



Packing Disks by Flipping and Flowing

Robert Connelly¹ · Steven J. Gortler²

Received: 11 November 2019 / Revised: 27 May 2020 / Accepted: 4 August 2020
© Springer Science+Business Media, LLC, part of Springer Nature 2020

Abstract

We provide a new type of proof for the Koebe–Andreev–Thurston (KAT) planar circle packing theorem based on combinatorial edge-flips. In particular, we show that starting from a disk packing with a maximal planar contact graph G , one can remove any flippable edge e^- of this graph and then continuously flow the disks in the plane, so that at the end of the flow, one obtains a new disk packing whose contact graph is the graph resulting from flipping the edge e^- in G . This flow is parameterized by a single inversive distance.

Keywords Circle packing · Rigidity · Inversive distance

Mathematics Subject Classification 52C26

1 Introduction

The well known Koebe–Andreev–Thurston (KAT) (planar) circle packing theorem states that for every planar graph G with n vertices, there is a corresponding packing of n disks (with mutually disjoint interiors) in the plane, whose contact graph is isomorphic to G . Moreover, if G is maximal (i.e., all faces triangular), then this packing is unique up to Möbius transformations and reflections [2,19,33]. This theorem is important in conformal geometry [27] and has been generalized in numerous directions

Editor in Charge: János Pach

R. Connelly: Partially supported by NSF Grant DMS-1564493. S. J. Gortler: Partially supported by NSF Grant DMS-1564473.

Robert Connelly
rc46@cornell.edu

Steven J. Gortler
sjg@cs.harvard.edu

¹ Department of Mathematics, Cornell University, Ithaca, USA

² School of Engineering and Applied Sciences, Harvard University, Cambridge, USA

[26,30,33]. In [14], a packing whose graph is a triangulation of the plane is called a *compact packing*.

There are a variety of techniques that have been used to establish the KAT theorem. These include methods based on conformal geometry [19], combinatorial analysis of hyperbolic polyhedra [2,28], circle geometry, cone-singularities and topology of continuous maps [24,33], variational methods [3,9,26], iterative "flattening" algorithms [10,25,32], and combinatorial Ricci flow [8].

In this paper, we present a novel simple and constructive approach for proving the KAT circle packing theorem based on ideas from (local) rigidity theory [6,13].

We start with any "trilaterated" packing; such a packing is constructed by starting with three fixed disks in mutual tangential contact, and then adding in, one-by-one, new disks that are tangential to three previously placed disks. It will be easy to establish that such a packing is unique up to Möbius transformation and reflection. (See Fig. 3.)

Next, we use the fact that one can always combinatorially transform any maximal planar graph to any desired target maximal planar graph G using a finite sequence of combinatorial "edge flips" [4,31,34]. (See Fig. 4.)

Then, at the heart of this paper, we show how we can continuously flow the planar packing in coordination with a single edge flip. More specifically, let $P(0)$ be a packing with a maximal planar contact graph G . Let H be a maximal planar graph obtained from G by removing one edge e^- and adding its "cross" edge e^+ (see Fig. 4). Also let G^- be the intermediate graph $G \setminus e^-$ such that G^- has one quadrilateral face. Then we can always find a continuous path of planar packings $P(t)$ for $t \in [0, 1]$ with the following properties:

- The contact graph of $P(0)$ is G .
- For t in the interval $(0, 1)$, the contact graph of $P(t)$ is G^- .
- The contact graph of $P(1)$ is H .

This means that we can continuously push the two disks corresponding to e^- apart, while maintaining a packing (no interior overlaps) and maintaining all of the other contacts in G^- , and we can keep doing this until the two disks across e^+ come together (see Fig. 1), giving us H as the contact graph. Moreover, up to the speed of the parameterization and up to Möbius transformations and reflections, this path is unique and is parameterized by a single inversive distance. The construction of this path is our main technical result, Theorem 9.9.

Together these ingredients can be used to establish the KAT circle packing theorem. Our construction requires the inversion of a certain Jacobian matrix at each time step of the flow, so we do not necessarily expect it to be faster than other practical circle packing techniques such as [10]. Interestingly, in our approach we always work with disks in the plane and never introduce or need to deal with any cone singularities in either the theory or the construction. Thus it provides a novel view of the circle packing problem.

