A Möbius-Invariant Power Diagram and Its Applications to Soap Bubbles and Planar Lombardi Drawing

David Eppstein

Received: 10 September 2013 / Revised: 1 March 2014 / Accepted: 8 July 2014 /
Published online: 5 September 2014
© Springer Science+Business Media New York 2014

Abstract We use three-dimensional hyperbolic geometry to define a form of power diagram for systems of circles in the plane that is invariant under Möbius transformations. By applying this construction to circle packings derived from the Koebe–Andreiev–Thurston circle packing theorem, we show that every planar graph of maximum degree three has a planar Lombardi drawing (a drawing in which the edges are drawn as circular arcs, meeting at equal angles at each vertex). We use circle packing to construct planar Lombardi drawings of a special class of 4-regular planar graphs, the medial graphs of polyhedral graphs, and we show that not every 4-regular planar graph has a planar Lombardi drawing. We also use these power diagrams to characterize the graphs formed by two-dimensional soap bubble clusters (in equilibrium configurations) as being exactly the 3-regular bridgeless planar multi-graphs, and we show that soap bubble clusters in stable equilibria must in addition be 3-connected.

Keywords Graph drawing · Lombardi drawing · Soap bubbles · Circle packing · Power diagram · Hyperbolic geometry

1 Introduction

Voronoi diagrams, the partitions of the Euclidean plane or other geometric spaces into cells within which one of a given set of point sites is closer than all others, have been fundamental to many algorithms in computational geometry. These diagrams may also be described by their planar dual Delaunay triangulations, in which two sites are connected by an edge if they lie on the boundary of a disk that is empty of sites. One of

D. Eppstein
Department of Computer Science, University of California, Irvine, CA, USA
e-mail: eppstein@uci.edu

the useful properties of Delaunay triangulations (viewed as abstract graphs rather than geometric objects and extended to include pairs of sites on the boundaries of empty disk complements) is that they are invariant under Möbius transformations of the plane (a set of circle-preserving transformations that includes circle inversions as well as the more usual translations, rotations, and dilations, and is closed under composition of transformations). This Möbius invariance property of Delaunay triangulations has been used, for instance, to define Möbius-invariant interpolation methods analogous to the Voronoi-based natural neighbor method [5].

Power diagrams [3] generalize Voronoi diagrams to sites that are circles rather than points, using the *power distance* from a circle, the length of a tangent line segment from a given point to the circle. Like Voronoi diagrams, power diagrams have simply-connected polygonal cells and linear total complexity. However, classical power diagrams and their duals do not have the same Möbius invariance property as Delaunay triangulations. In this paper we define an alternative power diagram for a system of circles that is Möbius invariant. Our diagram may be interpreted in terms of three-dimensional hyperbolic geometry, or as the minimization diagram of a distance function from points to circles; this distance can be defined geometrically in terms of the radii of certain tangent circles associated with the point and circle, or by an algebraic formula of their coordinates. The cell boundaries of this diagram are circular arcs. The diagram for an arbitrary system of circles may not have simply connected cells, but for circle packings (systems of tangent circles dual to maximal planar graphs) the cells are simply connected and the diagram may be constructed in linear time. We provide two applications of this construction, in graph drawing and in the mathematics of soap bubbles.

1.1 Lombardi Drawing

Lombardi drawing is a style of graph drawing, named after artist Mark Lombardi, in which the edges of a graph are drawn as circular arcs and in which every vertex is surrounded by edges that meet at equal angles at the vertex—that is, the drawing has *perfect angular resolution* [17, 18]. Several families of graphs are known to have Lombardi drawings, including regular graphs (under certain restrictions on their factorizations into regular subgraphs) and certain highly symmetric graphs [18]. Lombardi drawings have also been used to draw plane trees with perfect angular resolution in polynomial area, something that would be impossible for straight line drawings [17]. All graphs have drawings that relax the constraints of Lombardi drawing to allow pol-yarcs or unequal angles [9, 16], but the improved aesthetic quality of true Lombardi drawings makes it of interest to determine more precisely which graphs have such drawings.

Planarity, the avoidance of crossing edges, is of great importance both in Lombardi drawing and in graph drawing more generally. By Fáry’s theorem, every planar graph can be drawn planarly with straight line segments for its edges, and therefore it can also be drawn with circular arcs. Indeed, there exist *universal sets* of n points such that every n -vertex planar graph can be drawn using those points for vertices and circular arcs for edges, something that is not true for straight-line drawing [2]. However, those arcs

may not necessarily meet at equal angles. Not all planar graphs have planar Lombardi drawings [18] and indeed even the planar 3-trees (which are maximal planar and 3-connected) can fail to have planar Lombardi drawings [16]. Only a few positive results on planar Lombardi drawing are known: prior to our work, only trees, Halin graphs (the graphs formed from plane trees by adding a cycle connecting the leaves), outerpaths (graphs for which the adjacency structure of the bounded faces is a path), and the graphs of symmetric polyhedra were proven to have planar Lombardi drawings [18,40].

In this paper we take a major step forward in our knowledge of planar Lombardi drawings, and in the applicability of the Lombardi drawing style, by showing that all planar graphs of maximum degree three have planar Lombardi drawings. The heart of our method applies to 3-connected 3-regular planar graphs: we apply the Koebe–Andreiev–Thurston circle packing theorem to the dual graph, and then construct the Möbius-invariant power diagram of the circle packing, which we show to be a planar Lombardi drawing of the input. Our implementation of this method produces drawings in which the overall drawing and the individual faces are all approximately circular and in which the spacing of the vertices is locally uniform. We extend these results to graphs that are neither 3-regular nor 3-connected by using bridge-block trees and SPQR trees to decompose the graph into 3-connected subgraphs, drawing these subgraphs separately, and using Möbius transformations to glue them together into a single drawing. The same method may be applied to certain other non-3-regular graphs, including the Halin graphs. We also use a different type of circle packing (with two sets of orthogonal circles) to construct planar Lombardi drawings of a special class of 4-regular planar graphs, the medial graphs of polyhedral graphs. However, as we show, not every 4-regular planar graph has a planar Lombardi drawing.

1.2 Soap Bubbles

When soap bubbles are blown between closely spaced parallel glass plates, they form foams or clusters of bubbles in which the bubble walls are perpendicular to the sheets (Fig. 1). These clusters may be described by undirected graphs, with an edge for each bubble wall and a vertex for each junction where multiple bubbles meet. This physical phenomenon naturally raises combinatorial and algorithmic questions: which graphs can be formed in this way? And, if we are given as input a graph of the proper type, how can we construct as output a soap bubble cluster that could exist physically and that represents the given graph?

Analogous problems of characterizing graphs that represent other kinds of discrete geometric structure have long been studied. The prototypical example is Steinitz’s theorem, that the graphs of three-dimensional convex polyhedra are exactly the 3-vertex-connected planar graphs [51]; much subsequent research has focused on the algorithmic problem of constructing polyhedral representations with small coordinates [10,48,49]. The graphs of nonconvex polyhedra with axis-parallel edges may also be characterized in terms of planarity and connectivity [22], as may the graphs of subdivisions of rectangles into smaller rectangles [38,39,57].

In this paper we model soap bubble clusters mathematically using two long-established principles governing their behavior: Plateau’s laws describing the local

Fig. 1 Soap bubbles between two glass plates. Public domain negative image by Klaus-Dieter Keller from Wikimedia commons

geometry of the surfaces and junctions in a bubble cluster, and the Young–Laplace equation relating the curvature of bubble surfaces to the pressure in each bubble. We define a *planar soap bubble cluster* to be a family of arcs in the plane that obeys the planar versions of these principles. It turns out that soap bubbles, defined in this way, are Möbius invariant: Möbius transformation of a planar soap bubble cluster results in another physically realizable soap bubble cluster, despite the unphysicality of this transformation. This fact has been previously used to study planar triple bubbles [60], but as we show it has broader implications. Using Möbius invariance, we provide a simple proof that the graph of every planar soap bubble cluster is 2-vertex-connected, 3-regular, and planar. These necessary conditions turn out to be sufficient: every 2-vertex-connected 3-regular planar graph is the graph of a soap bubble cluster. To prove this, we use our Lombardi drawing algorithm: not every Lombardi drawing is a soap bubble cluster, because the drawing may not satisfy the Young–Laplace equation, but the drawings produced by our algorithm in the 2-vertex-connected case do satisfy this equation and form valid soap bubbles.

We are not aware of prior attempts to characterize the graphs of soap bubbles, but soap bubbles have been studied mathematically in many other ways. The Kelvin conjecture, stating that a foam with truncated-octahedron cells has the minimum total surface area among all soap bubble foams with equal-volume cells, was posed in the 19th century by Lord Kelvin and finally solved negatively in 1993 by the discovery of the Weaire–Phelan structure, a more parsimonious foam with two different cell shapes [55,59]; proof that this newer structure is optimal remains elusive. A related problem, the minimum-area cap for the ends of hexagonal honeycomb cells, also remains open [27]. It is conjectured that the minimum-surface-area four-bubble cluster, for given bubble volumes, is given by the stereographic projection of a four-dimensional simplex, but this has not been proven, and there are six-bubble clusters for which the numerically-computed surfaces separating the bubbles are not spheri-

cal [53]. Even as simple-sounding a problem as the double bubble conjecture (proving that the optimal structure for clusters of two bubbles with two fixed volumes has three spherical patches sharing a common circle) took many years to solve [30,34].

Two-dimensional soap films have also seen much prior research, both as a test case for three-dimensional theories and in the context of the “soap bubble computer”. In this device, pegs are placed between two glass plates at specified locations; a soap film that connects the pegs forms a minimal (though generally not minimum) Steiner network, providing a heuristic and approximate physical solution to an NP-complete problem [1,6,19,32,35]. Another result on two-dimensional soap films provides a planar analogue of the Kelvin conjecture: the minimum-length enclosure for an infinite set of bubbles of equal areas is a hexagonal tiling of the plane [27,28], and similar tilings arise also as the minimum-length enclosure for finite clusters of equal bubbles [12].

John Sullivan has suggested that “we might look for foams as relaxations of Voronoi decompositions” [53]. Our use of power diagrams to construct bubble clusters may be seen as a validation of Sullivan’s suggestion.

2 Preliminaries

2.1 Möbius Transformations

Let $\mathbb{S}^2$ denote the space formed by adding a single point ∞ “at infinity” to the Euclidean plane; this space is also known as the one-dimensional complex projective line $\mathbb{P}^1(\mathbb{C})$. In $\mathbb{S}^2$, straight lines may be interpreted as limiting cases of circles, with infinite radius and containing the point ∞ . A *Möbius transformation* [50] is a map from $\mathbb{S}^2$ to itself that transforms every circle (or line) into another circle (or line). Using complex-number coordinates, these transformations may be represented as the *fractional linear transformations*

$$z \mapsto \frac{az + b}{cz + d},$$

and their conjugates. Here a, b, c , and d are complex, and $ad - bc \neq 0$ to prevent the transformation from defining a constant function. If $c = 0$ then ∞ is mapped to itself; otherwise, $-d/c$ gets mapped to ∞ , and ∞ gets mapped to a/c . Multiplying all four of a, b, c , and d by the same complex number leaves the transformation unchanged, so the set of Möbius transformations has six real degrees of freedom. This number of degrees of freedom is sufficient to allow every triple of distinct points to be mapped by a Möbius transformation to every other triple of distinct points.

