

Balanced Circle Packings for Planar Graphs

Md. Jawaherul Alam¹, David Eppstein², Michael T. Goodrich²,
Stephen G. Kobourov¹, and Sergey Pupyrev^{1,3}

¹ Department of Computer Science, University of Arizona, Tucson, Arizona, USA

² Department of Computer Science, University of California, Irvine, California, USA

³ Institute of Mathematics and Computer Science, Ural Federal University, Russia

Abstract. We study *balanced* circle packings and circle-contact representations for planar graphs, where the ratio of the largest circle’s diameter to the smallest circle’s diameter is polynomial in the number of circles. We provide a number of positive and negative results for the existence of such balanced configurations.

1 Introduction

Circle packings are a frequently used and important tool in graph drawing [1, 5, 6, 18, 19, 27, 33, 36, 39]. In this application, they can be formalized using the notion of a *circle-contact representation* for a planar graph; this is a collection of interior-disjoint circles in \mathbb{R}^2 , corresponding one-for-one with the vertices of the graph, such that two vertices are adjacent if and only if their corresponding two circles are tangent to each other [23, 24]. Graphs with circle-contact representations are also known as “coin graphs” [40]. In a classic paper, Koebe [29] proved that every triangulated planar graph has a circle-contact representation, and this has been subsequently re-proved several times. Generalizing this, every planar graph has a circle-contact representation: we can triangulate the graph by adding “dummy” vertices connected to the existing vertices within each face, produce a circle-contact representation for this augmented graph, and then remove the circles corresponding to dummy vertices. It is not always possible to describe a circle-contact representation for a given graph by a symbolic formula involving radicals [4, 10], but they can nevertheless be constructed numerically and efficiently by polynomial-time iterative schemes [13, 35].

One of the drawbacks of some of these constructions, however, is that the sizes of the circles in some of these configurations may vary exponentially, leading to drawings with very high area or with portions that are so small that they are below the resolution of the display. For this reason, we are interested in *balanced* circle packings and circle-contact representations for planar graphs, where the ratio of the maximum and minimum diameters for the set of circles is polynomial in the number of vertices in the graph; see Fig. 1. Such drawings could be drawn with polynomial area, for instance, where the smallest circle determines the minimum resolution.

1.1 Related Work

There is a large body of work about representing planar graphs as contact graphs, where vertices are represented by geometrical objects and edges correspond to two objects touching in some pre-specified fashion. For example, Hliněný [23, 24] studies contact representations using curves and line segments as objects. Several authors have considered contact graphs of triangles of various types. For instance, de Fraysseix *et al.* [20] show that every planar graph has a triangle-contact representation, Badent *et al.* [2] show that partial planar 3-trees and some series-parallel graphs have contact representations with homothetic triangles, and Gonçalves *et al.* [22] prove that every 3-connected planar graph and its dual can be simultaneously represented by touching triangles (and they point out that 4-connected planar graphs also have contact representations with homothetic triangles). Also, Duncan *et al.* [17] show that every planar graph has a contact representation with convex hexagons all of whose sides have one of three possible slopes, and that hexagons




Fig. 1. Two planar graphs with possible circle-contact representations: (a) a representation that is not optimally balanced; (b) a perfectly-balanced representation

are necessary for some graphs, if convexity is required. With respect to balanced circle-contact representations, Breu and Kirkpatrick [9] show that it is NP-complete to test whether a graph has a perfectly-balanced circle-contact representation, in which every circle is the same size. Circle-contact graphs are related to *disk graphs* [25], which represent a graph by intersecting disks; unlike for circle contact graphs, the interiors of the disks are not required to be disjoint. Regarding the resolution of disk graphs, McDiarmid and Müller [34] show that there are n -vertex graphs such that in every realization by disks with integer radii, at least one coordinate or radius is $2^{2^{\Omega(n)}}$, and they also show that every disk graph can be realized by disks with integer coordinates and radii that are at most $2^{2^{O(n)}}$.

1.2 New Results

In this paper, we provide a number of positive and negative results regarding balanced circle-contact representations for planar graphs:

- Every planar graph with bounded maximum vertex degree and logarithmic outerplanarity admits a balanced circle-contact representation.
- There exist planar graphs with bounded maximum degree and linear outerplanarity, or with linear maximum degree and bounded outerplanarity, that do not admit a balanced circle-contact representation.
- Every tree admits a balanced circle-contact representation.
- Every outerpath admits a balanced circle-contact representation.
- Every cactus graph admits a balanced circle-contact representation.
- Every planar graph with bounded tree-depth admits a balanced circle-contact representation.

2 Bounded Degree and Logarithmic Outerplanarity

A plane graph (that is, a combinatorially fixed planar embedding of a planar graph) is *outerplanar* if all of its vertices are on the outer face. A k -*outerplanar graph* is defined recursively. As a base case, if a plane graph is outerplanar, then it is a 1-outerplanar graph. A plane graph is k -outerplanar, for $k > 1$, if the removal of all the outer vertices (and their incident edges) yields a graph such that each of the remaining components is $(k - 1)$ -outerplanar. The *outerplanarity* of a plane graph G is the minimum value for k such that G is k -outerplanar.

In this section, we show that every n -vertex plane graph with bounded maximum degree and $\mathcal{O}(\log n)$ outerplanarity admits a balanced circle-contact representation. We also show that these two restrictions are necessary, by demonstrating examples of planar graphs where either of these restrictions are violated and there is no balanced circle-contact representation.

2.1 Balanced Circle-Contact Representations

Theorem 1. *Every n -vertex k -outerplanar graph with maximum degree Δ admits a circle-contact representation where the ratio of the maximum and the minimum diameter is at most $f(\Delta)^{k+\log n}$, for some positive function f . In particular, when Δ is a fixed constant and k is $\mathcal{O}(\log n)$, this ratio is polynomial in n .*

In order to prove the theorem, we need the following result from [33]. See also [27].

Lemma 1 (Malitz-Papakostas). *The vertices of every triangulated planar graph G with the maximum degree Δ can be represented by nonoverlapping disks in the plane so that two disks are tangent to each other if and only if the corresponding vertices are adjacent, and for each two disks that are tangent to each other, the ratio of the radii of the smaller to the larger disk is at least $\alpha^{\Delta-2}$ with $\alpha = \frac{1}{3+2\sqrt{3}} \approx 0.15$.*

As a direct corollary, every maximal planar graph with maximum degree $\Delta = \mathcal{O}(1)$ and diameter $d = \mathcal{O}(\log n)$ has a balanced circle-contact representation. Theorem 1 goes beyond this.

Proof of Theorem 1: To prove the claim, it is sufficient to show how to augment a given k -outerplanar graph into a maximal planar graph with additional vertices so that its maximum degree remains $\mathcal{O}(\Delta)$ and its diameter becomes $\mathcal{O}(k + \log n)$. By Lemma 1, the resulting graph admits a circular contact representation with the given bounds on the ratio of radii. Removing the circles corresponding to the added vertices yields the desired balanced representation of the original graph.

Let G be an n -vertex k -outerplanar graph with the maximum degree Δ . If the outerplanarity k of G is bounded by a constant, we can easily augment G to logarithmic diameter, preserving its constant maximum degree, as follows. Inside each non-triangular face f of G , insert a balanced binary tree with $\lceil \log |f| \rceil$ levels and $|f|$ leaves and then triangulate the remaining non-triangular faces by inserting an *outerpath* (an outerplanar graph whose weak dual is a path) with constant maximum degree; see Fig. 2. However, such an augmentation results in a maximal planar graph with the diameter $d = \mathcal{O}(k \log n)$, which does not yield a balanced circle-contact representation when k is non-constant. For $k = \Omega(\log n)$, we present a different augmentation to achieve the diameter $d = \mathcal{O}(k + \log n)$ in the resulting graph.

