy flip operation of T in G produces no double edges and no new separating triangles.*

Step 1: 4-connectivity. Our main lemma to establish Theorem 4 is the following.

► **Lemma 14.** *Every maximal planar graph on $n \geq 6$ vertices can be transformed into a 4-connected maximal planar graph by a sequence of f flip and d dummy flip operations, for some $f, d \in \mathbb{N}$, such that $f + 2d \leq (n - 3)/2$.*

Recall that there are triangulations on n vertices that contain $\lfloor (3n - 9)/5 \rfloor$ pairwise edge-disjoint separating triangles [10, 17]. In this case, we need to flip away at least one edge from each separating triangle to reach 4-connectivity. Considering that a dummy flip operation flips three edges, the parameters in Lemma 14 satisfy $f + 3d \geq \lfloor (3n - 9)/5 \rfloor$. The crucial claim in Lemma 14 is that $f + 2d \leq (n - 3)/2$ is possible, and later we will show how to replace each dummy flip by two usual flips rather than three (Lemma 22).

The rest of this section is devoted to the proof of Lemma 14. We describe an algorithm that, given a triangulation G on $n \geq 6$ vertices, returns a sequence of f flip and d dummy flip operations that produces a 4-connected graph. The bound $6f + 12d \leq 3n - 9$ is established via the following charging scheme. Each edge of G , with the exception of the three edges of the outer face, receives one unit of credit. Each edge flip costs six units and each dummy flip costs fifteen units.

4-Block Decomposition. In our algorithm, we recursively process 4-connected subgraphs using the 4-block tree \mathcal{B} of G . By fixing an (arbitrary) plane embedding of G , we make \mathcal{B} a rooted tree such that the root is the 4-block that contains the boundary of the outer face of G . Every separating triangle T of G corresponds to an edge between two 4-blocks, where the parent lies in the exterior of T (plus T) and the child lies in the interior of T (plus T). For a 4-block G_i in \mathcal{B} denote by T_i the outer face of G_i , and denote by n_i the number of vertices of G_i minus three (the vertices of T_i). An edge of G_i is called an *interior* edge if it is not incident to the outer face T_i . For each 4-block G_i in \mathcal{B} we maintain counters f_i and d_i that denote the number of flips and dummy flips, respectively, that were used within G_i during the course of the algorithm. Initially $f_i = d_i = 0$, for every vertex G_i of \mathcal{B} .

The algorithm computes the sequence of flip and dummy flip operations incrementally, and maintains a current triangulation produced by the operations. Both the graph G and the 4-block decomposition \mathcal{B} change dynamically: when we flip an edge e of some separating triangle(s), all 4-blocks containing edge e merge into a single 4-block. At the end of the algorithm, the tree \mathcal{B} consists of a single 4-block that corresponds to the 4-connected graph G' . In order to avoid notational clutter, we always denote the current 4-block tree by \mathcal{B} . As an invariant (detailed below) we maintain that at each node of \mathcal{B} the number of interior edges (ignoring dummy edges) balances the cost of operations that were spent in this 4-block. As \mathcal{B} evolves, so does the graph $\mathcal{G}(\mathcal{B})$ represented by \mathcal{B} . This graph is the union of all nodes (4-blocks) in \mathcal{B} , where for any edge of \mathcal{B} the vertices and edges of the common triangle in the two endpoints (4-blocks) are identified.

Main loop. At every step, we take an arbitrary 4-block G_i on the penultimate level of \mathcal{B} , that is, G_i is not a leaf but all of its children are leaves. Let C_i denote the set of indices c

such that G_c is a child of G_i in \mathcal{B} , and denote $\mathcal{T}_i = \{T_c \mid c \in C_i\}$. The algorithm selects a sequence of edges of G_i to be flipped (or dummy flipped) in order to merge G_i with G_c , for all $c \in C_i$, into a new 4-block G_z . Denote the resulting 4-block tree by \mathcal{B}' . If no edge of T_i is flipped, then G_z is a leaf of \mathcal{B}' . But if an edge of T_i is flipped, then G_z may be an interior node of \mathcal{B}' .