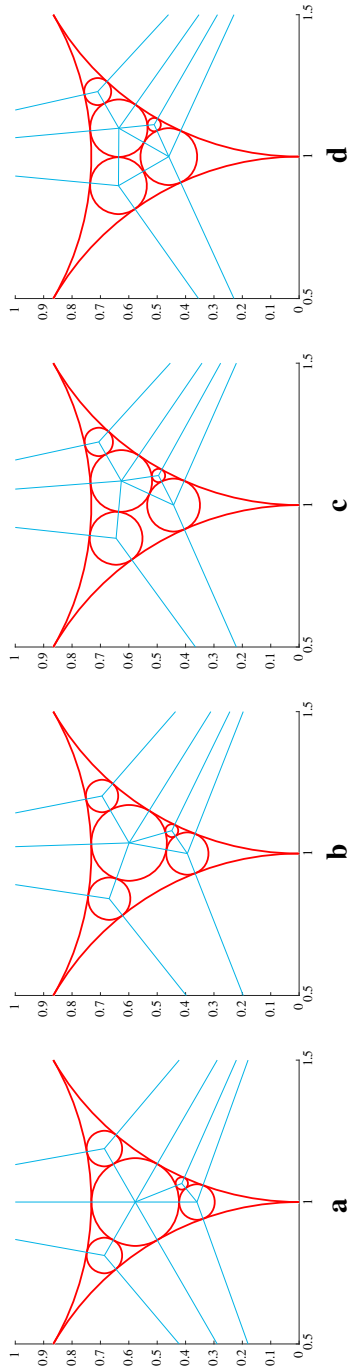


Fig. 1 **a** A trilaterated packing. We are about to flip the edge e^- connecting the central disk to the boundary disk on the lower left. **b**, **c** tridisk-contained almost triangulated packings (one quadrilateral) during the flow. **d** At the end of the flow, the cross-edge e^+ has made contact giving us a triangulated packing with a new graph

2 Preliminaries

Definition 2.1 A (planar) disk configuration (\mathbf{p}, \mathbf{r}) is a sequence (i.e., an indexed family) of n disks placed in \mathbb{R}^2 , with centers at $\mathbf{p} := (\mathbf{p}_1, \dots, \mathbf{p}_n)$ and with positive radii $\mathbf{r} := (r_1, \dots, r_n)$. A (planar) packing (\mathbf{p}, \mathbf{r}) is a disk configuration such that their interiors are mutually disjoint. The contact graph of a packing is the graph (over the same index set) that has one vertex corresponding to each disk and one edge corresponding to each pair of mutually tangent disks.

Definition 2.2 Let G be any graph with n vertices, and let a point configuration $\mathbf{p} := (\mathbf{p}_1, \dots, \mathbf{p}_n)$ be a sequence of n points in the plane. Define (G, \mathbf{p}) to be the straight line drawing of the graph G with vertex positions determined by \mathbf{p} .

Lemma 2.3 *Let G be the contact graph of some packing (\mathbf{p}, \mathbf{r}) . Then G must be planar and (G, \mathbf{p}) is a planar embedding.*

Proof This follows from the following two facts: Each of the line segments in (G, \mathbf{p}) must be contained within its associated two touching disks. None of the disks have overlapping interiors. \square

Definition 2.4 A Möbius transformation of the extended complex plane is a function of the form

$$\phi(z) := \frac{az + b}{cz + d},$$

where $ad - bc \neq 0$.

Möbius transformations are injective, orientation preserving. We can apply a Möbius transformation to \mathbb{R}^2 by identifying it with \mathbb{C} (the transformation may map one point to infinity, and infinity to another point). The generalized Möbius transformations are generated by the Möbius transformations and reflections through lines. An inversion about a circle, is a transformation where a point is sent to the point “one over its distance” from the center of the circle, times the square of the radius of the circle on the ray through the center of the circle and the point. It can be shown that a circle inversion is a generalized Möbius transformation.