We augment the graph using *weight-balanced binary trees*. Let T be a binary tree with leaves $l_1, l_2, \dots, l_{|f|}$ and a prespecified weight w_i assigned to each leaf l_i . The tree T is *weight-balanced* if the depth of each leaf l_i in T is $\mathcal{O}(\lceil \log(W/w_i) \rceil)$, where $W = \sum_{i=1}^{|f|} w_i$. There exist several algorithms for producing a weight-balanced binary tree with positive integer weights defined on its leaves [21, 38].




Fig. 2. (a) A face, (b) augmentation with a balanced binary tree, (c) triangulation with grey edges

To augment G , we label each vertex v of G with the number $l + 1$, where l is the number of outer cycles that need to be removed before v becomes an outer vertex. By our assumption that the outerplanarity of G is k , the label of every vertex is at most k . It follows from this labeling that, for each vertex v of G with label $l > 1$, there exists a face f containing v such that f has at least one vertex of label $l - 1$ and such that all the vertices on f have label either l or $l - 1$. We insert a weight-balanced binary tree inside f ; we choose an arbitrary vertex of f with label $l - 1$ as the root of the tree, and a subset of vertices with label l as the leaves; see Fig. 3. We construct these trees inside the different faces in such a way that each vertex of G with label $l > 1$ becomes a leaf in exactly one of the trees. Finally, we insert another weight-balanced tree T_0 on the outer face containing all the outer vertices as the leaves. Note that we have yet to specify the weights we assign to these leaves for producing the weight-balanced trees. By the construction, the union of all these trees forms a connected spanning tree of G ; we can consider the root of T_0 to be the root of the whole spanning tree.

Let us now specify the weights assigned to the leaves of the different weight-balanced trees. We label each tree with the label of its root, and define the weights for the leaves of each tree in a bottom-up ordering, by decreasing order of the labels of the trees. In a tree T with label $l = (k - 1)$, all the leaves have label k and are not the root of any other tree; we assign each of these leaves the weight 1. In this case, the total weight of T is the number of its leaves. Similarly, for a tree with label $l < k - 1$, we assign a weight of 1 to those leaves v that do not have any tree rooted at them; otherwise, if v is the root of a tree T_v with label $l + 1$, the weight of v is the total weight of T_v . The total weight of T is defined as the summation of the weights of all its leaves.

Now, for each vertex v of G , the distance to v from the root r of T_0 is $\mathcal{O}(k + \log n)$. Indeed, assume that $v = u_l$ is a vertex with label l and $u_{l-1}, \dots, u_1, u_0 = r$ are the root vertices of the successive weight-balanced trees $T_{u_{l-1}}, \dots, T_{u_1}, T_0$ with labels $l - 1, \dots, 1, 0$, respectively on the way from v to r ; see Fig. 3. Then the distance from v to r is

$$\mathcal{O}(\lceil \log w(r)/w(u_1) \rceil) + \mathcal{O}(\lceil \log w(u_1)/w(u_2) \rceil) + \dots + \mathcal{O}(\lceil \log w(u_{l-1})/w(v) \rceil) = \mathcal{O}(k + \log w(r)).$$

Here $w(u_i)$ denotes the weight of vertex u_i as the root; $w(r)$ is the weight of the root of T_0 , which is equal to the total number of vertices, n , in G . Therefore, the diameter of the augmented graph is $\mathcal{O}(k + \log n)$, where the first term, k , comes from the ceilings in the summation. Finally, we triangulate the graph by inserting outerpaths with constant maximum degree inside each non-triangular face to obtain a maximal planar graph with constant maximum degree and $\mathcal{O}(k + \log n)$ diameter. The result follows from Lemma 1. \square




Fig. 3. Augmentation of G with a weight-balanced binary trees




Fig. 4. Planar graphs with no balanced circle-contact representation: (a) the nested-triangles graph [16]; (b) a 2-outerplanar graph

2.2 Negative Results

Next we show that, for a graph with unbounded maximum degree or unbounded outerplanarity, there might not be a balanced circle-contact representation with circles.

Lemma 2. *There is no balanced circle-contact representation for the graphs in Fig. 4.*

Recall that Lemma 2 states that there is no balanced circle-contact representation for the graphs in Fig. 4. We need the following auxiliary lemma to prove this.




Fig. 5. Illustration for the proof of Lemma 3

Lemma 3. *Let C_1, C_2, C_3 be three interior-disjoint circles with radii r_1, r_2, r_3 , respectively that touch each other mutually at the three points p_1, p_2 and p_3 . Then every circle that lies completely inside the curvilinear triangle $p_1p_2p_3$ bounded by the three circles has radius at most half the median of r_1, r_2 and r_3 .*

Proof: Assume without loss of generality that $r_1 \leq r_2 \leq r_3$, so that the median radius is r_2 . Let C be a circle lying completely inside $p_1p_2p_3$ with radius r ; see Fig. 5. The radius of C is maximum when C touches all the three circles C_1, C_2 and C_3 . Indeed if C does not touch at least one of the three circles, say C_3 , then

it can be moved slightly towards C_3 and then its radius can be increased by a positive amount so that it still lies inside the region. Thus, we can assume without loss of generality that C touches all the three circles C_1, C_2, C_3 . Then by Descartes' Circle Theorem [14], we have

$$\left(\frac{1}{r_1} + \frac{1}{r_2} + \frac{1}{r_3} + \frac{1}{r}\right)^2 = 2\left(\left(\frac{1}{r_1}\right)^2 + \left(\frac{1}{r_2}\right)^2 + \left(\frac{1}{r_3}\right)^2 + \left(\frac{1}{r}\right)^2\right),$$

which simplifies to

$$\frac{1}{r} = \frac{1}{r_1} + \frac{1}{r_2} + \frac{1}{r_3} + 2\sqrt{\frac{1}{r_1 r_2} + \frac{1}{r_2 r_3} + \frac{1}{r_3 r_1}}.$$

(The other solution to this quadratic equation, with a negative sign on the square root term, gives the radius for the larger circle surrounding and tangent to C_1, C_2 , and C_3 .) Removing two positive terms from this equation gives the inequality $\frac{1}{r} \geq \frac{1}{r_1} + \frac{1}{r_2} \geq \frac{2}{r_2}$ or, equivalently, $r \leq r_2/2$. \square

Proof of Lemma 2: We prove the claim only for the 2-outerplanar graph in Fig. 4(b), since the proof for the nested-triangles graph in Fig. 4(a) follows a similar argument.

Because this graph is 3-connected, its planar embedding is unique up to the choice of the outer face. To begin with, assume that the outer face has been chosen as triangle xa_1b_1 . Let Γ be a contact representation for the graph with circles. We prove that Γ is not balanced. Let r be the radius of the circle C representing x and let r_i and r'_i , $1 \leq i \leq t$ be the radii of the circles C_i and C'_i representing a_i and b_i , respectively. Since Γ is a contact representation, for any value $i \in \{1, \dots, t-1\}$, C_i and C'_i are completely inside the region created by three circles C, C_{i+1} and C'_{i+1} . Therefore by Lemma 3, the values of r_i and r'_i are both less than $\frac{1}{2} \max(r_{i+1}, r'_{i+1})$. Since $t = (n-1)/2$, this implies that Γ is not balanced.