Möbius transformations may also be defined from *inversions*. If O is a circle centered at point o with radius r , inversion through O maps any point p to another point q on the ray from o through p , at distance $r^2/d(o, p)$ from o . O is fixed by the inversion, the inside of O becomes the outside and vice versa, and o trades places with ∞ . In the limiting case of a line, inversion becomes reflection across the line. Every Möbius transformation may be represented as a composition of a finite set of inversions.

Möbius transformations are *conformal mappings*: they preserve the angle of every two incident curves. Since they preserve both circularity and angles, they preserve the property of being a Lombardi drawing. Less obviously, we will see later that they also preserve the property of being a two-dimensional soap bubble cluster.

2.2 Hyperbolic Geometry

To understand our algorithms, it will be helpful to review some qualitative features of hyperbolic geometry, avoiding detailed calculations [8].

In the *upper halfspace model* of three-dimensional hyperbolic space (also called the Poincaré halfspace model), hyperbolic space is represented by an open halfspace of three-dimensional Euclidean space, but with a non-Euclidean distance metric. The boundary plane of the Euclidean halfspace does not belong to the hyperbolic space but may be thought of as the set of “points at infinity” for the hyperbolic space. It is convenient to add one more point ∞ to this boundary plane, so that it becomes a copy of $\mathbb{S}^2$. Hyperbolic lines are represented by Euclidean semicircles that meet the boundary plane at right angles, or by Euclidean rays perpendicular to the boundary plane; the vertical rays may be viewed as the limiting case of semicircles with one endpoint at ∞ . Hyperbolic planes are represented by Euclidean hemispheres that meet the boundary plane at right angles in a circle, or by halfplanes that touch the point ∞ and meet the boundary plane perpendicularly in a line.

Hyperbolic space is locally Euclidean: within sufficiently small balls, hyperbolic distances may be approximated arbitrarily well by Euclidean distances. For every two reference frames (a choice of a point within the space, and a system of Cartesian coordinates for the local Euclidean geometry near that point) a unique symmetry of the space takes one reference frame into the other. The symmetries of hyperbolic space may be extended to the plane at infinity, on which they act as Möbius transformations. Every Möbius transformation of $\mathbb{S}^2$ corresponds uniquely to a symmetry of hyperbolic space.

2.3 Circle Packing

The Koebe–Andreiev–Thurston circle packing theorem [52] states that the vertices of every maximal planar graph may be represented by circles with disjoint interiors, such that two vertices are adjacent if and only if the corresponding two circles are tangent. The representation is unique up to Möbius transformations. An alternative version of the theorem, which we use in some of our graph drawing results, states that the vertices of every 3-connected planar graph and of its dual graph may be represented by a combined system of circles in such a way that two circles are tangent if they represent adjacent vertices in the same graph, cross orthogonally if they represent an adjacent vertex-face pair (consisting of one vertex from each graph), and otherwise are disjoint. Again, this representation is unique up to Möbius transformations [7].

It is not known how to find circle packings in strongly polynomial time, but it is possible to approximate both types of packing by numerical algorithms that are polynomial in both the number of vertices and the accuracy of approximation [43]. We

implemented a numerical relaxation procedure detailed by Collins and Stephenson [11] for finding packings of triangulated planar graphs in which the outer face need not be a triangle and the outer circles have prespecified radii, a more general case than the maximal planar case that we need. Their procedure repeatedly performs the following steps:

- It chooses a circle in the packing, and computes the *angular defect* by which its neighboring circles either fail to surround it or surround it by a larger angle than 2π .
- It determines a representative radius for the neighboring circles such that, if they all had the representative radius (in place of their actual radii) they would have the same defect
- It replaces the radius of the central circle by a new value such that, if the surrounding circles all had the representative radius, the angular defect would be zero.

Each of these steps may be performed using simple trigonometric calculations. As Collins and Stephenson show, this method converges rapidly to a unique solution, the system of radii for a valid packing. Once a close approximation to the radii has been calculated, the positions of the circle centers are not difficult to determine. Collins and Stephenson observe that, with small changes, the same procedure can also be used to find orthogonal primal–dual circle packings, or more generally packings with specified overlap angles between adjacent circles.

The circle packings constructed by the Collins–Stephenson procedure (with equal outer radii) are unsatisfactory for some of our graph drawing applications, because the resulting drawings will have one vertex placed at ∞ . To improve our drawings, we find a Möbius transformation of the packing for which one chosen circle surrounds all the others. Among all transformations fixing the outer circle we choose the one that maximizes the radius of the smallest inner circle (Fig. 2). This problem of finding a radius-maximizing Möbius transformation may be expressed (using the connection between Möbius transformations and three-dimensional hyperbolic geometry) as a *quasiconvex program*, a problem of finding the minimum value of a pointwise maximum of quasiconvex functions. It may be solved either combinatorially in linear time using LP-type optimization algorithms or numerically using local improvement procedures; the theory of quasiconvex programs guarantees that there are no local optima in which the local improvement might get stuck [4, 20]. Since we are already using a numerical method to find circle packings, our implementation also takes the numerical approach to find the best transformation.

2.4 Triangle Centers

A *triangle center* is a function from Euclidean triangles to Euclidean points that is *equivariant* under similarities of the Euclidean plane, in the sense that performing a similarity transformation and then constructing the center gives the same result as constructing the center and then performing the similarity transformation. Hundreds of triangle centers are known, and they include many well-known points determined from a triangle, such as its centroid, circumcenter, incenter, and orthocenter [37]. It is convenient, in computing the position of a triangle center, to use *barycentric*

Fig. 2 *Left:* the central region of a circle packing constructed by the Collins–Stephenson procedure; the outer three circles, shown only partially in the figure, have equal radii. *Right:* a transformed packing with one circle exterior to the others, maximizing the minimum circle radius

coordinates, weights for which the center is the weighted average of the vertices. Alternatively, triangle centers may be specified using *trilinear coordinates*, triples of numbers proportional to the distances from each of the triangle’s sides. If the trilinear coordinates are $\alpha : \beta : \gamma$ then the barycentric coordinates are $\alpha a : \beta b : \gamma c$ where the factors a , b , and c are the lengths of the triangle sides opposite the vertex whose weight is α , β , or γ respectively.

There are two triangle centers that (as an unordered pair of points) are equivariant under Möbius transformations, and not just under Euclidean similarities. One of these two centers, the first *isodynamic point*, may be constructed by transforming the given triangle to an equilateral triangle (in such a way that the interiors of the circumcircles of the triangles map to each other), choosing the centroid of the equilateral triangle, and reversing the transformation. The second isodynamic point may be constructed similarly using ∞ in place of the centroid. Alternatively, the first isodynamic point may be calculated from its trilinear coordinates, which are

$$\sin(A + \pi/3) : \sin(B + \pi/3) : \sin(C + \pi/3),$$

where A , B , and C are the three angles of the given triangle. The same formula with $-\pi/3$ in place of $+\pi/3$ gives the trilinears of the second isodynamic point.

3 Power Diagram

In this section we use three-dimensional hyperbolic geometry to construct, from a given set $\mathcal{D}$ of disks (or complements of disks) in the plane, a diagram $M(\mathcal{D})$ that partitions the plane into regions surrounding each disk, separated from each other by circular arcs. This diagram will be equivariant under Möbius transformations in the sense that, for every Möbius transformation μ , $\mu(M(\mathcal{D})) = M(\mu(\mathcal{D}))$.

3.1 Construction

To construct $M(\mathcal{D})$, let Π be the plane containing $\mathcal{D}$, embed Π into a three-dimensional Euclidean space, let Π^+ be one of the two Euclidean halfspaces having Π as its boundary, and reinterpret Π^+ as a Poincaré halfspace model of three-dimensional hyperbolic space. Then, the circle bounding each disk $D_i \in \mathcal{D}$ models the set of limit points of a hyperbolic plane P_i in this space, and the disk itself is the set of limit points of a hyperbolic halfspace H_i . Define the *signed distance* between a hyperbolic point q and H_i , for $q \notin H_i$, to be the positive distance from q to P_i , and define the signed distance for $q \in H_i$ to be the negation of the distance from q to P_i .

The *minimization diagram* of this signed distance function is a partition of the hyperbolic space into cells C_i such that, within cell C_i , the signed distance from any point of the cell to H_i is smaller than the signed distance to every other halfspace H_j . This diagram is directly defined only for points within the hyperbolic space—the signed distance is undefined or infinite for points of the boundary plane Π . However, the minimization diagram may be extended to Π by letting the extended cell $\hat{C}_i$ be the topological closure of C_i in the halfspace model. We define $M(\mathcal{D})$ to be the restriction of this extended minimization diagram to the plane Π containing the given disks. That is, it is the partition of Π determined by the cells $\hat{C}_i \cap \Pi$.

3.2 Properties

Lemma 1 *For every Möbius transformation μ , $\mu(M(\mathcal{D})) = M(\mu(\mathcal{D}))$.*

Proof Because it is defined using only notions intrinsic to hyperbolic geometry (half-spaces and distance), the minimization diagram from which $M(\mathcal{D})$ is defined is invariant under hyperbolic isometries. The result follows from the correspondence between hyperbolic isometries and Möbius transformations: each hyperbolic isometry restricts to a Möbius transformation of the boundary plane of the halfspace model of hyperbolic geometry, and every Möbius transformation comes from a hyperbolic isometry in this way. $\square$

Lemma 2 *The boundaries of the cells in $M(\mathcal{D})$ are circular arcs or subsets of straight lines.*

Proof Let D_i and D_j be any two members of $\mathcal{D}$. First, suppose that D_i does not contain D_j and D_j does not contain D_i . In this case, we may choose an appropriate Möbius transformation μ such that $\mu(D_i)$ and $\mu(D_j)$ are both disks (rather than complements of disks) with equal radii. By symmetry, $M(\{\mu(D_i), \mu(D_j)\})$ must have two halfplane cells, separated from each other by the line that perpendicularly bisects the centers of $\mu(D_i)$ and $\mu(D_j)$.

Alternatively, suppose that one of the two sets D_i and D_j contains the other; without loss of generality, D_i contains D_j . In this case, we may choose an appropriate Möbius transformation μ such that $\mu(D_i)$ and $\mu(D_j)$ are both disks (rather than complements of disks) that are concentric. For any point q of hyperbolic space, the signed distance to H_j will always be smaller than the signed distance to H_i , so in this case $M(\{\mu(D_i), \mu(D_j)\})$ consists of a single cell with no boundary.

Fig. 3 Six disks and one disk complement (shaded) forming a set $\mathcal{D}$ whose diagram $M(\mathcal{D})$ (the thick curves) contains a disconnected region. The inner and outer regions of $M(\mathcal{D})$ both belong to a single cell, for the outer disk complement

This case analysis, together with Lemma 1 and the circle-preserving properties of Möbius transformations, shows that whenever two members of $\mathcal{D}$ determine adjacent cells of $M(\mathcal{D})$, the boundary separating these cells is a circular arc or a subset of a straight line, as the lemma states. $\square$

We remark that in the case of two disks that both have nonempty cells (i.e. the disks cross, are disjoint, or have disjoint complements) the boundary between the two cells lies on a *midcircle* (also known as a *circle of antisimilitude*) of the two circles that bound the disks. The midcircle is defined by the property that an inversion through the midcircle takes one disk to the other.