If an edge e of T_i is flipped and G_i is not the root of \mathcal{B} , then more blocks may merge into G_z : The edge e is definitely shared with the parent of G_i in \mathcal{B} , but it may be shared with further ancestors as well. In addition, the edge e may belong to (at most) one sibling G_s of G_i and some descendants of G_s as well. We denote by J the set of all j such that G_j is a leaf of \mathcal{B} that is merged into G_z . Similarly, denote by Q the set of all q such that G_q is an interior vertex of \mathcal{B} that is merged into G_z , and denote by Q^+ the set of indices $q \in Q$ such that $f_q + d_q > 0$. Note that neither J nor Q are empty, because $C_i \subseteq J$ and $i \in Q$. However, we may have $Q^+ = \emptyset$.

Algorithmic preliminaries. In each iteration, we flip the edges of a dual matching of G_i (a 4-connector, defined below), but if \mathcal{T}_i forms a checkerboard (defined below), we substitute three of these flip operations by one dummy flip.

A *4-connector* for G_i is a dual matching of G_i that contains precisely one edge from every triangle in \mathcal{T}_i . A 4-connector always exists because every perfect dual matching (Theorem 10) is a 4-connector (Lemma 11). A *minimum 4-connector* is a 4-connector of minimum cardinality. By Lemma 9 we can flip the edges of a 4-connector in an arbitrary order, and the 4-blocks G_c , for all $c \in C_i$, will merge into G_i . We say that \mathcal{T}_i is a *checkerboard* if the triangles in \mathcal{T}_i are pairwise edge-disjoint and every interior edge of G_i belongs to some triangle in \mathcal{T}_i . If \mathcal{T}_i is a checkerboard, then we perform a dummy flip on a triangle F that is selected according to the following lemma.

► **Lemma 15.** *If \mathcal{T}_i is a checkerboard, then G_i has a facial triangle F that is adjacent to three triangles in \mathcal{T}_i that are not all adjacent to T_i .*

Algorithm 4CONNECT(G). Given a triangulation G , fix an arbitrary embedding of G . This embedding defines a rooted 4-block tree \mathcal{B} . While \mathcal{B} is not a singleton, do: (1) Consider an arbitrary vertex G_i at the penultimate level of \mathcal{B} . (2) Find a minimum 4-connector M for G_i that contains a maximum number of edges of T_i (that is, one, if possible). (3) If \mathcal{T}_i is *not* a checkerboard, then flip the edges of M in an arbitrary order. (4) Otherwise, let F be an arbitrary facial triangle as in Lemma 15. First apply a dummy flip to F . For each of the three triangles from \mathcal{T}_i adjacent to F , remove the incident edge from M , and flip all remaining edges of M in an arbitrary order. (5) Finally, update \mathcal{B} and $\mathcal{G}(\mathcal{B})$.

Correctness of the Algorithm. We show that the above algorithm turns an input triangulation G on n vertices into a 4-connected triangulation using a sequence of f flip and d dummy flip operations, for some $f, d \in \mathbb{N}$, such that $f + 2d \leq (n - 3)/2$. By Lemmata 9, 13, and 15, the operations described in the algorithm can be performed. In every step of the algorithm at least two nodes of the 4-block tree are merged. Therefore, after a finite number of steps we are left with a block tree that consists of a single 4-block G' .

► **Observation 16.** *For each vertex v created by a dummy flip operation in algorithm 4CONNECT(G), subsequent operations do not modify the six facial triangles incident to v .*

Free edges. It remains to bound the number of flip and dummy flip operations performed by the algorithm. An edge within some 4-block G_i of \mathcal{B} is *free* if it is not incident to any separating triangle of $\mathcal{G}(\mathcal{B})$. Free edges are a good measure of progress for our algorithm because our final goal is to arrive at a state where all edges of $\mathcal{G}(\mathcal{B})$ are free.