In the general case, in which the outer face can be an arbitrary triangle of the graph, the proof is similar. No matter which triangle is chosen as the outer face, the remaining subgraph between that face and one of the two ends of the zigzag chain of triangles in the graph has the same form as the whole graph, with at least half as many vertices, and the result follows in the same way for this subgraph. \square

Lemma 2 shows the tightness of the two conditions for balanced circle-contact representations in Theorem 1. Note that the example of the graph in Fig. 4(b) can be extended for any specified maximum degree, by adding a simple path to the high-degree vertex. Furthermore, the example is a 2-outerplanar graph with no balanced circle-contact representation.

3 Trees and Outerplanar Graphs

Here we address balanced circle-contact representation for trees and outerplanar graphs. We prove that every tree, every outerpath, and every cactus graph admits a balanced circle-contact representation.

3.1 Trees

Theorem 2. *Every tree has a balanced circle-contact representation. Such a representation can be found in linear time.*

Proof: We first find a contact representation Γ of a given tree T with squares such that the ratio of the maximum and the minimum sizes for the squares is polynomial in the number of vertices n in T . To this end, we consider T as a rooted tree with an arbitrary vertex r as the root. Then we construct a contact representation of T with squares where each vertex v of T is represented by a square $R(v)$ such that $R(v)$ touches the square for its parent by its top side and it touches all the squares for its children by its bottom side; see Figs. 6(a) and 6(b). We choose the size of $R(v)$ as $l(v) + \varepsilon(n(v) - 1)$, where $\varepsilon > 0$ is a small




Fig. 6. Construction of a balanced circle-contact representation

positive constant and $n(v)$ and $l(v)$ denote the number of vertices and the number of leaves in the subtree of T rooted at v . In particular, the size of $R(v)$ is 1 when v is a leaf. If v is not a leaf, then suppose v_1, \dots, v_d are the children of v in the counterclockwise order around v . Then we place the squares $R(v_1), \dots, R(v_d)$ from left-to-right touching the bottom side of $R(v)$ such that for each $i \in \{1, \dots, d-1\}$, $R(v_{i+1})$ is placed ε unit to the right of $R(v_i)$; see Fig. 6(b). There is sufficient space to place all these squares in the bottom side of $R(v)$, since $n(v) = (\sum_{i=1}^d n(v_i)) - 1$ and $l(v) = \sum_{i=1}^d l(v_i)$. The representation contains no crossings or unwanted contacts since for each vertex v , the representation of the subtrees rooted at v is bounded in the left and right side by the two sides of $R(v)$, and all the subtrees rooted at the children of v are in disjoint regions ε unit away from each other. The size of the smallest square is 1, while the size of the largest square (for the root) has size $l(T) + \varepsilon(n-1) = \mathcal{O}(n)$, where $l(T)$ is the number of leaves in T .

Using Γ , we find a balanced circle-contact representation of T as follows. We replace each square $R(v)$, representing vertex v , by an inscribed circle of $R(v)$; see Fig. 6(c). The operation removes some contacts from the representation. We re-create these contacts by a top-down traversal of T and moving each circle upward until it touches its parent. Note that a given circle will not touch or intersect any circle other than the circles for its parent and its children, as for every vertex in the infinite strip between its leftmost and rightmost point for its circle, the closest circle in the upward direction is its parent's one. Thus, we obtain a contact representation of T with circles. The representation is balanced since the diameter for every circle is equal to the side-length for its square and we started with a balanced representation Γ .

The linear running time can be achieved by a linear-time traversal of T . First, by a bottom-up traversal of T , we compute the values $n(v)$ and $l(v)$ for each vertex v of T . Using the values for each vertex, we compute the square-contact representation for T by a linear-time top-down traversal of T . Finally, in another top-down traversal of T , for each vertex v of T , we can compute the exact translation required for the inscribed circles of $R(v)$ to touch the parent circle. \square

3.2 Cactus Graphs

Let us now describe how to compute a balanced circle-contact representation for a *cactus* graph, which is a connected graph in which every biconnected component is either an edge or a cycle. We use the algorithm described in the proof of Theorem 2, and we call it **Draw Tree**.

Let T be a rooted tree with a plane embedding. For each vertex v of T , add an edge between every pair of the children of v that are consecutive in the clockwise order around v . Call the resulting graph an *augmented fan-tree* for T . Clearly for any rooted tree T , the augmented fan-tree is outerplanar. We call an outerplanar graph a *fan-tree* graph if it is an augmented fan-tree for some rooted tree. A *star* is the complete




Fig. 7. (a) A star T and a contact representation of T with circles; (b) a subgraph of the fan for T and its contact representation with circles

bipartite graph $K_{1,n-1}$. The center of a star is the vertex that is adjacent to every other vertex. An augmented fan-tree for a star is obtained by taking the center as the root. Thus, an augmented fan-tree for a star is a *fan*. The *center* of a fan is again the vertex adjacent to all the other vertices.

Lemma 4. *Every subgraph of a fan admits a contact representation with circles in which, for each circle $c(v)$ representing a vertex v other than the center, the vertical strip containing $c(v)$ is empty above $c(v)$.*

Proof: Let G be a subgraph of a fan and let T be the star contained in the fan. We now use the contact representation Γ of T obtained by **Draw_Tree** to compute a representation for G . Consider the square-contact representation computed for T in the algorithm. This defines a vertical strip for each circle $c(v)$ in Γ representing a vertex v , and for all the vertices other than the center, these strips are disjoint; see Fig. 7(a). Call the left and right boundary of this strip the *left-* and *right-line* for c , respectively.

We now consider a set S of circles, one for each vertex of G other than the center, with the following properties:

- (P1) The circles are interior-disjoint.
- (P2) Each circle $c'(v)$ representing a vertex v spans the entire width of the vertical strip for v , and the vertical strip above $c'(v)$ is empty.
- (P3) For each vertex v , the circle $c'(v)$ touches the circle c_0 representing the center in Γ if v is adjacent to the center; otherwise, $c'(v)$ is exactly ε distance away from c_0 , for some fixed constant $\varepsilon > 0$.
- (P4) If a vertex v is not adjacent to the vertex on its left (or if v is the leftmost vertex), then the leftmost point of $c'(v)$ is on the left-line of v ; similarly, if v is not adjacent to the vertex on its right (or if v is the rightmost vertex), then the rightmost point of $c'(v)$ is on the right-line of v .
- (P5) The sizes for the circles are maximal with respect to the above properties.

Note that there exists a set of circles with the properties (P1)–(P4); in particular, the set of circles in Γ representing the vertices of T other than the center is such a set. We now claim that the set S of circles with properties (P1)–(P5) together with the circle c_0 gives a contact representation for G ; see Fig. 7(b). First note that a circle $c'(v)$ cannot touch any circle other than c_0 and the two circles $c(v_l)$ and $c(v_r)$ representing the




Fig. 8. Illustration for the proof of Lemma 4: if the circles for x and y do not touch each other, at least one can be increased in size

vertices v_l and v_r on its left and right, respectively. Indeed, it cannot pass the vertical strip for v_l and v_r above them due to (P2) and behind them due to (P3). Furthermore, the ε distance between c_0 and the circles for vertices non-adjacent to the center and the restriction on the left and right side in (P4) ensures that there is no extra contact. Hence, it is sufficient to show that for each edge in G , we have the contact between the corresponding circles.