In general the cells of $M(\mathcal{D})$ might not be connected; Fig. 3 shows an example in which they are not.

3.3 Radial Power Distance

Although the hyperbolic point of view makes Möbius invariance easy to prove, it is not easy to use in algorithms. As we show below, $M(\mathcal{D})$ can also be defined intrinsically in Euclidean geometry, as the minimization diagram of a particular function defined by a Euclidean construction.

Given a point q in the Euclidean plane, and a disk D_i defined by its center point c_i and radius r_i , let d_i be the Euclidean distance from q to c_i , and define the *radial power distance* from q to D_i to be $(d_i^2 - r_i^2)/2r_i$. We may similarly define the radial power distance for a disk complement $\bar{D}_i$, by using the same formula and letting r_i be the signed radius of the disk or disk complement. That is, r_i is a positive or negative number whose absolute value is the radius of the circle bounding the disk or disk complement and whose sign determines whether it is a disk or a disk complement.

Alternatively and more geometrically, for a point q and a disk D_i , the radial power distance from q to D_i is equal to the signed radius of a pair of congruent circles,

Fig. 4 *Left:* The power distance from the point outside the circle is the length of its tangent segment; the power distance from the point inside the circle is $-1/2$ the length of a chord bisected by the point. *Right:* The radial power distance from the point outside the disk is the radius of the two congruent circles tangent to each other at the point and both tangent to the disk; the radial power distance from the point inside the disk is the negative radius of the two congruent circles

tangent to each other at q and both tangent to D_i (Fig. 4). The sign of this distance is positive if q is outside D and negative if q is inside D . For a disk complement the definition is the same: the radial power distance is equal to the signed radius of a pair of congruent tangent circles, and is positive if q belongs to the disk complement and negative if it does not. It is straightforward to verify using the Pythagorean theorem that this geometric definition matches the algebraic one.

We note that this distance function has a similar form to the classical definition of the power of a circle, which may be defined algebraically as $d_i^2 - r_i^2$ or geometrically (for points q outside a given disk D_i) as the length of a line segment that has q as one endpoint and whose other endpoint is tangent to D_i . The power diagram [3] is the minimization diagram of the power distance, and we claim that our Möbius-invariant power diagram $M(\mathcal{D})$ coincides with the minimization diagram of the radial power distance.

To see the equivalence between the minimization diagrams for three-dimensional hyperbolic distance and for radial power distance, we use a geometric reformulation of the problem of finding nearest neighbors to a query point q . For a point q in hyperbolic space that is exterior to all of a set of halfspaces, the nearest halfspace may be found (just as in the Euclidean case) as the radius of the smallest sphere centered at q that touches at least one of the halfspaces, and this sphere is tangent to the halfspace or halfspaces it touches. Similarly, for a point q that is interior to some of the halfspaces, the minimum signed distance from q to a halfspace is the radius of the largest sphere centered at q that is contained entirely within at least one of the halfspaces, and again this sphere is tangent to the halfspace or halfspaces that minimize the signed distance. For a point q that is not within the hyperbolic space, but is instead a limit point of the space, the signed distance from q to the set of halfspaces defined from $\mathcal{D}$ is undefined (all curves in hyperbolic space from q to other points have infinite length) but we may still determine the identity of the nearest neighbor of q (that is, the cell of $M(\mathcal{D})$ containing q) in an analogous way. The geometric figure analogous to a set of

Fig. 5 Cross-section of a halfspace (*left*) and a horosphere externally tangent to it (*right*) in the halfspace model of hyperbolic space. The Euclidean plane onto which their shadows fall is the set of limit points of the model, and the *black curves* show the line qc (where q is the point at which the horosphere is tangent to the plane and c is the center of the Euclidean disk from which the halfspace was defined) and two circles tangent at q whose radius is the radial power distance

concentric spheres, for a limit point q of the hyperbolic space, is a set of *horospheres*, sets that appear in the halfspace model of hyperbolic space as spheres tangent at q to the boundary plane. If q is exterior to all of the given halfspaces, then a hyperbolic halfspace H_i is a nearest neighbor to q if there exists a horosphere centered at q that is tangent to H_i and disjoint from all of the other halfspaces. If q is contained in some of the halfspaces, then a hyperbolic halfspace H_i is a nearest neighbor to q if there exists a horosphere centered at q that is tangent to and contained in H_i , and that is not contained in any of the other halfspaces.

Now let D be any one of the input disks, and consider a planar cross-section of the halfspace model of three-dimensional hyperbolic space, defined by a cutting plane perpendicular to the boundary plane of the halfspace model that passes through q and through the (Euclidean) center c of D (Fig. 5). Let H be the hyperbolic plane defined from disk D ; then the intersection of H with the cross-sectional plane is a semicircle congruent to half of D . The cross-section of the horosphere centered at q and tangent to H is a circle, tangent to this semicircle. If this figure, consisting of a semicircle and tangent circle, is rotated by a right angle around line qc in either of two ways, the two rotated copies of the semicircle together form the boundary circle of D itself, and the two rotated copies of the horosphere section form two congruent circles tangent to each other at q and tangent to this boundary circle; these are the two circles used to define the radial power distance. Therefore, the nearest neighbor to q in the hyperbolic Voronoi diagram is the disk D for which we can find two congruent circles tangent to each other at q and tangent to D that are either internally tangent to D and as large as possible or (if no such disk exists) externally tangent to D and as small as possible. This is the same as the condition defining the nearest neighbor according to radial power distance, so the two diagrams coincide.

3.4 For Circle Packings

Our applications to graph drawing and soap bubbles use a special case of the diagram $M(\mathcal{D})$ in which the set $\mathcal{D}$ of disks and disk complements packs the plane, meaning

Fig. 6 The Möbius-invariant power diagram of a circle packing

that each connected component of $\mathbb{R}^2 \setminus (\cup \mathcal{D})$ is a triangular shape bounded by three mutually tangent circular arcs (Fig. 6).

Lemma 3 *Let $\mathcal{D}$ consist of three mutually tangent disks or disk complements with disjoint interiors. Then $M(\mathcal{D})$ is the image under a Möbius transformation of three infinite rays meeting at angles of $2\pi/3$ to each other. Each ray image is tangent to two of the three disks.*

Proof We may find a Möbius transformation μ taking $\mathcal{D}$ to three congruent (and still tangent) disks, with their centers on the vertices of an equilateral triangle. By symmetry, $M(\mu(\mathcal{D}))$ consists of three rays as described, bisecting the sides of the triangle. The result follows from Lemma 1. $\square$

Lemma 4 *Let $\mathcal{D}$ be a circle packing. Then for each disk D_i in $\mathcal{D}$, the corresponding cell C_i of $M(\mathcal{D})$ is simply connected, and is adjacent only to cells corresponding to disks tangent to D_i in the packing.*

Proof Let K_i be the intersection of the cells corresponding to D_i in the collection of diagrams $M(\{D_i, D_j\})$, where D_j ranges over the disks tangent to D_i . Clearly, $C_i \subset K_i$, because for every point q that does not belong to K_i , one of the disks D_j has smaller signed distance to q than D_i .

However, it is not possible for C_i to be a proper subset of K_i , because the sets of cells C_i and K_i both form a tessellation of the plane (a family of closed sets whose union covers the plane and whose interiors are disjoint). The sets C_i form a tessellation by their construction as the cells of a minimization diagram. The sets K_i form a tessellation by inspection: each of these sets consists of the disk D_i together with some portion of the neighboring triangular cusps between disks, so it cannot overlap with a set K_j that is not associated with a disk D_j tangent to D_i , each point within a disk D_i is covered by the corresponding set K_i , and each point within one of the triangular cusps is covered by one of the three corresponding sets K_i . If any C_i were a proper subset of K_i , the points in the difference of the two sets would not be covered by any cell C_i , a contradiction of the tessellation property.

Because C_i is a subset but not a proper subset of K_i , it must equal K_i , which by construction has the properties claimed in the lemma. $\square$

When $\mathcal{D}$ is a circle packing, each vertex of the drawing lies within the cusp formed by three mutually tangent circles. By Möbius invariance, this vertex must be placed at the isodynamic point of the triangle formed by the three points of tangency of this cusp. We may compute its position as the weighted average of the three tangency points, using the known formula for the barycentric coordinates of the isodynamic point to calculate its weights. Each circular arc of the drawing has two vertices as its end-points and passes through one tangency point, which together determine its location. These facts allow a simple and direct linear-time construction of $M(\mathcal{D})$ (given as input a combinatorial description of the neighbors to each disk D_i of $\mathcal{D}$ in clockwise order around D_i) using only Euclidean geometry, without any need to perform computations that involve hyperbolic coordinates or distances.

4 Lombardi Drawing

In this section we apply our Möbius-invariant power diagram to construct planar Lombardi drawings of all subcubic planar graphs, and of some four-regular planar graphs. However, as we show, not every four-regular planar graph has a planar Lombardi drawing.

4.1 Cubic Polyhedral Graphs

We first explain the simplest case of our Lombardi drawing algorithm, in which the planar graph to be drawn is 3-connected and 3-regular. Our drawings may be constructed by the following steps:

- Construct the dual graph of the given input graph, a maximal planar graph, and its (unique) planar embedding.
- Apply the Collins–Stephenson procedure to realize the dual maximal planar graph as the intersection graph of a collection of tangent circles C_i .
- Use quasiconvex programming to find a Möbius transformation of the circles taking them to a configuration in which one circle C_0 is exterior to all the others, and maximizing the minimum radius of the internal circles.
- Let $\mathcal{D}$ be the set of disks and disk complements associated with the transformed circle packing, with one disk complement for the exterior circle C_0 and a disk for each remaining circle. Construct the Möbius-invariant power diagram $M(\mathcal{D})$, place a vertex of the drawing at each point where three power diagram cells meet, and place an edge of the drawing on each arc separating two adjacent power diagram cells (Fig. 7).

Theorem 1 *If we are given as input an n -vertex 3-regular 3-connected planar graph, we can produce a planar Lombardi drawing for the graph in time $T + O(n)$, where T denotes the time needed to find a circle packing dual to the given graph.*

Proof The fact that this is a Lombardi drawing (with circular arcs meeting at angles of $2\pi/3$) follows from Lemma 3. Once a circle packing is constructed, the phase of the algorithm that computes a Möbius transformation maximizing the minimum circle

radius can be performed in linear time [4]. The linear time bound for the construction of $M(\mathcal{D})$ is as described in Sect. 3.4. $\square$

We implemented in Python the algorithm for 3-connected graphs, including the Collins–Stephenson circle packing algorithm and a numerical improvement method for finding optimal Möbius transformations. Our implementation takes as input a text file with one line per vertex; each line lists the identifiers for a vertex and its three neighbors in clockwise order. The output drawing is represented in the SVG vector graphics file format. Figure 8 shows some drawings created by our implementation.