Invariants. As an invariant we maintain that every vertex G_i of \mathcal{B} satisfies the following condition:

- (F1) If G_i is the only vertex of \mathcal{B} , then it has at least $6f_i + 15d_i + 3$ free edges.
- (F2) If G_i is a leaf of \mathcal{B} that is not the root of \mathcal{B} , then G_i has at least $6f_i + 15d_i + 3$ free interior edges.
- (F3) If G_i is an interior vertex of \mathcal{B} , then either $f_i = d_i = 0$ or G_i has at least $6f_i + 15d_i + 1$ free interior edges.

Initially, (F2) holds for every leaf G_i of \mathcal{B} because all of the interior $3(n_i + 3) - 6 - 3 = 3n_i$ edges are free, $n_i \geq 1$, and $f_i = d_i = 0$. Trivially, (F3) holds for every interior vertex G_i of \mathcal{B} because $f_i = d_i = 0$. Having a certain number of edges in a plane graph implies having a certain number of vertices, as quantified by the following lemma.

Invariant maintenance. It remains to show that each step of the algorithm maintains invariants (F1)–(F3). We start bounding the size of a minimum 4-connector M of G_i .

► **Lemma 17.** *Let M be a minimum 4-connector for G_i that contains the maximum number of edges of T_i .*

- (1) *If s edges of G_i are each incident to two triangles from T_i , then $|M| \leq |C_i| - \lceil s/3 \rceil$.*
- (2) *If the triangles in T_i are pairwise edge-disjoint and $f_i + d_i > 0$, then $|M| \leq n_i - 2f_i - 5d_i$.*

In both cases, equality is possible only if M contains an edge of T_i .

Proof. (1) Partition the edge set of G_i into three dual matchings by Theorem 10. One of them, say D , contains at least $\lceil s/3 \rceil$ of the s edges that are incident to two triangles from T_i . If every triangle in T_i selects a unique incident edge from D , we obtain a 4-connector $R \subseteq D$ of size at most $|C_i| - \lceil s/3 \rceil$. The minimality of M yields $|M| \leq |R|$.

(2) By (F3), G_i has at least $6f_i + 15d_i + 1$ free interior edges. Therefore, at least one of the three perfect dual matchings guaranteed by Theorem 10 contains at least $\lceil (6f_i + 15d_i + 1)/3 \rceil = 2f_i + 5d_i + 1$ free interior edges. After removal of all those edges, the resulting dual matching R is still a 4-connector. By the minimality of M we have $|M| \leq |R| \leq (n_i + 1) - (2f_i + 5d_i + 1) = n_i - 2f_i - 5d_i$.

If M contains no edge of T_i , consider again the 4-connector D from above. It is obtained by removing free interior edges from a perfect dual matching of G_i . Hence D contains an edge of T_i . But then by the choice of M we know that D is not a minimum 4-connector and so we have $|M| \leq |C_i| - \lceil s/3 \rceil - 1$ and $|M| \leq n_i - 2f_i - 5d_i - 1$, respectively. ◀

► **Lemma 18.** *Let M be a minimum 4-connector for G_i and suppose that T_i is not a checkerboard. Then after flipping the edges in M , the resulting 4-block G_z contains at least $6f_z + 15d_z + 3\lceil s/3 \rceil + |Q^+|$ free interior edges, where s denotes the number of edges of G_i that are incident to two triangles from T_i .*

► **Lemma 19.** *Suppose that G_i together with all its children in \mathcal{B} is merged into a leaf G_z of \mathcal{B}' using f flips and d dummy flips. Then G_z contains at least $6(f_z - f_i - f) + 15(d_z - d_i - d) + 3n_i + 3|C_i| + 3|Q| - 3$ free interior edges.*

► **Lemma 20.** *Suppose that $f_i = d_i = 0$ and G_i along with all its children in \mathcal{B} is merged into an **interior node** G_z of \mathcal{B}' using f flips and d dummy flips. Then G_z contains at least $6(f_z - f) + 15(d_z - d) + 3n_i + 3|C_i| + 1$ free interior edges.*

► **Lemma 21.** *Suppose that \mathcal{T}_i is a **checkerboard**. Then G_z fulfills invariants (F1)–(F3).*