Since each circle $c'(v)$ is maximal in size, it must touch at least three objects. One of them is either the circle c_0 or the ε offset line for c_0 . Thus, if v_l and v_r are the left and right neighbors of v (if any), then $c'(v)$ must touch two of the followings: (i) $c'(v_l)$ (or the left line of v if v_l does not exist), (ii) the right line for v_l , (iii) $c'(v_r)$ (or the right line of v if v_r does not exist), and (iv) the left line for v_r . Assume without loss of generality that both v_l and v_r exist for v . Then if $c'(v)$ touches both $c'(v_l)$ and $c'(v_r)$, we have the desired contacts for v . Therefore, for a desired contact of $c'(v)$ to be absent, either $c'(v)$ touches both $c'(v_l)$ and the right-line of v_l (and misses the contact with $c'(v_r)$), or it touches both $c'(v_r)$ and the left-line of v_r (and misses the contact with $c'(v_l)$).

Assume, for the sake of a contradiction, that there are two consecutive vertices x and y that are adjacent in G but $c'(x)$ and $c'(y)$ do not touch each other. Let l and r be the vertices to the left of x and to the right of y , respectively. Then it must be the case that x touches both $c'(l)$ and the right line for l and y touches both $c'(r)$ and the left line of r ; see Fig. 8(a). One can then increase the size of either $c'(x)$ or $c'(y)$ (say $c'(y)$) such that it now touches $c'(x)$ and the left-line for r (but not $c'(r)$), a contradiction to the maximality for the circles; see Fig. 8(b). \square

Using the lemma, we can obtain a quadratic-time algorithm as follows. Given a subgraph G of a fan, compute the balanced circle-contact representation Γ for the corresponding star T using **Draw_Tree**. Then pick the vertices of T other than the center in an arbitrary order and for each vertex v , replace the circle $c(v)$ in Γ by a circle of maximum size that does not violate any of the properties (P1)–(P4) in the proof of Lemma 4. This takes a linear time. Now for every edge (x, y) for which $c(x)$ and $c(y)$ do not touch, replace one of the two circles (say, $c(y)$) with a circle that touches $c(x)$ as in Fig. 8(b). Note that this may result in a loss of a contact between $c(y)$ and the circle to its right. We perform a similar operation for the circle to the right of $c(y)$, then possibly for the circle on its right and so on, until all missing contact are repaired. This process requires linear time per edge; hence, the total running time to compute the desired contact representation is quadratic. The contact representation is balanced since the representation obtained by **Draw_Tree** is balanced and afterwards we only increase the size of circles that are not of the largest size.

Theorem 3. *Every n -vertex fan-tree graph has a balanced circle-contact representation. Such a representation can be found in $\mathcal{O}(n^2)$ time.*




Fig. 9. (a) A cactus graph G ; (b) augmenting G to a fan-tree so that the directed edges form a rooted tree and are oriented towards the root

Proof: Let G be a fan-tree graph and let T be the corresponding tree for which G is the augmented fan-tree. Using **Draw_Tree**, we first obtain a balanced circle-contact representation of T . As in the proof of Lemma 4, this defines a vertical strip for each vertex in T . In a top-down traversal of T , we can find a contact representation of G with circles by repeating the quadratic-time algorithm for the subgraphs of fans. Hence, the total complexity is $\sum_{v \in V(T)} \deg_T^2(v) = \mathcal{O}(n^2)$. \square

As a corollary of Theorem 3, we obtain an algorithm for creating balanced circle-contact representation of a cactus graph.

Corollary 1. *Every n -vertex cactus graph has a balanced circle-contact representation. Such a representation can be found in $\mathcal{O}(n^2)$ time.*

Proof: Given cactus graph G , choose a root vertex v arbitrarily. For each cycle C of G , add an edge from each vertex of C to the (unique) closest vertex to v in C (Fig. 9). The resulting supergraph of G is a fan-tree; the result follows by Theorem 3. \square

3.3 Outerpaths

We end this section with a linear-time algorithm for balanced circle-contact representation of outerpaths. Recall that *outerpaths* are outerplanar graphs whose weak dual is a path.

Theorem 4. *Every outerpath has a balanced circle-contact representation. Such a representation can be found in linear time.*

Proof: Let G be an outerpath. Assume for now that all the faces of G are triangles. Then the consecutive faces of G along the weak dual path can be partitioned into a sequence of fans. Define a path P in G that consists of all the edges that are shared by two maximal fans in G , along with the two terminal edges; see Fig. 10(a). We will find a contact representation of G in which each vertex of P is realized by a unit circle. Number the vertices from one end to the other along P as $1, 2, \dots, p$. Then for each vertex i on P , draw a unit circle C_i with the center at (i, y) , where $y = 0$ if i is even, and otherwise $y = \sqrt{3}/2$; see Fig. 10(b).

In this drawing, each circle C_i touches the circles C_{i-2} , C_{i-1} , C_{i+1} and C_{i+2} . Thus we have some additional and unwanted tangencies in this realization: namely, for some circles C_i , the tangency with C_{i+2}




Fig. 10. (a) An outerpath G ; (b)–(c) construction of a balanced circle-contact representation for G

(or symmetrically with C_{i-1}) does not correspond to an edge in G . When i is odd, we may remove an extra tangency between circles C_i and C_{i+2} by rotating C_{i+2} , and all circles to the left of it, clockwise by a small angle θ around the center of circle C_{i+1} . Similarly, when i is even, we may remove an extra tangency between circles C_i and C_{i+2} by rotating C_{i+2} , and all circles to the left of it, counterclockwise by a small angle θ around the center of C_{i+1} . We choose θ to be $90^\circ/n$ so that these rotations do not create any overlap between the other circles.

Next, we draw the circles representing the remaining vertices of each fan, within the regions created by these rotations; see Fig. 10(c). It is easy to see that the sizes for these circles are at least $\Omega(1/n^2)$. Finally, if G has some non-triangular faces, then we can add additional edges to augment G into triangulated outerpath G' , and perform the construction above. A small perturbation of the construction suffices to remove its extra adjacencies. Clearly, the running time is linear. \square

4 Bounded Tree-Depth

The *tree-depth* of a graph [37] is a measure of its complexity, related to but weaker than its treewidth. Thus, problems that are hard even for graphs of bounded treewidth, such as testing 1-planarity, can be solved efficiently for graphs of bounded tree-depth [3], and constructions that might be impossible for graphs of bounded treewidth, such as finding balanced circle-contact representations even when the vertex degree is unbounded, might become possible for bounded tree-depth.

One definition of tree-depth involves a certain kind of forest, which we call an *ancestral forest* for a given graph G ; see Fig. 11(a). An ancestral forest is a forest F having the vertices of G as its vertices, such that every edge of G connects a vertex v with one of the ancestors of v in the forest. For instance, every depth-first-search forest is ancestral; however, unlike depth-first search forests, an ancestral forest is allowed

to have edges that do not belong to G . The tree-depth of G is then the minimum height of an ancestral forest for G , where the height of a tree is the maximum number of vertices on a root-leaf path. Tree-depth can also be defined inductively (a connected graph has tree-depth $\leq k$ if it is possible to remove a vertex so that each remaining connected component has tree-depth $\leq k - 1$) or alternatively it can be defined in terms of the existence of certain graph colorings, called *centered colorings* [37].




Fig. 11. (a) A planar graph of tree-depth 3 and one of its possible ancestral trees, (b) the longest path in a graph of tree-depth 3 can have at most $7 = 2^3 - 1$ vertices

If the number of vertices in the longest simple path in a graph G is ℓ , then G has tree-depth at most ℓ , for this will be the maximum possible height of a depth-first-search forest. In the other direction, the tree-depth of G must be at least $\log_2(\ell + 1)$, because the path itself requires this height in any ancestral forest; see Fig. 11(b). Therefore, for any infinite family of graphs, it is equivalent to state that either the family has bounded tree-depth or that it has bounded longest path length [37].