4.2 Two-Connected Subcubic Graphs

We next describe how to extend our Lombardi drawing technique to 2-connected graphs. A 2-connected graph may have vertices of degree two or three, but the degree-2 vertices may be suppressed, forming a 3-regular multigraph with the same connectivity. If the multigraph has a planar Lombardi drawing, so does the original graph with the degree-2 vertices, as the vertices may be restored by subdividing edges (placing new vertices in the interior of the circular arc representing the edge) without changing the Lombardi property. Therefore, within this section we assume that the given graph is 3-regular.

We use a standard tool for decomposing 2-connected graphs, the SPQR tree [13, 14, 26, 33, 41]. An SPQR tree for a graph G is a tree structure in which each tree node is associated with a graph C_i , known as a 3-connected component of G . In each 3-connected component, some of the edges are labeled as “virtual”, and each edge of the SPQR-tree is labeled by a pair of oriented virtual edges from its two endpoints; each virtual edge of a component is associated in this way with exactly one tree edge. The given graph G may be formed by gluing the components C_i together, by identifying pairs of endpoints of virtual edges according to the labeling of the tree edges, and then deleting the virtual edges themselves. The nodes of an SPQR tree have three types: R nodes, in which the associated graph is 3-connected, S nodes, in which the associated graph is a cycle, and P nodes, in which the associated graph is a *bond graph*, a multigraph with two vertices and three or more parallel edges. With

Fig. 7 Lombardi drawing of the Frucht graph [24], one of the two smallest 3-regular asymmetric graphs, derived from the circle packing of Fig. 2 and the power diagram of Fig. 6

Fig. 8 Sample drawings from our implementation. (a) Markström's 24-vertex graph with no 4- or 8-cycles [42]. (b) Došlić's 38-vertex graph with girth five and cyclic edge connectivity three [15]. (c) Tutte's 46-vertex non-Hamiltonian cubic planar 3-connected graph [56]. (d) Grinberg's 46-vertex non-Hamiltonian graph with cyclic edge connectivity five [25]. (e) 46-vertex Halin graph formed from a complete ternary free tree. (f) 252-vertex hexagonal mesh. (g) 60-vertex buckyball or truncated icosahedron. (h) Irregular 69-vertex graph

Fig. 9 *Left:* The SPQR tree of a planar 3-regular and bridgeless but not 3-connected graph; the *dashed black* edges connect pairs of virtual edges to be glued together. *Right:* A Lombardi drawing formed by transforming drawings of each 3-connected component so that sets of virtual edges connected to the same S node all lie on the same circle, removing the arcs of this circle coming from the virtual edges in the P and R nodes, and adding arcs for the real edges of the S node

the additional constraint that no two S nodes and no two P nodes may be adjacent, the SPQR tree is uniquely determined from G , and may be constructed in linear time.

The SPQR trees of 3-regular graphs have an additional structure that will be helpful for us:

Lemma 5 (Potheri [46], Eppstein and Mumford [22]) *A 2-connected graph G is 3-regular if and only if each edge in its SPQR tree has exactly one S node as an endpoint, each S node is associated with an even cycle that alternates between virtual and non-virtual edges, each P node is associated with a three-edge bond graph, and each R node is associated with a graph that is itself 3-regular.*

Lemma 6 *Let G be a graph with a planar Lombardi drawing, and let e be any edge of G . Then, for any $\varepsilon > 0$ there is a planar Lombardi drawing of G in which e lies on the outer face of the drawing and forms a circular arc that covers an angle of its circle greater than or equal to $2\pi(1 - \varepsilon)$.*

Proof Transform the given drawing by an inversion whose center is near but not on the midpoint of e . $\square$

Theorem 2 *If we are given as input an n -vertex 2-connected planar graph with maximum degree 3, we can produce a planar Lombardi drawing for the graph in time $T + O(n)$, where T denotes the time needed to find and optimally transform a family of circle packings with total cardinality $O(n)$.*

Proof Suppress the degree-two vertices of the given graph to produce a 3-regular 2-connected graph, decompose it into an SPQR tree with the additional structure of Lemma 5, and use Theorem 1 to construct a planar Lombardi drawing of the 3-connected component associated with each R node. In each S node of the tree, glue the drawings from the adjacent tree nodes together by using Lemma 6 to expand the virtual edges, move these edges to the outer face of their drawings, and shrink the rest of the drawings, and then align the circular arcs representing their virtual edges so that they all lie on a common circle (Fig. 9). Finally, subdivide the edges of the drawing as necessary to restore the suppressed degree-two vertices. $\square$

Fig. 10 Modifying a 3-connected Lombardi drawing to attach bridges to subdivision points along an edge

4.3 Subcubic Graphs with Bridges

We are finally ready to describe our algorithm for constructing planar Lombardi drawings for arbitrary subcubic planar graphs, possibly including bridge edges.

Our algorithm begins by deleting all the bridge edges from the graph, and suppressing all degree two vertices, leaving a collection of isolated vertices and 2-connected 3-regular subgraphs. It then constructs the SPQR tree of each 2-connected subgraph, forming its 3-connected components, and applies Theorem 1 to construct a planar Lombardi drawing for each R node in each SPQR tree.

Within each face of each of these drawings, form a sequence of circular arcs forming 30° angles to the arcs of the face and 60° angles to each other, forming a smaller inset polygon (Fig. 10, left). These arcs cannot intersect each other, as can be seen by considering separately each triangular cusp of the circle packing from which the drawing was constructed. Within each circle of the circle packing (viewing the circle as a disk model for the hyperbolic plane) the new arcs are confined to disjoint halvespaces. For each edge e of one of these drawings that corresponds to a path of suppressed degree-two vertices some of which were bridge endpoints, let k be the number of bridge endpoints of e . Choose arbitrarily one or the other of the two circular arcs A forming 30° angles to e , and replace e by a sequence of $k + 1$ circular arcs, all of which form the same 30° angle to A . These arcs meet each other at 120° angles, and we may extend a bridge from each of the vertices formed in this way to a new degree-one vertex (Fig. 10, right). By a similar construction we unsuppress and attach bridge edges to the suppressed bridge endpoints within the P nodes and S nodes of the SPQR tree.

Next, we glue the components within each SPQR tree together, as in Theorem 2. At the same time, for each isolated vertex created by the bridge deletion step, we create a drawing of a claw $K_{1,3}$ consisting of three unit-length line segments meeting at 120° angles. Additionally, we unsuppress the remaining degree-two vertices of the original graph, by subdividing edges of these drawings. After this step, we have separate drawings for each block (biconnected component) of the original graph, in which each bridge incident to the block has a degree-one vertex at its other end.

Finally, for each two blocks that should be connected by a bridge, we apply an inversion centered on the degree-one endpoint of the bridge, causing the bridge edge to become an infinite ray exterior to the drawing. After this transformation the two blocks may be glued together by making their two copies of the bridge edge lie on the same line. This completes the proof of the main result of this section:

Theorem 3 *If we are given as input an n -vertex planar graph with maximum degree 3, we can produce a planar Lombardi drawing for the graph in time $T + O(n)$, where T denotes the time needed to find and optimally transform a family of circle packings with total cardinality $O(n)$.*

4.4 Halin Graphs

The method of constructing Lombardi drawings from a packing of disjoint tangent circles, used above for 3-regular polyhedral graphs, applies to some other graphs as well. We define a system of interior-disjoint disks $\mathcal{D}$ to be *well-centered* if the Möbius-invariant power diagram $M(\mathcal{D})$ forms a system of arcs and vertices that forms a Lombardi drawing of the planar dual to the intersection graph of $\mathcal{D}$. Equivalently, $\mathcal{D}$ is well-centered when its diagram has one vertex within each connected component of the complement $\mathbb{R}^2 \setminus \mathcal{D}$ of $\mathcal{D}$, and when all cell boundaries incident to that vertex meet at equal angles. If a planar graph G has a dual graph that can be represented by a well-centered system of disks $\mathcal{D}$, then $M(\mathcal{D})$ will form a Lombardi drawing of G .

As an example of a system of disks that is not a disk packing but is well-centered, consider k congruent disks with centers all on a common circle, sized so that each disk is tangent to its two neighbors on the circle. In this case, the intersection graph of $\mathcal{D}$ is a k -cycle, and $M(\mathcal{D})$ is a system of k rays meeting at equal angles. By performing an appropriate Möbius transformation of $\mathcal{D}$, we may perturb it so that $M(\mathcal{D})$ forms a system of k circular arcs, meeting at equal angles at two points, one of which is the transformation of the center of the circle and the other of which is the transformation of ∞ . This is a Lombardi drawing of a multigraph with k edges and two vertices, the planar dual of the k -cycle.

As a less trivial example, let G be any *Halin graph* [29], a 3-connected planar graph formed from a plane tree (with no degree-two vertices) by adding a cycle connecting its leaves in the order given by the plane embedding of the tree. As a polyhedral graph, G can be represented (uniquely up to Möbius transformation) by a pair of circle packings $\mathcal{D}$ and $\mathcal{D}'$, where the intersection graph of $\mathcal{D}$ is the dual graph of G , the intersection graph of $\mathcal{D}'$ is G itself, and adjacent pairs of vertices in G and its dual (corresponding to an incident vertex-face pair in G) are represented by orthogonal circles in $\mathcal{D}$ and $\mathcal{D}'$. We claim that $\mathcal{D}$ is well-centered. For, if v is a vertex on the outer face of G , it has degree three and is surrounded by three mutually tangent disks, for which well-centeredness follows from Lemma 3. Alternatively, if v is an interior vertex of G , then the k circles surrounding v (where k is the degree of v in G) are all orthogonal to the circle representing v and tangent to the circle representing the outer face of G . It follows that the set of circles surrounding v is Möbius-equivalent to the system of k concyclic circles described above for the k -cycle, and so is well-centered. Therefore, $M(\mathcal{D})$ is a Lombardi drawing of G .

It was previously known that every Halin graph had a planar Lombardi drawing [18] but the previous construction involved certain arbitrary choices. In contrast, this method, based as it is on the unique orthogonal circle packing for 3-connected planar graphs and their duals, provides a principled and unique way of making those choices for every Halin graph.

Fig. 11 Dual circle packings for a cube and octahedron, with a Lombardi drawing of their medial graph, the cuboctahedron

4.5 Four-Regular Medial Graphs

A variant of the circle packing theorem states that every 3-connected planar graph G and its dual can be simultaneously represented by tangent circles, so that two circles from the two packings are orthogonal if they represent a vertex v of G and a face of G containing v , and disjoint otherwise. The packing is unique up to Möbius transformation, and may be found by a numerical procedure similar to the one for circle packing of maximal planar graphs [43]. These dual packings allow us to construct a planar Lombardi drawing of the *medial graph* of G , the 4-regular graph formed by placing a vertex on the midpoint of each edge of G and connecting two of these new vertices by an edge whenever they are the midpoints of consecutive edges on the same face. Each pair of orthogonal circles in the packing and dual packing of G form a lune, which may be bisected by a circular arc forming a $\pi/4$ angle with both circles. When two of these bisecting circular arcs meet at the same point of tangency of the circle packing, they form an angle of $\pi/2$ or π to each other. Therefore, we have a Lombardi drawing of the four-regular graph that has a vertex at each tangency point of the packing and an edge for each of these bisecting arcs. This graph is the medial graph of G (Fig. 11). As was the case for 3-regular polyhedral graphs, this drawing has an interpretation as the Möbius-invariant power diagram $M(\mathcal{D})$ of a collection of disks $\mathcal{D}$ that includes a disk for each circle in the circle packing for G and its dual.