Case analysis. We now use Lemmata 18–21 to show that every step of algorithm 4CONNECT maintains (F1)–(F3). By Lemma 18 we may suppose in the following that the triangles in \mathcal{T}_i are pairwise edge-disjoint. Because if they are not, then the 4-block G_z obtained by flipping the edges in M by Lemma 18 fulfills one of (F1), (F2), or (F3), depending on the status of G_z in \mathcal{B}' . If \mathcal{T}_i is a checkerboard, then we are done by Lemma 21. Hence suppose that \mathcal{T}_i is not a checkerboard. Together with the fact that the triangles in \mathcal{T}_i are pairwise edge-disjoint, it follows that G_i has at least one free interior edge. As this edge appears in one of the three perfect dual matchings of Theorem 10, we conclude that $|M| \leq n_i$. For the remainder of the analysis we distinguish four cases.

Case 0: G_z is the only node of \mathcal{B}' . Then by Lemma 18 there are at least $6f_z + 15d_z + 3\lceil s/3 \rceil + |Q^+| \geq 6f_z + 15d_z$ free interior edges in G_z . Together with the three free non-interior edges of T_z this proves (F1).

Case 1: G_z is an interior vertex of \mathcal{B}' . If $Q^+ \neq \emptyset$, then (F3) holds by Lemma 18. Hence suppose that $Q^+ = \emptyset$ and so in particular $f_i = d_i = 0$. Using Lemma 20 with $f = |M|$ and $d = 0$ we obtain at least $6(f_z - f) + 15(d_z - d) + 3n_i + 3|C_i| + 1 \geq 6(f_z - f) + 15d_z + 3f + 3f + 1 = 6f_z + 15d_z + 1$ free interior edges in G_z , which proves (F3).

Case 2: G_z is a leaf of \mathcal{B}' (but not the only node) and $f_i + d_i > 0$. We distinguish two subcases. If M contains no edge of T_i , then Lemma 17(2) yields $f = |M| \leq n_i - 2f_i - 5d_i - 1$. By Lemma 19, we find at least

$$\begin{aligned} & 6(f_z - f_i - f) + 15(d_z - d_i) + 3n_i + 3|C_i| + 3|Q| - 3 \\ & \geq 6f_z - 6f_i - 3f - 3(n_i - 2f_i - 5d_i - 1) + 15(d_z - d_i) + 3n_i + 3f + 3|Q| - 3 \\ & = 6f_z + 15d_z + 3|Q| \end{aligned}$$

free interior edges in G_z . Since $i \in Q$, we have $|Q| \geq 1$ and (F2) follows.

Otherwise, M contains an edge of T_i . Then the parent G_p of G_i in \mathcal{B} is merged into G_z . Lemma 17(2) yields $f = |M| \leq n_i - 2f_i - 5d_i$. By Lemma 19 we find at least

$$\begin{aligned} & 6(f_z - f_i - f) + 15(d_z - d_i) + 3n_i + 3|C_i| + 3|Q| - 3 \\ & \geq 6f_z - 6f_i - 3f - 3(n_i - 2f_i - 5d_i) + 15(d_z - d_i) + 3n_i + 3f + 3|Q| - 3 \\ & = 6f_z + 15d_z + 3(|Q| - 1) \end{aligned}$$

free interior edges in G_z . Since $\{i, p\} \subseteq Q$, we have $|Q| \geq 2$ and (F2) follows.

Case 3: G_z is a leaf of \mathcal{B}' (but not the only node) and $f_i = d_i = 0$. We distinguish two subcases. If $|M| \leq n_i - 1$, then Lemma 19 guarantees $6f_z - 3f - 3f + 15d_z + 3n_i + 3|C_i| + 3|Q| - 3 \geq 6f_z - 3(n_i - 1) - 3f + 15d_z + 3n_i + 3f + 3|Q| - 3 = 6f_z + 15d_z + 3|Q|$ free interior edges in G_z , which together with $i \in Q$ proves (F2).