4.1 Tree structures for graph connectivity

Our results in this section depend on a characterization of the planar graphs of bounded tree-depth, which we state in terms of two standard data structures for graph connectivity, the block-cut tree and the SPQR tree.

In any undirected graph G , one can define an equivalence relation on the edges of G according to which two distinct edges are equivalent whenever they belong to a simple cycle of G . The equivalence classes of this relation form subgraphs of G , called the *blocks* or *2-vertex-connected components* of G , and are separated from each other by *cut vertices*, vertices that belong to more than one block. The collection of blocks and cut vertices may be represented by a tree structure (or, in a disconnected graph, a forest) called the *block-cut tree* [28]: this tree has a node for each block and each cut vertex of G , and it has an edge connecting each cut vertex to the blocks that contain it. A graph is *2-vertex-connected* if it has only a single block (and no isolated vertices).

Similarly, a graph is *3-vertex-connected* if it does not have any *separation pairs*, pairs of vertices whose removal would disconnect the graph. If a graph G is 2-connected but not 3-connected, it may be represented by a structure called an SPQR tree [15]. In the formulation of SPQR trees that we use here, there are nodes of three types, each of which is associated with a *triconnected component* of the given graph G , graphs derived from G and having as their vertices subsets of the vertices of G [26, 32]. In an *S*-node of an SPQR tree, the associated graph is a simple cycle. In a *P*-node, the associated graph is a *dipole graph*, a two-vertex multigraph with three or more edges connecting its two vertices. And in an *R*-node of an SPQR tree, the associated graph is a 3-vertex-connected graph that is neither a cycle nor a dipole. In each of the graphs associated with the nodes of an SPQR tree, some of the edges are *real* (belonging to the original graph G)

and some are *virtual*. Each edge of the SPQR tree identifies a pair of virtual edges from the graphs associated with its two nodes, and each virtual edge is identified in this way by one SPQR tree edge. The original graph G can be reconstructed by gluing together each identified pair of virtual edges (merging each endpoint of one edge with an endpoint of the other edge to form a supervertex), and then deleting these edges. With the additional constraints that S nodes not be adjacent to other S nodes, and that P nodes not be adjacent to other P nodes, this decomposition is uniquely determined from G . A 2-connected graph G is planar if and only if each of the graphs associated with the nodes of its SPQR tree is planar, and the SPQR tree can be used to represent all of its possible planar embeddings.

4.2 Characterization of planar bounded tree-depth

We use the block-cut tree and SPQR tree structures to characterize the planar graphs of bounded tree-depth, as follows.

Lemma 5. *A family, \mathcal{F} , of planar graphs has bounded tree-depth if and only if the following three conditions are satisfied:*

- *The block-cut trees of the graphs in \mathcal{F} have bounded height.*
- *The SPQR trees of the blocks of the graphs in \mathcal{F} have bounded height.*
- *For each node of an SPQR tree of a block of a graph in \mathcal{F} , the number of vertices in the graph associated with the node is bounded.*

Proof: The requirements that the block-cut trees and SPQR trees have bounded height are clearly both necessary for the tree-depth to be bounded. For, if the graphs in \mathcal{F} had block-cut trees or SPQR trees that could be arbitrarily high, we could follow the structure of these trees to find paths in the graphs that were arbitrarily long, which in turn would imply that the tree-depth (which is at least logarithmic in the path length) would necessarily itself be arbitrarily high. In an SPQR tree of any graph, a P-node cannot have an unbounded number of vertices. An S-node with many vertices would necessarily again lead to a long path, as any virtual edge in the cycle associated with the S-node can be replaced by a path of real edges. To show that it is also necessary to have a bounded number of vertices in the graphs associated with R nodes, we invoke a result of Chen and Yu [12], who proved that a 3-connected planar graph with k nodes necessarily contains a cycle of length $\Omega(k^{\log_2 3})$. If the graphs associated with R-nodes could have nonconstant values of k , this would again necessarily imply the existence of paths of non-constant length in the graphs of \mathcal{F} , and therefore it would also imply that the tree-depth of \mathcal{F} would not be bounded.

In the other direction, suppose that we have a family \mathcal{F} of graphs whose block-cut tree height is at most b , SPQR tree height is at most h , and SPQR node vertex count is at most s . Then, any simple path Π in such a graph can pass through at most b blocks: Π must follow a path in the block-cut tree, alternating between blocks and cut vertices, and the total number of blocks and cut vertices in any such path is at most $2b - 1$. Within any block of the block-cut tree, Π must follow a walk through the SPQR tree that passes through each node at most s times, because each time it passes through a node it uses up one of the vertices in the associated graph. Therefore, the subtree of the SPQR tree followed by Π has maximum degree at most $s + 1$, and can include at most $\mathcal{O}((s + 1)^h)$ nodes. Within each of these nodes, Π can have at most s vertices per node. Therefore, the total number of nodes in Π can be at most $\mathcal{O}(sb(s + 1)^h) = \mathcal{O}(1)$. Any depth-first-search forest of a graph in \mathcal{F} has height bounded by this quantity, showing that all graphs in \mathcal{F} have bounded tree-depth. \square




Fig. 12. Lemma 6: Any two configurations of two tangent circles and a marked point on one of the two circles can be Möbius-transformed into each other

4.3 Möbius gluing

For any planar graph G of bounded tree-depth, we will form a circle-contact representation of G by decomposing G into its block-cut tree, decomposing each block into its SPQR tree, using Koebe’s circle packing theorem to find a circle-contact representation for each node of each SPQR tree, and gluing these packings together to give a representation of G' . Because each SPQR tree node has an associated graph of bounded size, the packings that we glue together will be balanced. When we glue two packings together, we will ensure that we only form a polynomial imbalance in the ratio of sizes of their circles. This, together with the bounded height of the trees that guide our overall gluing strategy, will ensure that the eventual packing we construct is balanced.

To glue two packings together, we use *Möbius transformations*, a group of geometric transformations of the plane (plus one point at infinity) that includes inversion by a circle as well as the familiar Euclidean translations, rotations, and scaling transformations. These transformations take circles to circles (allowing for some degenerate cases in which a line must be interpreted as an infinite-radius circle through the point at infinity) and preserve contacts between circles. The specific transformation we need is described by the following lemma.

Lemma 6. *Let X and Y be two geometric configurations, each consisting of two tangent circles and a marked point on one but not both of the circles. Then there exists a Möbius transformation that takes X to Y .*

Proof: It is equivalent to state that there exists a Möbius transformation that takes both of these configurations to the same configuration as each other. Both X and Y may be transformed to a pair of parallel lines, by an inversion through a circle centered on their point of tangency. A scaling transformation makes these two pairs of parallel lines the same distance apart from each other, a rotation makes all four lines parallel, a reflection (if necessary) makes both transformed configurations have the marked point on the upper of the two lines, and then a translation matches the positions of the marked points. \square

Lemma 6 is illustrated in Fig. 12. To control the sizes of the transformed packings generated by this lemma, we use the following lemma:

Lemma 7. *Let two circles C_0 and C_1 be given, and suppose that we wish to apply Lemma 6 to transform a system of at most n balanced packings, each containing at most n circles, so that some two tangent circles in each packing are placed at the positions of C_0 and C_1 . Then there exists a sequence of marks on circle C_0 such that, if each packing is placed with one of its points of contact on one of these marks, then no two of the transformed packings intersect each other, and the transformed size of the circles in each packing is smaller than the size of C_0 and C_1 by at most a polynomial factor.*




Fig. 13. An inverted view of the construction in Lemma 7, for $n = 3$

Proof: Place a line tangent to the two circles and place the first mark at the point where it touches C_0 . Then, to place each successive mark, pack a sequence of $2n$ circles into the triangular gap between C_0 , C_1 , and the circle or line defining the previous mark, in such a way that each circle is tangent to C_0 , C_1 , and the previous circle in the sequence. Place the next mark at the point where the last of these circles touches C_0 . These circles are as large as possible for their position in any packing that lies between C_0 and C_1 , so it is not possible for any of the transformed packings to stretch more than half-way from one mark to the next; therefore, no two transformed packings interfere with each other.