As with our technique for finding Lombardi drawings of subcubic graphs, the drawings constructed in this way can be extended to certain 2-connected graphs using the SPQR tree. However, this technique does not apply to 3-connected 4-regular planar graphs that are not medial graphs of polyhedral graphs, such as the one shown in Fig. 12. More strongly, in the next section we show that there exist 4-regular planar graphs that have no planar Lombardi drawing.

4.6 A 4-Regular Graph with No Planar Lombardi Drawing

The 4-regular 2-connected planar graph G_{18} in Fig. 13 has no planar Lombardi drawing, as we now prove.

Fig. 12 A 4-regular polyhedron that is not the medial graph of a polyhedral graph, described by Dillencourt and Eppstein [21]

Fig. 13 A 4-regular 2-connected planar graph G_{18} that has no planar Lombardi drawing

The important parts of this graph, for our purposes, are the six labeled vertices; the remaining vertices are present only to make the graph 4-regular, and could be contracted to single vertices to produce a simpler eight-vertex counterexample to the existence of planar Lombardi drawings while sacrificing regularity. Note that (up to automorphisms) the labeled part of G_{18} has a unique planar embedding, and that the two quadrilateral faces $abfc$ and $adfe$ have their two opposite vertices a and f in common. Because G_{18} is 4-regular, it has a Lombardi drawing with a circular layout, with the arcs representing its edges forming 45° angles to the circle on which the vertices are drawn [18]. However, a circular Lombardi drawing of this type is not guaranteed to be planar, and we will show that G_{18} does not have a planar Lombardi drawing.

A key idea in the proof is to relate quadrilateral faces of any possible drawing, such as faces $abfc$ and $adfe$, with circles, as the following lemma shows.

Lemma 7 *Let $pqrs$ be a quadrilateral face of a planar Lombardi drawing, and suppose that all four of its vertices have degree four. Then there exists a circle or line through all four of the points p , q , r , and s , that does not cross any edges of $pqrs$.*

Proof Perform an inversion through a circle centered at p , sending p to ∞ . The transformed copy of $pqrs$ includes two perpendicular rays qp and sp , which we may assume without loss of generality are axis-parallel, as shown in Fig. 14.

If both qr and rs are drawn as curved arcs (rather than straight line segments) in the transformed drawing, let θ be the angle qr subtends on its circle; in order for all the angles of the figure to be right, rs must subtend the complementary angle $\pi - \theta$. The angle between arc qr and line segment qr (at either of the two intersection points of

Fig. 14 Figure for proof of Lemma 7

these two curves) is $\theta/2$, and symmetrically the angle between arc rs and line segment rs is $(\pi - \theta)/2$. Angle qrs is equal to the sum of these two angles together with the right angle between the two arcs; that is, it is $\theta/2 + \pi/2 + (\pi - \theta)/2 = \pi$. Thus, q , r , and s are collinear.

If one of qr and rs (say qr) is not a curved arc, but a straight line segment, it must be axis-parallel in order to form a right angle at q , and the other arc rs must be a semicircle in order to form right angles at r and s . Again, in this case, q , r , and s are collinear.

In either case, the line O through the three points q , r , and s also passes through $p = \infty$ (because every line passes through ∞). Reversing the inversion we initially performed, O becomes the desired circle. It can't cross the circular arcs of pqr because two circles can only intersect at two points and each of these arcs already intersects O at its two endpoints. $\square$

Theorem 4 *Graph G_{18} has no planar Lombardi drawing.*

Proof If G_{18} had a planar Lombardi drawing, then by Lemma 7 its quadrilateral faces $adfe$ and $abfc$ would have circumscribing circles. If they do not coincide, these circles intersect at the two vertices a and f ; the other four vertices must be placed between them on the four circular arcs connecting these two intersection points. By performing an inversion centered within face $abfc$ we may assume without loss of generality that the drawing resembles the schematic layout shown in Fig. 15, in the sense that face $abfc$ lies outside of its circle. The other circle, $adfe$, must contain one but not both of the vertices b and c ; without loss of generality we assume it contains c and does not contain b , because the other case is symmetric under a relabeling of the vertices. Then (in order for the drawing to be planar) vertex d may be placed anywhere along the arc of circle $adfe$ closest to b . The ray from c radially outwards from the center of circle $abfc$ lies entirely within face $abfc$, and crosses circle $adfe$; vertex e may be placed on its arc either above or below this ray, but by symmetry (flipping the diagram and relabeling the vertices if necessary) we assume that it is above the ray, as depicted schematically in the figure.

Fig. 15 Figure for proof of Theorem 4

For this placement, the edge ce must exit vertex c towards the interior of circle $abfc$, because it forms a 180° angle with either edge ac or edge fc , both of which are exterior to the circle. It must then cross circle $abfc$ again, turning clockwise as it does, in order to stay within quadrilateral $ae fc$ and reach vertex e . By a symmetric argument it must exit vertex e towards the exterior of circle $adfe$, because it forms a 180° angle with either edge ae or edge fe , both of which are interior to the circle. Thus, coming from c , it must cross circle $adfe$, turning counterclockwise as it does, to reach the circle a second time. But in a Lombardi drawing, edge ce must be drawn as a circular arc, which cannot turn clockwise for part of its length and counterclockwise for part of its length. Therefore, this placement is impossible.

The remaining case arises when circles $abfc$ and $adfe$ coincide. But in this case, by similar reasoning, edge ce must exit c towards the interior of the single circle, cross the circle, and reach e from the exterior, an impossibility since a circular arc cannot have three intersection points with a circle. $\square$

5 Soap Bubbles

As we show in this section, our Möbius-invariant power diagram can be used to construct planar soap bubble clusters with any realizable topology, leading to a graph-theoretic characterization of these clusters.

5.1 A Mathematical Model of Soap Bubbles

Two classical results describe the geometry of soap bubbles. *Plateau's laws*, observed experimentally in the 19th century by Joseph Plateau [45] and proven rigorously for minimal surfaces in 1976 by Jean Taylor [54], state that in a three-dimensional soap bubble cluster,

- Each two-dimensional surface has constant *mean curvature*. That is, the two principal curvatures of the surface take the same average value at all points in the surface.

- At each one-dimensional junction of surfaces, exactly three surfaces meet, and they form dihedral angles of $2\pi/3$ with each other. Such a junction is called a *Plateau border*. A Plateau border may either be a simple closed space curve or a curve with two endpoints, in either case touching the same three surfaces for its entire length.
- At each endpoint of a Plateau border, exactly four Plateau borders and six two-dimensional surfaces meet, and the borders form angles of $\cos^{-1}(-1/3)$ with each other (the same angle that is formed by rays from the center of a regular tetrahedron through its vertices).

The *Young–Laplace equation*, also formulated in the 19th century by Thomas Young and Pierre-Simon Laplace, states that the mean curvature of a two-dimensional surface in a soap bubble cluster is proportional to the difference in pressure on the two sides of the surface, with a constant of proportionality determined by the surface tension of the fluid forming the soap bubbles [36].

In the case of a soap bubble cluster formed between two flat plates, the bubble walls are perpendicular to the plates, and all cross-sections of the bubbles by a plane parallel to one of the plates are congruent to each other. In this case, the principal curvature in the direction perpendicular to the plates is zero, and we may simplify both laws to describe the planar figure formed by the cross-sections.

Thus, the two-dimensional restriction of Plateau's laws state that:

- The figure consists of one-dimensional curves of constant curvature; that is, circular arcs or line segments that do not cross each other.
- At each endpoint of one of these arcs or segments, exactly three curves meet, and they form angles of $2\pi/3$ with each other.

Correspondingly, the two-dimensional restriction of the Young–Laplace equation states that:

- The curvature of any one of the circular arcs (the inverse of its radius) is proportional to the difference in pressure between the bubbles it separates. Bubbles with the same pressure as each other are separated by line segments, with zero curvature.

We define a planar soap bubble cluster to be a finite collection of curves, obeying the planar restrictions of Plateau's laws and the Young–Laplace equation (for some assignment of pressures to each bubble). Equivalently, in graph-theoretic terminology, it is a planar embedding of a graph with circular-arc edges, forming angles of $2\pi/3$ at each vertex, for which it is possible to find a pressure assignment to the bubbles that is consistent with all of the arc curvatures. As we describe in Sect. 5.2, we can replace this existential condition (the existence of a consistent pressure assignment) with a simpler statement about the curvatures of triples of arcs meeting at vertices. In a planar soap bubble cluster, the local forces on each vertex or arc caused by pressure and surface tension will sum to zero, so it will be in an equilibrium (though possibly an unstable one [23,58]).

As our results will later show, the Lombardi drawings depicted in Figs. 6 and 8 (minus the vertex disks) form examples of soap bubble clusters. In contrast, Fig. 16 depicts a planar 3-regular Lombardi drawing that is not a soap bubble cluster, because

Fig. 16 A collection of segments and arcs that obeys Plateau's laws, but does not form a planar soap bubble cluster, because it has no valid pressure assignment consistent with its arc curvatures. Three of its segments do not separate different regions, impossible for soap bubbles

it disobeys the Young–Laplace equation. There is no assignment of pressures to the bubbles in this figure that is consistent with the curvatures of its arcs, and at the six asymmetric vertices in the figure, the curvatures of the three incident arcs do not sum to zero, violating the local reformulation of the Young–Laplace equation described in Sect. 5.2.

5.2 Localization of the Young–Laplace Equation

It is convenient to replace the Young–Laplace equation with a more local criterion, using a signed variation of curvature. Given a circular arc ending at a point p , of curvature $c \geq 0$, we define the *signed curvature* of the arc at p to be c whenever the arc curves clockwise as it leaves p , and to be $-c$ whenever it curves counterclockwise. For line segments, of course, the signed curvature is zero.

Lemma 8 *A finite collection of circular arcs and line segments forms a planar soap bubble cluster if and only if it obeys Plateau's laws and, at each endpoint of an arc or segment, the sum of the signed curvatures of the three incoming curves is zero.*

Proof In one direction, in a planar soap bubble cluster, let x , y , and z be the pressures of three bubbles that meet at a point. By assumption, the cluster obeys the Young–Laplace equation, according to which the three signed curvatures of the arcs at the point are proportional to the pressure differences $x - y$, $y - z$, and $z - x$. The fact that these three signed curvatures sum to zero follows from the simple algebraic fact that $(x - y) + (y - z) + (z - x) = 0$.

In the other direction, suppose that finitely many circular arcs obey Plateau's laws and have signed curvatures adding to zero at each point where three arcs meet. We must assign pressures to the bubbles formed by the arcs, such that the curvatures obey the Young–Laplace equation. Assign zero pressure to the outside region. For each bounded bubble B , choose arbitrarily a curve c that starts in the outside region, ends within B , meets the arcs only at proper crossing points interior to an arc, and has finitely many crossing points. Set the pressure within B to the sum of the pressure differences determined by the curvatures of each crossed arc.