A generalized Möbius transformation ϕ maps an input circle or a line to an output circle or a line. If a circle C_i maps to a circle C_o with the same orientation under ϕ , then ϕ will map the interior of C_i (i.e., a disk) to the interior of C_o (i.e., a disk). If a circle C_i maps to a circle C_o with the opposite orientation under ϕ , then ϕ will map the interior of C_i (i.e., a disk) to the exterior of C_o (i.e., the complement of a closed disk). From its injectivity, we see that a Möbius transformation maps a packing to a packing (but up to one of the disks may become a half plane or the complement of a disk). The following theorem is standard:

Theorem 2.5 *There is a unique Möbius transformation mapping any sequence of three distinct “input” points to any other sequence of three distinct “output” points. This map varies smoothly as we alter the desired output point locations.*




Fig. 2 A set of three mutually externally touching circles, with red center points, have three contact points in blue. These contact points lie at a circle that is called the incircle (dashed) of the triangle determined by the centers of the circles. The blue contact points are not contained in any half of this incircle. Going the other way, given any three points that are not contained in any half of their incident circle, we can associate a unique set of three mutually externally touching circles

Definition 2.6 The inversive distance between two disks, (p_i, r_i) and (p_j, r_j) , is defined as

$$\frac{\|p_i - p_j\|^2 - r_i^2 - r_j^2}{2r_i r_j}.$$

It evaluates to 1 when the disks have external tangential contact.

Inversive distances play well with Möbius transformations as the following standard theorem states:

Theorem 2.7 *The inversive distance between two disks is invariant for generalized Möbius transformations.*

3 Tridisk-Contained Packings

In this section, we will use Theorem 2.5 and a standard construction (e.g., [24]) to mod out generalized Möbius transformations. The basic idea is to use a transformation to always put three chosen disks into a canonical form. We go through the details somewhat carefully.

Definition 3.1 A tridisk is a sequence of three non-overlapping disks in mutual tangential contact. (See Fig. 2.) The canonical tridisk is the tridisk made of three disks of unit radius, where the disk centers lie at three points of a fixed equilateral triangle. (And to make this unique, we fix an indexing of these three disks.) (See Fig. 3.) The incircle of a tridisk is the circle touching the three contact points of the tridisk. (See Fig. 2.)

Lemma 3.2 *Given an input tridisk, there is a unique Möbius transformation mapping it to the canonical tridisk.*

Proof The input tridisk has a sequence of three contact points. There is a unique Möbius transformation ϕ mapping these to the sequence of three contact points of

the canonical tridisk (Theorem 2.5). A sequence of three externally touching circles are uniquely determined by such a sequence of points (as in Fig. 2). So ϕ must map the three input touching circles to the canonical three circles (with external tangential contact). Now from injectivity, ϕ must map the interiors of the three distinct input disks to three disjoint regions. This disjointness ensures that ϕ must map the interiors of the three input circles to the interiors of the three canonical circles. \square

Note that the above transformation may map the interior of the tridisk's incircle to either the interior or the exterior of the canonical tridisk's incircle. This depends on the relative orientations of the three contact points.

Corollary 3.3 *Given an input tridisk, let us select either the interior or the exterior of the tridisk's incircle. Then there is a unique **generalized** Möbius transformation mapping the tridisk to the canonical tridisk and which maps the selected side of the tridisk's incircle to the interior of the canonical tridisk's incircle.*

Proof We start with the Möbius transformation of Lemma 3.2. If it has mapped the wrong side of the input incircle to the output incircle's exterior we follow this by inverting through the canonical tridisk incircle. This must preserve the contact points, and thus must leave each of the three circles invariant. \square

The following will be needed later.

Corollary 3.4 *Given an input tridisk, with centers $(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3)$ and radii (r_1, r_2, r_3) , and given any fixed displacement vector $\mathbf{v} \in \mathbb{R}^2$, then for sufficiently small t , there is a unique Möbius transformation f_t that maps this tridisk to an output tridisk that has centers at $(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3 + t\mathbf{v})$ (the radii can change). Moreover, this transformation is smooth in t .*

Proof Given the desired three (non-colinear) centers $(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3 + t\mathbf{v})$, there is always a unique set of radii (r'_1, r'_2, r'_3) , such that this creates an output tridisk. This output tridisk has a sequence of three contact points, giving us the unique Möbius transformation of Lemma 3.2. All of the above steps are smooth. \square


Definition 3.5 Given a packing P with contact graph G that includes some triangular face T , we can mark the packing by assigning, in some order, the labels 1, 2, and 3, to the three disks corresponding to the vertices of T . This specifies a tridisk in P . We refer to the other disks as unmarked.