To analyze the size of the circles used to generate this system of marks, it is helpful to construct the same sequence of circles in a different way. Place two parallel lines in the plane, a unit distance apart, and pack a sequence of $2n(n-1) + 1$ unit-diameter circles between them, so that each of the circles in this sequence is tangent to both lines and to its neighbors in the sequence. Then, invert this system of lines and circles through a unit-radius circle centered on a boundary point of the first circle in the sequence of unit-diameter circles; see Fig. 13. This first circle will invert to a line, tangent to the circles C_0 and C_1 formed by the inverted images of the two parallel lines, and the remaining circles will invert to a sequence of circles between C_0 and C_1 as constructed above. By choosing the inversion center appropriately, it is possible to choose arbitrarily the ratio of radii of the circles C_0 and C_1 formed by inverting the two parallel lines, so we can construct a system of circles in this way that is similar to the previous construction for any arbitrary pair of circles C_0 and C_1 .

When we invert by a unit circle, the distances of each point from the inversion center become inverted: the closer of the two parallel lines becomes the larger of the two circles C_0 and C_1 , which may be arbitrarily large. Because the other one of the two parallel lines has all points at distance at least $1/2$ from the inversion center, some points within distance 1 of the inversion center, and some points arbitrarily far from the inversion center, it becomes transformed into a circle that passes from the inversion center to some points outside the inversion circle, but that remains within a radius-2 disk centered at the inversion center. Thus, its radius is necessarily between $1/2$ and 1. The i th of the $2n(n-1) + 1$ circles in the sequence of circles has its nearest point to the inversion center at a distance between $i - 1/2$ and i , and its farthest point to the inversion center at a distance between $i + 1/2$ and $i + 1$, so when these distances are inverted its transformed diameter becomes $\Theta(1/i^2)$. Therefore, the smallest of the $2n(n-1)$ circles used to generate this system of marks is only polynomially small: its radius is $\Theta(1/n^4)$ times the smaller of the radii of C_0 and C_1 . \square

4.4 Adjacent separating pairs

As a warm-up to our main result on bounded tree-depth, we show here the existence of balanced circle-contact representations for a special subclass of the planar graphs of bounded tree-depth, the graphs G in which the separation pairs are all adjacent. That is, if we construct the block-cut tree of G , and for each block construct the SPQR tree, then each separation pair of each SPQR tree should be the endpoints of an edge somewhere in G . Equivalently, among each two adjacent nodes in any of the SPQR trees for the blocks of G , one of the two nodes should be a P-node, and one of the edges of this P-node should be a real edge.

Lemma 8. *For every constant bound d , every planar graph G with tree-depth at most d in which all separation pairs are adjacent has a balanced circle-contact representation.*

Proof: We form the block-cut tree of G , form the SPQR tree of each block, and form a circle-contact representation for the (constant size) graph associated with each SPQR tree node. For the R nodes, we use the circle packing theorem; for the S nodes, we pack equal-size circles in a circular layout; and for the P nodes, we place two equal circles in contact with each other (representing the one real edge of the P node). We must then glue these packings together so that, when two or more nodes contain the same separation pair, this pair is represented by the same two circles; because of our assumption about the adjacency of separation pairs, we can assume that these circles are tangent to each other.

To perform this gluing step, at one of the P-nodes of the SPQR tree, we consider all of the neighboring nodes connected to it (corresponding to the virtual edges of the P node). For each neighboring node X , we form a configuration consisting of the two circles representing the separation vertices in its packing, and a marked point where another circle is tangent to these two circles. For the P-node itself, we have the corresponding two circles, on which we may place a mark at an arbitrary point. By Lemma 6, we may transform the packing representing X (and anything else that has already been glued to it) so that its separation pair is represented by the same two circles as the P node and its marked point of contact with these two circles is at the given point. By placing the marked point on the P-node sufficiently close to its two circles' point of tangency, the packing representing X will also be transformed to be close to the point of tangency, allowing it to avoid overlapping with any other part of the circle packing that has already been glued to the P-node.

In this way, the packings for each node of an SPQR tree can be glued together, step by step, to form a single packing representing the whole SPQR tree. At each P-node, in order to keep the packings glued to that node from interfering with each other, it is sufficient to place the marks polynomially close to each other and to the point of tangency, causing the glued-in packings to be reduced in size (compared to the two circles representing the separation vertices) by a polynomial factor. Because the SPQR tree has bounded height, at most a constant number of these polynomial factors are compounded together, resulting in polynomial balance overall.

The process for gluing together blocks of the block-cut tree is very similar, but simpler, because each gluing step must match only one circle (the articulation vertex between two blocks) rather than two circles. The analysis is the same. \square

4.5 Inversive-distance circle packings

To apply the same method to graphs with non-adjacent separation pairs, we would like to modify the graphs associated with each SPQR tree node by removing the virtual edges between each non-adjacent separation pair, before finding a circle-contact representation for each of these graphs and then gluing these representations together by Möbius transformations. In the case of adjacent separation pairs, this gluing used Lemma 6, in which a one-parameter family of Möbius transformations allows any tangent pair of circles to be mapped to any other tangent pair of circles, with an additional degree of freedom that allows two marked points to be aligned with each other. However, for non-adjacent separation pairs, the problem is made more complicated by the fact that pairs of disjoint circles cannot always be mapped into each other by a Möbius transformation. Instead, an invariant of these pairs called the *inversive distance* controls the existence of such a mapping. Two Euclidean circles with radii r and R , whose centers are at distance ℓ from each other, have inversive distance

$$I = \frac{\ell^2 - r^2 - R^2}{2rR}.$$

This is an invariant under Möbius transformations, and can also be defined for circles in hyperbolic and spherical geometry [7]. The inversive distance of two tangent circles is one; it is less than one for circles that cross each other and greater than one for disjoint circles. Two pairs of circles that both have equal inversive distances can be mapped into each other in the same way as Lemma 6 (with the freedom to align two marked points) but two pairs with unequal inversive distances cannot be mapped into each other.

Bowers and Stephenson [8] suggested the use of inversive distance in circle packings. An inversive-distance circle packing is specified by a planar graph G and a label $I(uv)$ for each edge uv of G ; the goal is to represent the vertices of G by circles such that, for each edge uv , the circles representing u and v are at inversive distance exactly $I(uv)$ from each other. Bowers and Stephenson write that “the theoretical underpinnings are not yet in place” for this type of packing. In particular it is still not known under what conditions an inversive-distance circle packing exists, although a numerical method briefly suggested by Collins and Stephenson [13] and detailed by Bowers and Hurdal [7] is reported to work well in practice for finding such packings. Notwithstanding this gap in our theoretical knowledge, we will prove that, in the case of interest for us, inversive-distance circle packings do exist and can be used to find balanced packings of graphs of bounded tree-depth.

One thing that is known about inversive-distance circle packings is that, when they do exist for a given maximal planar graph, they are unique.