To show that this system of pressures obeys the Young–Laplace equation, consider any two bubbles B_1 and B_2 separated by an arc A , and let c_1 and c_2 be the curves from the outside region to B_1 and B_2 by which their pressures were determined. We may link c_1 and c_2 by another curve that remains entirely within the outside region, forming a single (possibly self-crossing) curve that connects B_1 to B_2 , and we may continuously deform this curve, crossing finitely many triple points of the system of arcs as we do, until it forms a line segment crossing arc A . Each time this deformation causes the curve to cross a triple point, the assumption that the curvatures at that point add to zero ensures that the sum of the pressure differences of crossings along the curve remains constant. Before the deformation, this sum was the pressure difference between B_1 and B_2 in our system of pressures, and after the deformation, it is the pressure difference required between bubbles B_1 and B_2 for the curvature of A to obey the Young–Laplace equation. Since these two pressure differences are equal, A (and by the same argument every arc) obeys the equation. $\square$

An analogous result for bubbles on the sphere instead of the plane is in Quinn Maurmann’s appendix to [47].

When three three-dimensional bubbles form a cluster in which all surface patches are spherical, it was known to Plateau that the three inner surface patches have centers of curvature that all lie on a single line [31]. Although centers of circles and collinearity of points are not Möbius-invariant properties, we can restate the same collinearity property for planar bubble clusters in a Möbius-invariant way, in terms of the crossing points of the circles. This restatement will allow us to prove that the Möbius transformation of a planar soap bubble cluster is itself another planar soap bubble cluster.

Lemma 9 *Let three circular arcs meet at angles of $2\pi/3$ at a point X , with centers of curvature C_i for $i \in \{1, 2, 3\}$, and let $r_i = |XC_i|$ be the radii of each of the three arcs. Then the following three conditions are equivalent:*

1. *The sum of the three signed curvatures of the arcs is zero.*
2. *The three center points C_i are collinear.*
3. *The three circles with centers C_i and radii r_i have two triple crossing points.*

Proof We partition the proof into four implications between the three conditions of the lemma.

(1) $\Rightarrow$ (2): Suppose that the three signed curvatures sum to zero; then one must have a different sign than the other two. Without loss of generality (by permuting the indices and by mirroring the configuration, if necessary) we may assume that the signed curvatures are $1/r_1$, $-1/r_2$, and $1/r_3$. By assumption, the sum of these three quantities is zero; by multiplying each term of this sum by $r_1r_2r_3$ and rearranging, we obtain the equation $r_1r_2 + r_2r_3 = r_1r_3$.

Note that the three lines XC_i form angles of $\pi/3$ with each other, because they are perpendicular to the arcs, which meet at angles of $2\pi/3$. The assumption on the signs of the curvatures implies that angles C_1XC_2 and C_2XC_3 must both equal $\pi/3$, and angle C_1XC_3 must equal $2\pi/3$. Now consider the areas of the three triangles C_1XC_2 , C_2XC_3 , and C_1XC_3 . In any triangle, the area can be computed by the side-angle-side

Fig. 17 Arcs meeting the conditions of Lemma 9

formula as half the product of two adjacent side lengths with the sine of the angle formed by the same two sides. Thus, these triangle areas are $\frac{1}{2}r_1r_2 \sin \frac{\pi}{3}$, $\frac{1}{2}r_2r_3 \sin \frac{\pi}{3}$, and $\frac{1}{2}r_1r_3 \sin \frac{2\pi}{3}$ respectively. But $\sin \pi/3 = \sin 2\pi/3$, so the already-obtained equation $r_1r_2 + r_2r_3 = r_1r_3$ implies that triangles C_1XC_2 and C_2XC_3 together have the same total area as triangle C_1XC_3 . This could only happen if the three centers C_1 , C_2 , and C_3 are collinear, for otherwise the sum of the areas of C_1XC_2 and C_2XC_3 would differ from the area of C_1XC_3 by the area of triangle $C_1C_2C_3$, which is zero only when these three points are collinear.

(2) $\Rightarrow$ (1): By (2) the three center points C_i are collinear; assume without loss of generality that C_2 lies between C_1 and C_3 on their common line. Because the three centers C_1 , C_2 , and C_3 form angles of $\pi/3$ rather than $2\pi/3$, the signed curvature of the middle center C_2 has the opposite sign to the signed curvature of the other three curvatures; we may assume without loss of generality (by mirror reversing the configuration if necessary) that these three signed curvatures are $1/r_1$, $-1/r_2$, and $1/r_3$. As in the previous case, the three lines XC_i form angles of $\pi/3$ with each other so angles C_1XC_2 and C_2XC_3 must both equal $\pi/3$, angle C_1XC_3 must equal $2\pi/3$, and by collinearity the two triangles C_1XC_2 and C_2XC_3 together disjointly cover the same region of the plane as the single triangle C_1XC_3 .

We can apply the side-angle-side formula to this region in two different ways, giving the equation

$$\frac{1}{2}r_1r_2 \sin \frac{\pi}{3} + \frac{1}{2}r_2r_3 \sin \frac{\pi}{3} = \frac{1}{2}r_1r_3 \sin \frac{2\pi}{3}.$$

But the factors of $1/2$ cancel, as do the factors of $\sin \pi/3 = \sin 2\pi/3$, leaving the simpler equation $r_1r_2 + r_2r_3 = r_1r_3$. Dividing all terms by $r_1r_2r_3$ gives $1/r_3 + 1/r_1 = 1/r_2$, and rearranging gives $1/r_1 - 1/r_2 + 1/r_3 = 0$ as desired.

(2) $\Rightarrow$ (3): The three circles have at least one triple crossing point, at X . Let ℓ be the line through the three centers, assumed to exist by (2). Then because ℓ passes through each circle center, a reflection across ℓ is a symmetry of each circle and therefore of the whole configuration of three circles. The point X cannot lie on ℓ , because two circles centered on a line cannot cross at a point that is also on the line, they can only meet at a point of tangency, violating the assumption that the arcs meet at angles of $2\pi/3$. Therefore, the reflection of X across ℓ is also a triple crossing point, and the three circles have two triple crossings as (3) states.

(3) $\Rightarrow$ (2): For any two intersecting circles, it is necessarily true that their centers lie on the perpendicular bisector of the chord connecting their two intersection points. But by the assumption of (3) that there are two triple crossing points, the chords defined by each pair of the three given circles coincide; therefore, their perpendicular bisectors also coincide in a line ℓ that contains all three circle centers. $\square$

Figure 17 shows a set of arcs, circles, centers, and radii meeting the conditions of the lemma. The lemma applies directly only to circular arcs and not to straight line segments; for line segments we need a more careful statement.

Lemma 10 *Let three circular arcs or straight line segments meet at angles of $2\pi/3$ at a point X . Then the following two conditions are equivalent:*

1. *The sum of the three signed curvatures of the arcs and segments is zero.*
2. *The three circles or lines containing the three arcs or line segments have two triple crossing points (possibly including the point at infinity).*

Proof We consider different cases for each possible number of straight line segments in the input configuration.

- When there are no segments and three circular arcs, the lemma is just a simplified formulation of Lemma 9.
- When there is one line segment and two circular arcs, the curvature sum is zero if and only if the two arcs are reflections of each other across the line segment (otherwise their two curvatures cannot cancel) and there is a second triple crossing point if and only if the same reflection condition is met (otherwise the two circles will cross the line in different points).
- When there are two line segments and one circular arc, the sum of the curvatures is always nonzero (it equals the curvature of the circular arc), and there is only one triple crossing point (the two lines have their second intersection at infinity and their intersections with the circle at finite points).
- When there are three line segments, the sum of the curvatures is always zero, as all curvatures are themselves zero, and the three lines meet at X and at the point at infinity, so both conditions are true.

Since the equivalence holds in each possible case, it holds in general. $\square$

This triple crossing property is closely related to Lemma 5.1 of [60] and to the existence and uniqueness of the standard double bubble, for which see e.g. Proposition 14.1 of [44].

5.3 Möbius Invariance

Möbius transformations do not preserve physically meaningful quantities of soap bubbles such as their pressures or the total length or area of soap film. They do not even preserve the finiteness of the bubbles: it is possible to find a transformation that takes the unbounded region of the plane to a bounded region and vice versa. Nevertheless, surprisingly, they transform soap bubble clusters into other soap bubble clusters. We outline a proof of Möbius invariance below.

Lemma 11 *Let B be a planar soap bubble cluster, and let μ be a Möbius transformation with the property that no curve or arc of B is transformed into a ray or double ray. Then $\mu(B)$ is also a planar soap bubble cluster.*

Proof By Lemmas 8 and 10, B consists of arcs or segments meeting at angles of $2\pi/3$ such that each three arcs that meet have circles forming two triple crossings. The meeting angles and the triple crossings of circles are preserved by Möbius transformation, so $\mu(B)$ is also a collection of arcs or segments meeting at angles of $2\pi/3$ such that each three arcs that meet have circles forming two triple crossings. By Lemma 10 again, the signed curvatures at each meeting point of $\mu(B)$ add to zero, and by Lemma 8 again, $\mu(B)$ is a planar soap bubble cluster. $\square$

5.4 Bridgelessness

Based on the Möbius invariance of planar soap bubble clusters, we prove in this section the necessity of our graph-theoretic characterization: every planar soap bubble cluster comes from a graph that is planar, 3-regular, and 2-connected.

It is clear that a planar soap bubble cluster must form a planar graph, and Plateau's laws imply that it must be 3-regular, so the nontrivial part of our characterization of soap bubble graphs is the claim that these graphs are 2-vertex-connected, or equivalently (since they are three-regular) that they have no bridges. In the language of soap bubbles, this means that every arc or segment in the cluster separates two different bubbles. As we now prove, this is necessarily true for a planar soap bubble cluster.

Lemma 12 *In a planar soap bubble cluster, no segment or arc can have the same bubble on both of its sides.*

Proof Assume for a contradiction that the cluster includes a segment or arc e that has the same bubble on both sides. Nonzero curvature of e would contradict the Young–Laplace equation, as the pressure is equal on both sides of e . But if e is a line segment, then it is straightforward to find a Möbius transformation that transforms e into a curved arc. The resulting cluster again contradicts the Young–Laplace equation, and by Lemma 11, the untransformed cluster also contradicts the same equation. $\square$

5.5 Realizing Graphs as Soap Bubbles

As we now show, the steps of our Lombardi drawing algorithm that apply to 2-connected or 3-connected graphs generate soap bubble clusters, giving us an algo-

rithmic proof that our necessary conditions for a graph to be the graph of a soap bubble cluster are also sufficient.

Lemma 13 *Let $\mathcal{D}$ be a set of three mutually tangent circles. Then their Möbius-invariant power diagram $M(\mathcal{D})$ is a double bubble formed by three circular arcs, line segments, or rays, meeting in two triple crossing points at angles of $2\pi/3$, with signed curvatures summing to zero.*

Proof By Lemma 3, $M(\mathcal{D})$ is the image under a Möbius transformation of three rays meeting at angles of $2\pi/3$. These rays meet in two triple crossing points, the common apex of the rays and the point at infinity, and this triple crossing property is preserved under Möbius transformation. The claim that the curvatures of the transformed rays sum to zero follows from Lemma 10. $\square$

Theorem 5 *A multigraph G can be realized as the arcs and junctions of a planar soap bubble cluster if and only if G is planar, 3-regular, and bridgeless.*

Proof Planarity and 3-regularity follow from Plateau's laws. The requirement that a planar soap bubble cluster be bridgeless is Lemma 12. It remains to show that every planar, 3-regular, and bridgeless graph G may be realized as a soap bubble.