Otherwise, we have $f = |M| = n_i$. We claim that M contains an edge of T_i (otherwise \mathcal{T}_i would be a checkerboard). Suppose that M does not contain any edge of T_i . Then no

edge of T_i is incident to any triangle from \mathcal{T}_i . (Otherwise, we could replace the edge in M that is incident to such a triangle by the edge shared with T_i . As the triangles in \mathcal{T}_i are pairwise edge-disjoint, the replaced edge is incident to only one triangle from \mathcal{T}_i . The result is a 4-connector of the same size as M , but with an edge of T_i . This contradicts our choice of M .) Therefore, all $3n_i$ interior edges of G_i are incident to triangles in \mathcal{T}_i , and so \mathcal{T}_i is a checkerboard. This proves our claim.

As M contains an edge of T_i , the parent G_p of G_i in \mathcal{B} is merged into G_z as well. By Lemma 19 we find at least $6(f_z - n_i) + 15d_z + 3n_i + 3n_i + 3|Q| - 3 \geq 6f_z + 15d_z + 3(|Q| - 1)$ free interior edges in G_z , which together with $\{i, p\} \subseteq Q$ proves (F2).

Summary. In all cases we have shown that the resulting 4-block tree \mathcal{B}' satisfies our invariant. Thus the resulting 4-connected graph G' has $n + d$ vertices and at least $6f + 15d + 3$ edges, where f and d denote the number of flip and dummy flip operations, respectively, that were executed during the algorithm. Being a maximal planar graph, G' contains exactly $3(n + d) - 6$ edges. Therefore, $6f + 15d + 3 \leq 3(n + d) - 6$ and so $2f + 4d \leq n - 3$, as required. This completes the proof of Lemma 14.

Step 2: Eliminating dummy vertices. At this stage we have a 4-connected planar graph G' . By Tutte's Theorem such a graph is Hamiltonian, so consider some Hamiltonian cycle H' of G' . It remains to argue how G' and H' can be used to obtain a short sequence of edge flips that transform the original graph G into a Hamiltonian graph G'' . The following lemma in combination with Lemma 14 completes the proof for Theorem 4.

► **Lemma 22.** *Suppose that G' has been obtained from G using f flips and d dummy flips. Then G can be transformed into a Hamiltonian maximal planar graph using at most $f + 2d$ edge flips.*

Proof of Corollary 6. The following analogue of Lemma 22, combined with Lemma 14 and the characterization of Bernhart and Kainen [5], proves Corollary 6.

► **Lemma 23.** *Suppose that G' has been obtained from G using f subdivisions and d dummy flips. Then G can be subdivided into a subhamiltonian graph using at most $f + 2d$ subdivision vertices (at most one per edge).*

Acknowledgements This work began at the 12th Gremo's Workshop on Open Problems (GWOP), June 30–July 4, 2014, in Val Sinestra (GR), Switzerland. We thank all participants for the productive and positive atmosphere, and in particular Radoslav Fulek, Anna Gundert, Malte Milatz, Bettina Speckmann, Sebastian Stich, and Tibor Szabó for inspiring discussions.

References

- 1 Bernardo M. Ábrego, Oswin Aichholzer, Silvia Fernández-Merchant, Pedro Ramos, and Gelasio Salazar. Shellable drawings and the cylindrical crossing number of K_n . *Discrete Comput. Geom.*, 52(4):743–753, 2014.
- 2 Oswin Aichholzer, Clemens Huemer, and Hannes Krasser. Triangulations without pointed spanning trees. *Comput. Geom. Theory Appl.*, 40(1):79–83, 2008.
- 3 T. Alastair and J. Nicholson. Permutation procedure for minimising the number of crossings in a network. *Proc. IEE*, 115(1):21–26, 1968.