Lemma 9 (Luo [30]). *Let G be a maximal planar graph, with outer face Δ , let G have an assignment I of inversive distances to its edges, and fix a non-collinear placement of the three vertices of Δ in the Euclidean plane. If an inversive-distance circle packing exists for G and I , with the circles representing Δ centered at the fixed placement of the vertices of Δ , then this packing is uniquely determined by G , I , and the placement of Δ .*

The fixed placement of Δ simplifies the statement of this result by eliminating the possibility of Möbius transformations that would change the position of the packing without changing its inversive distances. Luo actually proved this result in greater generality, for triangulated topological surfaces with Euclidean or hyperbolic geometry; the analogous statement for spherical geometry turns out not to be true [31]. In Luo’s more general setting, the packing is determined (up to scale) by the system of angular defects at each vertex of the surface. In the version that we need for the Euclidean plane, the angular defects are zero except at the vertices of Δ , and the fixed placement of Δ determines both the angular defects at these vertices and the scale of the packing, so the result as stated above follows from Luo’s more general theorem.

Corollary 2. *Let G be a maximal planar graph, with outer face Δ , and fix a non-collinear placement of the three vertices of Δ in the Euclidean plane. Let m be the number of edges in G . Then the set of assignments of distances to the edges of G that have inversive-distance circle packings is an open subset of \mathbb{R}^m , and there exists a continuous function from these distance assignments to the circle centers and radii of the (unique) corresponding packings.*

Proof: The space of distance assignments is m -dimensional, as is (by Euler’s formula) the space of possible placements of circle centers (obeying the fixed placement of Δ) and radii in a packing. Thus, by Lemma 9, the function from circle placements to inversive distances is a continuous injective function between two sets of the same dimension. By Brouwer’s theorem on invariance of domain [11], this function has an open image and a continuous inverse function. \square

By using the fact that the set of allowable inversive-distance assignments is open, we can find the packing we need:

Lemma 10. *Let $G = (V, E)$ be a planar graph, and let $F \subset E$ be a subset of “virtual” edges in G . Then, for all sufficiently small $\varepsilon > 0$, there exists a collection of circles in the plane, and a one-to-one correspondence between these circles and the vertices of G , such that:*

- For each edge $e \in E \setminus F$, the two circles representing the endpoints of e are tangent (inversive distance 1), and
- for each edge $e \in F$, the two circles representing the endpoints of e have inversive distance exactly $1 + \varepsilon$.

Additionally, this circle packing can be chosen to vary continuously with ε , in such a way that as $\varepsilon \rightarrow 0$ the packing converges to a circle-contact representation for G .

Proof: Add extra vertices to G , and extra edges incident to those vertices, to make the augmented graph maximal planar; include as part of this augmentation a triangle Δ of added vertices. Choose a placement of Δ at the vertices of an equilateral triangle in the plane. Let D be the distance assignment that sets all inversive distances on edges of the augmented graph to equal 1; then there exists a circle packing for D , by the original circle packing theorem of Koebe. By Corollary 2, there exists an open neighborhood of D within which all distance assignments are achievable; let ε be sufficiently small that all vectors within L_∞ distance ε of D are achievable. In particular, the distance assignment that sets all distances in F equal to $1 + \varepsilon$, and all remaining distances equal to 1, has an inversive circle packing. The subset of this packing that corresponds to the original vertices in V then satisfies the conditions of the lemma. Once the augmentation of G and the placement of Δ is fixed, the continuous variation of the packing on ε follows from the continuity of the function from distance assignments to packings stated in Corollary 2. \square

4.6 Circle packings for graphs of bounded tree-depth

Theorem 5. *For every constant bound d , every planar graph with tree-depth at most d has a balanced circle-contact representation.*

Proof: We follow the same outline as the proof of Lemma 8, except that we choose a sufficiently small value of ε (the same ε for all nodes of all SPQR trees) and use Lemma 10 to construct circle packings for each node in which the real edges of the graph associated with the node, and the virtual edges corresponding to adjacent separation pairs, have inversive distance one (that is, their circles are tangent) while the remaining edges in each node (non-adjacent separation pairs) have inversive distance exactly $1 + \varepsilon$ (their circles are disjoint).

The analogue of Lemma 6 holds for any two pairs of tangent circles with the same inversive distance as each other. By continuity of the packing with respect to ε , for small enough values of ε , the analogues of Lemma 1 and Lemma 7 also remain valid. Additionally, although inversive-distance circle packings in general are not guaranteed to have disjoint circles for pairs of vertices that are not connected by an edge in the underlying graph, this property also follows for our circle packings, for sufficiently small ε , by continuity. With these ingredients in hand, the proof proceeds as before. \square

We remark that the proof described above does not require ε to be significantly smaller than the balance of the resulting packing. Indeed, we only need to apply the inversive circle packing method of Lemma 10 to the R-nodes of SPQR trees, as the other nodes have associated graphs that are easy to represent directly. Because each R-node has $\mathcal{O}(1)$ vertices, there are only $\mathcal{O}(1)$ combinatorially distinct choices of an associated graph and set of virtual edges for which we need packings, and in order to apply Lemma 10 we can choose any ε that is small enough to work for each of these inputs, a value that is independent of n . The limiting factor controlling our actual choice of ε is the variant of Lemma 7 that applies to inversive-distance circle packings, which only requires ε to be polynomially small.

Although our proof is not constructive, the report by Collins and Stephenson [13] that their inversive-distance circle packing algorithm works well in practice gives us hope that the same would be true of its application to this problem.

5 Conclusion

We studied balanced circle packings for planar graphs, showing that several rich classes of graphs have balanced circle packings, while several simple classes of graphs do not. One interesting open problem is whether or not every outerplanar graph has a balanced circle packing representation. While we identified several subclasses of outerplanar graphs that admit such representations, the question remains open for general outerplanar graphs.

Acknowledgments. This work is supported in part by the National Science Foundation under grants CCF-1228639, CCF-1115971, DEB 1053573, and by the Office of Naval Research under Grant No. N00014-08-1-1015.