Definition 3.6 A marked packing P , where the mark-specified tridisk forms the canonical tridisk, and all of the unmarked disks lie inside of its "tricuspid" region is called canonical-tridisk-contained. We will shorten this to simply tridisk-contained.

The next two lemmas tell us that we can use tridisk-containment as an effective modding process.

Lemma 3.7 *Let P be a packing of $n \geq 4$ disks with a 3-connected contact graph G that includes some triangular face T . Let P be marked, specifying a tridisk. Then there is a unique generalized Möbius transformation that maps P to a tridisk-contained packing.*

Fig. 3 A trilaterated packing of five disks. The outer three (yellow) disks form the canonical tridisk



Proof From Lemma 2.3, (G, p) is a planar embedding. Since the tridisk corresponds to a triangular face of G , then in a packing, either all of the other disks lie inside this triangle, or they all lie outside of it. The lemma then follows using Corollary 3.3. \square

Lemma 3.8 *Let P be a tridisk-contained marked packing and with contact graph G . If P is the only packing with contact graph G , up to a generalized Möbius transformation, then P is the only tridisk-contained packing with this marking and with contact graph G . The converse holds as well.*

Proof The case with $n = 3$ is clear, so we will now assume that $n \geq 4$. Suppose that P is the only packing with contact graph G , up to a generalized Möbius transformation. Then, from uniqueness of the transform in Lemma 3.7, it is the only tridisk-contained packing with this contact graph and marking.

For the other direction, suppose there was a packing Q unrelated to P by a generalized Möbius transformation, but with the same contact graph, G . Let us take the marking of P , and use this to mark Q . Then from Lemma 3.7, there is a generalized Möbius transformation taking Q to some tridisk-contained marked packing Q' . By the assumption of no-Möbius-relationship, Q' cannot equal P , but both are tridisk-contained with the same contact graph and marking. This contradicts the assumed marked uniqueness of P . \square

4 Trilaterated Packings

Next we describe a simple family of packings where we can easily argue existence and uniqueness.

Definition 4.1 To create a trilaterated packing with $n \geq 3$ vertices we start with the canonical tridisk. We then apply $n - 3$ subdivision steps. Each subdivision step adds one packing disk, externally tangent to exactly three already-existing disks. We can create such a packing for any $n \geq 3$. (See Fig. 3.) Such a packing is naturally marked by the canonical tridisk, and is tridisk-contained.

Lemma 4.2 *Let P be a trilaterated packing and C its contact graph. Then P is the unique packing with contact graph C , up to a generalized Möbius transformation.*




Fig. 4 An edge flip

Proof Using Lemma 3.8, all we need to show is that P is the only tridisk-contained packing with this marking and with contact graph C .

The proof then follows by induction starting with the unique canonical tridisk. In particular, assume that $n \geq 4$ and that, by induction, the lemma holds for all trilaterated packings with $n - 1$ disks. A trilaterated packing on n disks is obtained from a trilaterated packing on $n - 1$ disks by adding in one final disk. To result in a packing, the final disk must be placed in the tricusp of its three neighbors (which are fixed by induction). (If the new disk center were placed outside of the tricusp of its three contacting circles, but inside of the canonical tricusp, then this would create a non-embedding for the straight-line drawing, (G, \mathbf{p}) , violating Lemma 2.3.) There is only one center and radius for this last disk that creates the three required contacts, and is inside the tricusp of its three neighbors. \square

5 Edge Flips

Definition 5.1 A maximal planar graph is a planar graph such that adding any extra edge on the same vertex set results in a non-planar graph.


A 3-connected planar graph with $n \geq 3$ vertices has a well-defined set of faces. A maximal planar graph is always 3-connected, and each of its faces is a triangle.

Definition 5.2 Let G be a maximal planar graph with $n \geq 5$ vertices. Let e^- be an edge of G connecting two vertices, denoted by a and c . This edge bounds two triangles in G . Denote by b and d the other two vertices of these two triangles. We say that e^- is flippable if there is no edge in G connecting b and d . (See Fig. 4.)

Lemma 5.3 Let G be a maximal planar graph and e^- a flippable edge of G . Then $G^- := G \setminus e^-$ is 3-connected.