If G is planar, 3-regular, and 3-connected, then (as in Sect. 4.1) G may be drawn as the Möbius-invariant power diagram $M(\mathcal{D})$ for a circle packing $\mathcal{D}$ representing the dual graph of G . At each vertex of the diagram, the regions for three mutually tangent circles of $\mathcal{D}$ meet. By Lemma 13, the vertex is incident to three circular arcs meeting at angles of $2\pi/3$ (satisfying Plateau's laws), with signed curvatures summing to zero. Therefore, by Lemma 8, the diagram forms a planar soap bubble cluster representing G .

Finally, if G is bridgeless, planar, and 3-regular but not 3-connected, it must be 2-connected. In this case, as in Sect. 4.2 and Fig. 9, it may be formed by constructing an SPQR tree representing the 3-connected components of G , separately drawing each 3-connected component as a double bubble or as the Möbius-invariant power diagram of a circle packing, transforming these drawings so that they do not cross each other and so that the virtual edges that connect to each S node lie on a common circle, deleting the virtual edges, and drawing the real edges of the S node along the shared circle. As in the 3-connected case, the drawing of each component obeys Plateau's laws and has zero curvature sum at each vertex. By Lemma 11, these properties are preserved by the step of the drawing algorithm that transforms each component. The step of the algorithm that deletes the virtual edges and replaces them by real edges does not change the local geometry at any vertex of the drawing, so again it preserves the properties of obedience to Plateau's laws and zero local curvature sum. Thus, after these steps we have a drawing that continues to obey these properties. By Lemma 8, the diagram forms a planar soap bubble cluster representing G . $\square$

We note that, although our circle packing method is capable of realizing all graphs of planar soap bubble clusters, it is not capable of generating all valid geometries of planar soap bubble clusters. In particular, in the 3-connected case, the clusters we generate have the property that within each bubble there is a circle tangent to all the arcs of the bubble, but this property does not hold for arbitrary soap bubble clusters. For instance, not all of the bubbles in Fig. 1 have this property.

5.6 Stability

So far, we have been considering only whether a bubble cluster is in equilibrium or not, but we may differentiate what kind of equilibrium a cluster is in by its stability, describing the relation between the energy of the equilibrium to the energy of nearby configurations. To do this, we need some additional definitions. In the case of bubble clusters, we may define a *configuration* to be a family of locally-smooth curves that partitions space into a fixed number of regions (where each region may not necessarily be connected). A neighborhood of a given configuration B is the set of all configurations for which the Hausdorff distance of each region to the corresponding region of B is within some threshold.

There are two natural ways of defining the energy of a configuration. If each bubble in the cluster contains a fixed quantity of a compressible gas, then the energy may naturally be defined as the sum of two terms, one for the tension of the bubble walls and the other for the pressure within the bubbles. That is, the energy is $\gamma L + \sum P_i A_i$, where A_i is the area of the i th bubble, P_i is the difference of its pressure (modeled according to the ideal gas law) from ambient, L is the total length of the curves in the cluster, and γ is the surface tension, the constant of proportionality of the Young–Laplace equation. Alternatively, we may restrict the neighborhood of a configuration to maintain a fixed area for each bubble, and use only the first term γL for the energy; this describes a limiting case in which each bubble contains a fixed quantity of incompressible liquid.

With these definitions for neighborhoods and energy, a *stable equilibrium* is a bubble cluster that uniquely attains the minimum energy. A *metastable equilibrium* is a bubble cluster that is not stable, but that has a neighborhood within which it uniquely attains the minimum energy. An *unstable equilibrium* is a bubble cluster for which all neighborhoods contain a cluster with lower energy. And a *neutral equilibrium* is a bubble cluster that is neither metastable nor unstable: every neighborhood contains a different cluster with equal or lower energy, but there exists a neighborhood that does not contain any lower-energy state.

When all four kinds of equilibrium are allowed, Theorem 5 characterizes the graphs of equilibrium bubble clusters. However, disallowing unstable and neutral equilibria changes this characterization, as we now show. To do so, we use Möbius invariance again, to prove that non-3-connected clusters necessarily include cocircular features that are incompatible with stability.

Lemma 14 *Let G be the graph of a planar soap bubble cluster B , having two edges e_1 and e_2 such that $G \setminus \{e_1, e_2\}$ is disconnected. Then in B , e_1 and e_2 form cocircular arcs or collinear segments.*

Proof Suppose for a contradiction that e_1 and e_2 lie on distinct circles or lines. Then, we could find a point q that lies on the circle or line containing e_1 (but does not belong to e_1 itself), is disjoint from the circle or line containing e_2 , and is also disjoint from all other arcs of B . By Lemma 11, a Möbius transformation μ defined by inversion through q would transform B into another bubble cluster $\mu(B)$, in which $\mu(e_1)$ becomes a line segment while $\mu(e_2)$ remains curved. But this contradicts the Young–Laplace equation, because both arcs must separate the same two bubbles and therefore must have the same curvature. This contradiction shows that e_1 and e_2 are cocircular. $\square$

Fig. 18 An unstable soap bubble cluster in which many of the bubbles do not have inscribed circles, showing that it cannot be generated as the Möbius-invariant power diagram of a circle packing

Theorem 6 *Let G be a planar graph that is not 3-connected. Then G cannot be the graph of a stable or metastable bubble cluster, regardless of whether the energy is measured as γL or $\gamma L + \sum P_i A_i$.*

Proof We may assume that G is 3-regular and bridgeless, for otherwise it is not the graph of a bubble cluster regardless of stability. Let B be a bubble cluster representing G . The SPQR tree of G includes an S node with edges that alternate between real and virtual (Lemma 5). The removal of any two of its real edges e_1 and e_2 would split the graph into two separated components, so by Lemma 14 these two edges must be represented in B by arcs of the same circle. Because of this cocircularity, every neighborhood of B contains another isomorphic bubble cluster B' in which one of the two components of $G \setminus \{e_1, e_2\}$ is rotated by some small angle ε with respect to the other around the circle containing e_1 and e_2 . B' has the same total arc length as B , and its bubbles all have the same volume and pressure as in B , so it has the same energy as B . For any neighborhood of B , there will be a sufficiently small but nonzero rotation amount ε such that the rotated cluster belongs to the neighborhood. Therefore, B cannot be the unique minimum energy configuration in any neighborhood: it is at best in a neutral equilibrium, and cannot be stable or metastable. $\square$

For graphs that are 3-connected, we do not know whether our construction always produces bubble clusters that are stable or metastable, or alternatively whether the graphs or stable or metastable bubble clusters form restricted subclasses of the 3-regular 3-connected planar graphs. There exist 3-connected bubble clusters in unstable equilibria [23, 58]; however, the known constructions for these clusters involve long thin bubbles that do not contain a circle tangent to each side of the bubble (Fig. 18), so these unstable clusters cannot be generated by our circle packing method.

6 Conclusions

We have developed a new Möbius-invariant power diagram of circles. Using it, we showed that all planar graphs with maximum degree three (and some planar graphs of

degree four) have planar Lombardi drawings, greatly extending the classes of graphs for which such drawings are known to exist. We have implemented our algorithm for the special case of 3-connected 3-regular graphs. We have also used the same diagrams to prove that two-dimensional soap bubble clusters, obeying Plateau's laws and the Young–Laplace equation, may be characterized graph-theoretically as the planar 3-regular 2-connected graphs, and we have made a preliminary investigation of the stability of the clusters constructed by our method.

The Lombardi drawings constructed by our new algorithm have a natural, organic shape, in which the outer face and all the interior faces are approximately circular. If the maximum number of edges per face of the input graph is bounded, then adjacent circles in the circle packing will have radii whose ratio is also bounded, from which it follows that the vertex spacing of the drawing is locally uniform. Our drawings automatically display any global cyclic or dihedral symmetries of the input graph, and may be used even in cases such as the truncated icosahedron which our previous methods, designed for graphs with a high degree of symmetry [18], are incapable of handling. They also display certain local symmetries of the graph, in the sense that, if G is a subgraph of a 3-connected graph that is connected to the rest of the graph by exactly three edges, then every copy of G in one of our drawings can be transformed into every other copy of G by a Möbius transformation. Possible flaws in this approach, from the graph drawing point of view, are that the area of the drawing (scaled so that the shortest edge has unit length) may be exponential, that the coordinates of the vertices may not be possible to represent precisely (as integers or bounded-degree algebraic numbers) but instead must be approximated numerically, and that the time for computing these drawings depends on the accuracy of this numerical approximation rather than varying strongly polynomially with the input size.

Although our results greatly extend our knowledge of planar Lombardi drawings, other fundamental questions in this area remain unsolved. For instance, what is the complexity of testing whether a planar graph has a Lombardi drawing? It would also be of interest to determine the maximum number d such that all graphs of maximum degree d , or all planar graphs of maximum degree d , have Lombardi drawings that are not necessarily planar.

Much about the combinatorial geometry of soap bubbles still remains to be studied. Most immediately, it would be of interest to fully characterize the graphs of stable or metastable planar soap bubbles; we have shown that they must be 3-connected, but do not know whether that is a sufficient condition. Infinite bubble clusters in the plane with translational symmetry may be viewed as covering spaces of finite clusters on a torus; can these clusters be characterized? This problem is made more difficult than the planar 3-connected case by the inability to use Möbius transformations and the need to consider multigraphs and non-polyhedral embeddings. As a far-off goal, we would eventually like to be able to characterize the cell complexes that can represent soap bubble clusters in three dimensions. However, even much simpler questions in this direction remain unanswered. For instance, is a three-dimensional analogue of Lemma 12 valid, or can stable three-dimensional soap bubble clusters include surfaces that have the same bubble on both sides?

Acknowledgments This work was supported in part by NSF Grants 0830403 and 1217322 and by the Office of Naval Research under grant N00014-08-1-1015. Portions of this work appeared in “Planar Lombardi drawings for subcubic graphs”, 20th International Symposium on Graph Drawing (GD 2012), LNCS 7704, Springer-Verlag, 2013, pp. 126–137, and in “The graphs of planar soap bubbles”, 29th ACM Symp. on Computational Geometry, 2013, pp. 27–36. The author is grateful to Frank Morgan and John Sullivan for helpful comments.