References

1. O. Aichholzer, G. Rote, A. Schulz, and B. Vogtenhuber. Pointed drawings of planar graphs. *Comput. Geom. Th. Appl.* 45(9):482–494, 2012, doi:10.1016/j.comgeo.2010.08.001.
2. M. Badent, C. Binucci, E. D. Giacomo, W. Didimo, S. Felsner, F. Giordano, J. Kratochvíl, P. Palladino, M. Patrignani, and F. Trotta. Homothetic triangle contact representations of planar graphs. *Canadian Conference on Computational Geometry (CCCG’07)*, pp. 233–236, 2007.
3. M. J. Bannister, S. Cabello, and D. Eppstein. Parameterized complexity of 1-planarity. *13th Int. Symp. Algorithms and Data Structures (WADS 2013)*, pp. 97–108. Springer, Lect. Notes Comput. Sc. 8037, 2013, doi:10.1007/978-3-642-40104-6_9.
4. M. J. Bannister, W. E. Devanny, D. Eppstein, and M. T. Goodrich. The Galois complexity of graph drawing: Why numerical solutions are ubiquitous for force-directed, spectral, and circle packing drawings. *22nd Int. Symp. Graph Drawing (GD 2014)*. Springer, Lect. Notes Comput. Sc., 2014. to appear.
5. M. A. Bekos and C. N. Raftopoulou. Circle-representations of simple 4-regular planar graphs. *20th Int. Symp. Graph Drawing (GD 2012)*, pp. 138–149. Springer, Lect. Notes Comput. Sc. 7704, 2013, doi:10.1007/978-3-642-36763-2_13.
6. M. Bern and D. Eppstein. Optimal Möbius transformations for information visualization and meshing. *7th Int. Worksh. Algorithms and Data Structures (WADS 2001)*, pp. 14–25. Springer, Lect. Notes Comput. Sc. 2125, 2001, doi:10.1007/3-540-44634-6_3, arXiv:cs.CG/0101006.
7. P. L. Bowers and M. K. Hurdal. Planar conformal mappings of piecewise flat surfaces. *Visualization and Mathematics III*, pp. 3–34. Springer, Mathematics and Visualization, 2003, doi:10.1007/978-3-662-05105-4_1.
8. P. L. Bowers and K. Stephenson. 8.2 Inversive distance packings. *Uniformizing dessins and Belyi maps via circle packing*, pp. 78–82, *Memoirs of the American Mathematical Society* 805, 2004, doi:10.1090/memo/0805.
9. H. Breu and D. G. Kirkpatrick. Unit disk graph recognition is NP-hard. *Comput. Geom. Th. Appl.* 9(1–2):3–24, 1998, doi:10.1016/S0925-7721(97)00014-X.
10. G. Brightwell and E. Scheinerman. Representations of planar graphs. *SIAM J. Discrete Math.* 6(2):214–229, 1993, doi:10.1137/0406017.
11. L. E. J. Brouwer. Beweis der Invarianz des n -dimensionalen Gebiets. *Mathematische Annalen* 71:305–315, 1912.
12. G. Chen and X. Yu. Long cycles in 3-connected graphs. *J. Comb. Theory B* 86(1):80–99, 2002, doi:10.1006/jctb.2002.2113.
13. C. R. Collins and K. Stephenson. A circle packing algorithm. *Comput. Geom. Th. Appl.* 25(3):233–256, 2003, doi:10.1016/S0925-7721(02)00099-8.
14. R. Descartes. *Oeuvres de Descartes. Correspondance IV*, 1901.
15. G. Di Battista and R. Tamassia. On-line graph algorithms with SPQR-trees. *17th Int. Colloq. Automata, Languages and Programming (ICALP 1990)*, pp. 598–611. Springer, Lect. Notes Comput. Sc. 443, 1990, doi:10.1007/BFb0032061.
16. D. Dolev, T. Leighton, and H. Trickey. Planar embedding of planar graphs. *Advances in Computing Research* 2:147–161, 1984.
17. C. A. Duncan, E. R. Gansner, Y. F. Hu, M. Kaufmann, and S. G. Kobourov. Optimal polygonal representation of planar graphs. *Algorithmica* 63(3):672–691, 2012, doi:10.1007/s00453-011-9525-2.
18. D. Eppstein. Planar Lombardi drawings for subcubic graphs. *20th Int. Symp. Graph Drawing (GD 2012)*, pp. 126–137. Springer, Lect. Notes Comput. Sc. 7704, 2013, doi:10.1007/978-3-642-36763-2_12.
19. D. Eppstein, D. Holten, M. Löffler, M. Nöllenburg, B. Speckmann, and K. Verbeek. Strict confluent drawing. *21st Int. Symp. Graph Drawing (GD 2013)*, pp. 352–363. Springer, Lect. Notes Comput. Sc. 8242, 2013, doi:10.1007/978-3-319-03841-4_31, arXiv:1308.6824.
20. H. de Fraysseix, P. O. de Mendez, and P. Rosenstiehl. On triangle contact graphs. *Combinatorics, Probability & Computing* 3(2):233–246, 1994, doi:10.1017/S0963548300001139.
21. E. N. Gilbert and E. F. Moore. Variable-length binary encodings. *Bell System Technical Journal* 38(4):933–967, 1959, doi:10.1002/j.1538-7305.1959.tb01583.x.

22. D. Gonçalves, B. Lévêque, and A. Pinlou. Triangle contact representations and duality. *Discrete Comput. Geom.* 48(1):239–254, 2012, doi:10.1007/s00454-012-9400-1.
23. P. Hliněný. Contact graphs of curves. *3rd Int. Symp. Graph Drawing (GD 1995)*, pp. 312–323. Springer, Lect. Notes Comput. Sc. 1027, 1996, doi:10.1007/BFb0021814.
24. P. Hliněný. Classes and recognition of curve contact graphs. *J. Comb. Theory B* 74(1):87–103, 1998, doi:10.1006/jctb.1998.1846.
25. P. Hliněný and J. Kratochvíl. Representing graphs by disks and balls (a survey of recognition-complexity results). *Discrete Math.* 229(1–3):101–124, 2001, doi:10.1016/S0012-365X(00)00204-1.
26. J. Hopcroft and R. E. Tarjan. Dividing a graph into triconnected components. *SIAM J. Comput.* 2(3):135–158, 1973, doi:10.1137/0202012.
27. B. Keszegh, J. Pach, and D. Pálvölgyi. Drawing planar graphs of bounded degree with few slopes. *SIAM J. Discrete Math.* 27(2):1171–1183, 2013, doi:10.1137/100815001.
28. S. Khuller and R. Thurimella. Approximation algorithms for graph augmentation. *J. Algorithms* 14(2):214–225, 1993, doi:10.1006/jagm.1993.1010.
29. P. Koebe. Kontaktprobleme der konformen Abbildung. *Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl.* 88:141–164, 1936.
30. F. Luo. Rigidity of polyhedral surfaces, III. *Geometry & Topology* 15(4):2299–2319, 2011, doi:10.2140/gt.2011.15.2299.
31. J. Ma and J.-M. Schlenker. Non-rigidity of spherical inversive distance circle packings. *Discrete Comput. Geom.* 47(3):610–617, 2012, doi:10.1007/s00454-012-9399-3.
32. S. Mac Lane. A structural characterization of planar combinatorial graphs. *Duke Mathematical Journal* 3(3):460–472, 1937, doi:10.1215/S0012-7094-37-00336-3.
33. S. M. Malitz and A. Papakostas. On the angular resolution of planar graphs. *SIAM J. Discrete Math.* 7(2):172–183, 1994, doi:10.1137/S0895480193242931.
34. C. McDiarmid and T. Müller. Integer realizations of disk and segment graphs. *J. Comb. Theory B* 103(1):114–143, 2013, doi:10.1016/j.jctb.2012.09.004.
35. B. Mohar. A polynomial time circle packing algorithm. *Discrete Math.* 117(1–3):257–263, 1993, doi:10.1016/0012-365X(93)90340-Y.
36. B. Mohar. Drawing graphs in the hyperbolic plane. *7th Int. Symp. Graph Drawing (GD 1999)*, pp. 127–136. Springer, Lect. Notes Comput. Sc. 1731, 1999, doi:10.1007/3-540-46648-7_13.
37. J. Nešetřil and P. Ossona de Mendez. *Sparsity: Graphs, Structures, and Algorithms*. Springer, 2012, pp. 115–144, doi:10.1007/978-3-642-27875-4.
38. J. Nievergelt and E. M. Reingold. Binary search trees of bounded balance. *SIAM J. Comput.* 2:33–43, 1973, doi:10.1137/0202005.
39. G. Rote. Realizing planar graphs as convex polytopes. *19th Int. Symp. Graph Drawing (GD 2011)*, pp. 238–241. Springer, Lect. Notes Comput. Sc. 7034, 2012, doi:10.1007/978-3-642-25878-7_23.
40. H. Sachs. Coin graphs, polyhedra, and conformal mapping. *Discrete Math.* 134(1–3):133–138, 1994, doi:10.1016/0012-365X(93)E0068-F.