- 4 Patrizio Angelini, David Eppstein, Fabrizio Frati, Michael Kaufmann, Sylvain Lazard, Tamara Mchedlidze, Monique Teillaud, and Alexander Wolff. Universal point sets for drawing planar graphs with circular arcs. *Journal of Graph Algorithms and Applications*, 18(3):313–324, 2014.
- 5 Frank Bernhart and Paul C. Kainen. The book thickness of a graph. *J. Combin. Theory, Ser. B* 27:320–331, 1979.
- 6 J. Blažek and M. Koman. A minimal problem concerning complete plane graphs. In M. Fiedler, editor, *Theory of graphs and its applications*, pages 113–117. Czech. Acad. of Sci., 1964.
- 7 John Adrian Bondy and U. S. R. Murty. *Graph Theory*, volume 244 of *Graduate texts in Mathematics*. Springer, Berlin, 2008.
- 8 Prosenjit Bose, Jurek Czyzowicz, Zhicheng Gao, Pat Morin, and David R. Wood. Simultaneous diagonal flips in plane triangulations. *J. Graph Theory*, 54(4):307–330, 2007.
- 9 Prosenjit Bose and Ferran Hurtado. Flips in planar graphs. *Comput. Geom. Theory Appl.*, 42(1):60–80, 2009.
- 10 Prosenjit Bose, Dana Jansens, André van Renssen, Maria Saumell, and Sander Verdonschot. Making triangulations 4-connected using flips. *Comput. Geom. Theory Appl.*, 47(2):187–197, 2014.
- 11 Prosenjit Bose and Sander Verdonschot. A history of flips in combinatorial triangulations. In *Computational Geometry—XIV Spanish Meeting on Computational Geometry, EGC 2011*, volume 7579 of *LNCS*, pages 29–44. Springer, Berlin, 2012.
- 12 Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Berlin, 4 edition, 2010.
- 13 David Eppstein. Universal point sets for planar graph drawings with circular arcs. Presentation at the 25th Canadian Conference on Computational Geometry, Waterloo, Canada, <http://www.ics.uci.edu/~eppstein/pubs/AngEppFra-CCCG-13-slides.pdf>, 2013.
- 14 Emilio Di Giacomo, Walter Didimo, and Giuseppe Liotta. Spine and radial drawings. In Roberto Tamassia, editor, *Handbook of Graph Drawing and Visualization*, chapter 8, pages 247–284. CRC press, Boca Raton, FL, 2013.
- 15 Emilio Di Giacomo, Walter Didimo, Giuseppe Liotta, and Stephen K. Wismath. Curve-constrained drawings of planar graphs. *Comput. Geom. Theory Appl.*, 30(1):1–23, 2005.
- 16 Branko Grünbaum. *Convex Polytopes*, volume 221 of *Graduate texts in Mathematics*. Springer, Berlin, 2003.
- 17 Jochen Harant, Mirko Horňák, and Zdzisław Skupień. Separating 3-cycles in plane triangulations. *Discrete Math.*, 239(1):127–136, 2001.
- 18 Goos Kant. A more compact visibility representation. *Int. J. Comput. Geometry Appl.*, 7(3):197–210, 1997.
- 19 Michael Kaufmann and Roland Wiese. Embedding vertices at points: Few bends suffice for planar graphs. *J. Graph Algorithms Appl.*, 6(1):115–129, 2002.
- 20 Hideo Komuro. The diagonal flips of triangulations on the sphere. *Yokohama Math. J.*, 44:115–122, 1997.
- 21 Ryuichi Mori, Atsuhiko Nakamoto, and Katsuhiko Ota. Diagonal flips in Hamiltonian triangulations on the sphere. *Graphs Combin.*, 19(3):413–418, 2003.
- 22 János Pach and Rephael Wenger. Embedding planar graphs at fixed vertex locations. *Graphs Combin.*, 17:717–728, 2001.
- 23 Lionel Pournin. The diameter of associahedra. *Adv. Math.*, 259:13–42, 2014.
- 24 Thomas L. Saaty. The minimum number of intersections in complete graphs. *Proceedings of the National Academy of Sciences of the United States of America*, 52:688–690, 1964.
- 25 Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1:647–682, 1988.

- 26 Avi Wigderson. The complexity of the Hamiltonian circuit problem for maximal planar graphs. Technical Report 298, Princeton University, 1982.