References

1. Aaronson, S.: NP-complete problems and physical reality. *SIGACT News* **36**(1), 30–52 (2005)
2. Angelini, P., Eppstein, D., Frati, F., Kaufmann, M., Lazard, S., Mchedlidze, T., Teillaud, M., Wolff, A.: Universal point sets for planar graph drawings with circular arcs. In: Proceedings of the 25th Canadian Conference on Computational Geometry, pp. 117–122 (2013)
3. Aurenhammer, F.: Power diagrams: properties, algorithms and applications. *SIAM J. Comput.* **16**(1), 78–96 (1987)
4. Bern, M., Eppstein, D.: Optimal Möbius transformations for information visualization and meshing. In: 7th Workshop Algorithms and Data Structures. Lecture Notes in Computer Science, vol. 2125, pp. 14–25. Springer, Berlin (2001)
5. Bern, M., Eppstein, D.: Möbius-invariant natural neighbor interpolation. In: Proceedings of 14th ACM-SIAM Symposium. Discrete Algorithms (SODA’03), pp. 128–129 (2003). <http://portal.acm.org/citation.cfm?id=644130>
6. Bern, M., Graham, R.: The shortest-network problem. *Sci. Am.* **260**, 84–89 (1989)
7. Brightwell, G.R., Scheinerman, E.R.: Representations of planar graphs. *SIAM J. Discrete Math.* **6**(2), 214–229 (1993)
8. Cannon, J.W., Floyd, W.J., Kenyon, R., Parry, W.R.: Hyperbolic geometry. In: Levy, S. (ed.) *Flavors of Geometry*, vol. 31, pp. 59–115. Cambridge University Press, MSRI, Cambridge (1997)
9. Chernobelskiy, R., Cunningham, K., Goodrich, M.T., Kobourov, S.G., Trott, L.: Force-directed Lombardi-style graph drawing. In: van Kreveld, M.J., Speckmann, B. (eds.) *Proceedings of 19th International Symposium on Graph Drawing. Lecture Notes in Computer Science*, vol. 7034, pp. 320–331. Springer, Berlin (2011)
10. Chrobak, M., Goodrich, M.T., Tamassia, R.: Convex drawings of graphs in two and three dimensions (preliminary version). In: *Proceedings of 12th Symposium on Computational Geometry*, pp. 319–328 (1996).
11. Collins, C.R., Stephenson, K.: A circle packing algorithm. *Comput. Geom. Theory Appl.* **25**(3), 233–256 (2003)
12. Cox, S.J., Morgan, F., Graner, F.: Are large two-dimensional clusters of perimeter-minimizing bubbles of equal-area hexagonal or circular? (2012). Preprint
13. Di Battista, G., Tamassia, R.: Incremental planarity testing. In: 30th Symposium on Foundations of Computer Science, pp. 436–441. IEEE (1989).
14. Di Battista, G., Tamassia, R.: On-line graph algorithms with SPQR-trees. In: 17th International Colloquium Automata, Languages and Programming. *Lecture Notes in Computer Science*, vol. 443, pp. 598–611. Springer, Berlin (1990).
15. Došlić, T.: On some structural properties of fullerene graphs. *J. Math. Chem.* **31**(2), 187–195 (2002)
16. Duncan, C.A., Eppstein, D., Goodrich, M.T., Kobourov, S.G., Löffler, M.: Planar and poly-arc Lombardi drawings. In: M.J. van Kreveld, B. Speckmann (eds.) *Proceedings of 19th International Symposium on Graph Drawing. Lecture Notes in Computer Science*, vol. 7034, pp. 308–319. Springer, Berlin (2011).
17. Duncan, C.A., Eppstein, D., Goodrich, M.T., Kobourov, S.G., Nöllenburg, M.: Drawing trees with perfect angular resolution and polynomial area. In: U. Brandes, S. Cornelsen (eds.) *Proceedings of 18th International Symposium on Graph Drawing. Lecture Notes in Computer Science*, vol. 6502, pp. 183–194. Springer, Berlin (2011).
18. Duncan, C.A., Eppstein, D., Goodrich, M.T., Kobourov, S.G., Nöllenburg, M.: Lombardi drawings of graphs. *J. Graph Algorithms Appl.* **16**(1), 85–108 (2012)
19. Dutta, P., Khastgir, S.P., Roy, A.: Steiner trees and spanning trees in six-pin soap films. *Am. J. Phys.* **78**(2), 215–221 (2010)

20. Eppstein, D.: Quasiconvex programming. In: Goodman, J.E., Pach, J., Welzl, E. (eds.) *Combinatorial and Computational Geometry*, vol. 52, pp. 287–331. Cambridge University Press, MSRI, Cambridge (2005)
21. Eppstein, D., Dillencourt, M.B.: Uninscribable 4-regular polyhedron. *Electronic Geometry Models* p. 2003.08.001 (2003). <http://www.eg-models.de/2003.08.001/>
22. Eppstein, D., Mumford, E.: Steinitz theorems for orthogonal polyhedra. In: *26th Symposium on Computational Geometry*, pp. 429–438. ACM (2010).
23. Fortes, M.A., Vaz, M.F., Cox, S.J., Teixeira, P.I.C.: Instabilities in two-dimensional flower and chain clusters of bubbles. *Colloids Surf. A* **309**(1–3), 64–70 (2007)
24. Frucht, R.: Herstellung von Graphen mit vorgegebener abstrakter Gruppe. *Compos. Math.* **6**, 239–250 (1939).
25. Grinberg, È.J.: Plane homogeneous graphs of degree three without Hamiltonian circuits. In: *Latvian Math. Yearbook* 4, pp. 51–58. Izdat. “Zinatne”, Riga (1968)
26. Gutwenger, C., Mutzel, P.: A linear time implementation of SPQR-trees. In: *8th International Symposium on Graph Drawing. Lecture Notes in Computer Science*, vol. 1984, pp. 77–90. Springer, Berlin (2001)
27. Hales, T.C.: Cannonballs and honeycombs. *Notices Am. Math. Soc.* **47**(4), 440–449 (2000)
28. Hales, T.C.: The honeycomb conjecture. *Discrete Comput. Geom.* **25**(1), 1–22 (2001)
29. Halin, R.: Studies on minimally n -connected graphs. In: *Combinatorial Mathematics and Its Applications (Proc. Conf., Oxford, 1969)*, pp. 129–136. Academic Press, London (1971)
30. Hass, J., Schlafly, R.: Double bubbles minimize. *Ann. Math.* **151**(2), 459–515 (2000)
31. Herdtle, T., Aref, H.: On the geometry of composite bubbles. *Proc. Math. Phys. Sci.* **434**(1891), 441–447 (1991).
32. Hoffman, D.T.: Smart soap bubbles can do calculus. *Math. Teach.* **72**(5), 377–385 (1979).
33. Hopcroft, J., Tarjan, R.: Dividing a graph into triconnected components. *SIAM J. Comput.* **2**(3), 135–158 (1973)
34. Hutchings, M., Morgan, F., Ritoré, M., Ros, A.: Proof of the double bubble conjecture. *Ann. Math.* **155**(2), 459–489 (2002)
35. Isenberg, C.: The soap film: an analogue computer. *Am. Sci.* **64**(5), 514–518 (1976).
36. Isenberg, C.: Chapter 5. The Laplace–Young Equation. In: *The Science of Soap Films and Soap Bubbles*. Dover, New York (1978)
37. Kimberling, C.: Triangle Centers and Central Triangles, vol. 129. *Utilitas Math*, Winnipeg (1998). <http://faculty.evansville.edu/ck6/tcenters/tcct.html>
38. Koźmiński, K., Kinnen, E.: An algorithm for finding a rectangular dual of a planar graph for use in area planning for VLSI integrated circuits. In: *Proceedings of 21st Conference on Design Automation (DAC’84)*, pp. 655–656 (1984).
39. Leinwand, S.M., Lai, Y.T.: An algorithm for building rectangular floor-plans. In: *Proceedings of 21st Conference on Design Automation (DAC’84)*, pp. 663–664 (1984).
40. Löffler, M., Nöllenburg, M.: Planar Lombardi drawings of outerpaths. In: *20th International Symposium, GD 2012, Redmond, WA, September 19–21, 2012, Revised Selected Papers. Lecture Notes in Computer Science*, vol. 7704, pp. 561–562. Springer, Berlin (2013).
41. Mac Lane, S.: A structural characterization of planar combinatorial graphs. *Duke Math. J.* **3**(3), 460–472 (1937)
42. Markström, K.: Extremal graphs for some problems on cycles in graphs. *Congr. Numer.* **171**, 179–192 (2004)
43. Mohar, B.: A polynomial time circle packing algorithm. *Discrete Math.* **117**(1–3), 257–263 (1993)
44. Morgan, F.: *Geometric Measure Theory: A Beginner’s Guide*, 4th edn. Academic, Boston (2009)
45. Plateau, J.A.F.: *Statique Experimentale et Theorie des Liquides Soumis aux Seules Forces Moleculaires*. Gauthier-Villiard, Paris (1873)
46. Pootheri, S.K.: Decomposition characterizations of classes of 2-connected graphs. In: *39th Southeast Conference ACM* (2001)
47. Quinn, C.: Least-perimeter partitions of the sphere. *Rose–Hulman Inst. Tech. Undergrad. Math. J.* **8**(2) (2007). <http://www.rose-hulman.edu/mathjournal/archives/2007/vol8-n2/paper11/v8n2-11pd.pdf>. Appendix by Quinn Maurmann
48. Ribó Mor, A., Rote, G., Schulz, A.: Small grid embeddings of 3-polytopes. *Discrete Comput. Geom.* **45**(1), 65–87 (2011)

49. Schulz, A.: Drawing 3-polytopes with good vertex resolution. *J. Graph Algorithms Appl.* **15**(1), 33–52 (2011)
50. Schwerdtfeger, H.: *Geometry of Complex Numbers: Circle Geometry, Moebius Transformation, Non-Euclidean Geometry*. Dover, New York (1979)
51. Steinitz, E.: Polyeder und Raumeinteilungen. In: *Encyclopädie der mathematischen Wissenschaften, Band 3 (Geometries)*, pp. 1–139. B.G. Teubner, Leipzig (1922)
52. Stephenson, K.: *Introduction to Circle Packing: The Theory of Discrete Analytic Functions*. Cambridge University Press, Cambridge (2005)
53. Sullivan, J.M.: The geometry of bubbles and foams. In: Sadoc, J.F., Rivier, N. (eds.) *Foams and Emulsions: Proceedings of NATO Advanced Study Institute on Foams and Emulsions, Emulsions and Cellular Materials, Cargèse, Corsica, 12–24 May, 1997*, NATO Adv. Sci. Inst. Ser. E vol. 354, pp. 379–402. Kluwer, Dordrecht (1999)
54. Taylor, J.E.: The structure of singularities in soap-bubble-like and soap-film-like minimal surfaces. *Ann. Math.* **103**(3), 489–539 (1976)
55. Thompson, W.: On the division of space with minimum partitional area. *Philos. Mag.* **24**(151), 503 (1887)
56. Tutte, W.T.: On Hamiltonian circuits. *J. Lond. Math. Soc.* **21**(2), 98–101 (1946)
57. Ungar, P.: On diagrams representing maps. *J. Lond. Math. Soc.* **28**, 336–342 (1953)
58. Weaire, D., Cox, S.J., Graner, F.: Uniqueness, stability and Hessian eigenvalues for two-dimensional bubble clusters. *Eur. Phys. J. E* **7**(2), 123–127 (2002)
59. Weaire, D., Phelan, R.: A counter-example to Kelvin’s conjecture on minimal surfaces. *Philos. Mag. Lett.* **69**, 107–110 (1994)
60. Wichiramala, W.: *The Planar Triple Bubble Problem*. Ph.D. thesis, University of Illinois at Urbana-Champaign (2002). <http://www.math.sc.chula.ac.th/~wacharin/old/math/thesisss.pdf>