Proof The graph G^- is a triangulation of the quadrilateral on vertices a, b, c, d . Such a triangulation must be 3-connected unless there is an edge connecting two non-

Fig. 5 All the inequalities of Lemma 6.1 are equalities here. Note that the distance from p_0 to the circumference of each of the disks of radius r is $(2/\sqrt{3} - 1)r$



consecutive vertices of the quadrilateral (e.g. [21]). But we have removed the edge connecting a and c , and flippability presumes no edge between b and d . \square

Definition 5.4 Let G be a maximal planar graph and let e^- be a flippable edge. Then we call the graph $G^- := G \setminus e^-$ an almost-maximal planar graph. All of its faces are triangles, except for a single quadrilateral, with vertices a, b, c , and d .

Definition 5.5 Let G^- be an almost maximal planar graph as above. Let H be the graph obtained from G^- by adding the edge e^+ connecting b and d . Then we say that H was obtained from G by an edge flip on e^- .

If H is obtained from G by an edge flip, then H is a maximal planar graph. The following theorem [31,34] lets us use edge flips to navigate between two graphs (and this can be done while preserving the vertex indices).

Theorem 5.6 We can transform any maximal planar graph on n vertices to any other maximal planar graph on n vertices through a finite sequence of edge flips.

Definition 5.7 Let G be a maximal planar graph. We call any packing (p, r) with contact graph G , a triangulated packing. Let G^- be an almost-maximal planar graph. We call any packing (p, r) with contact graph G^- , an almost-triangulated packing.


Given an almost-maximal packing with contact graph G^- , we can mark the three disks corresponding to any triangular face of G^- and transform the packing to obtain a tridisk-contained, almost-maximal marked packing.

6 Bounding Radii

Next we will argue that in a tridisk-contained, almost-triangulated packing, all of our radii are bounded away from 0. We start with

Lemma 6.1 There is no configuration of four points, p_0, p_1, p_2, p_3 , in the Euclidean plane and real number $r > 0$ with all the following properties:

Fig. 6 Here the point p_0 is within ϵ of each of three disks all of radius at least r , the radius of the smallest disk. The other disks have their radii shrunk to r , while keeping the nearest points to p_0 the same. The p_i are the centers of the shrunken disks



- For i and j in $\{1, 2, 3\}$ with $i \neq j$, we have $|p_i - p_j| \geq 2r$.
- For $i \in \{1, 2, 3\}$, we have $|p_i - p_0| \leq (2/\sqrt{3})r$.
- At least one of the above six inequalities is strict.

The proof is by constructing a positive semi-definite (stress-energy) function whose global minimum occurs only when the above inequalities hold as equalities. The energy is constructed so that for any configuration where the inequalities are satisfied, the energy value is at least as small as this minimum value. See [11, 12] for a discussion for this easy case and many other similar ones. See Fig. 5.

Lemma 6.2 *Let there be three packing disks, and a point p_0 , whose distance to the boundary of each of the disks is less than some ϵ . Then at least one of the disks has radius $\leq \epsilon / (2/\sqrt{3} - 1) = \epsilon \cdot 6.46\dots < 7\epsilon$.*

Proof Suppose that all of the disks had radii $> \epsilon / (2/\sqrt{3} - 1)$. Then we could shrink the radii of the two largest disks to the radius r of the smallest disk, as in Fig. 6 while maintaining a packing. Let us denote these new disk centers as p_1, p_2 , and p_3 . For i and j in $\{1, 2, 3\}$ with $i \neq j$, we would have $|p_i - p_j| \geq 2r$. Meanwhile we would have, for each $i \in \{1, 2, 3\}$, $|p_0 - p_i| < \epsilon + r \leq (2/\sqrt{3} - 1)r + r = 2r/\sqrt{3}$. This would contradict Lemma 6.1. \square

And now we can prove our proposition:

Proposition 6.3 *Let G be a fixed 3-connected graph on n vertices (such as an almost-maximal planar graph). Let (p, r) be a packing with contact graph G . Suppose that at least three of the disks have radii ≥ 1 . Then, all of the radii are bounded away from 0 (where the bound depends only on n).*

Proof Define M_ϵ to be a maximal connected set of vertices where all radii are $< \epsilon$ for any chosen ϵ . Claim: If $\epsilon < 1$, M_ϵ is non-empty, and $M_{7n\epsilon}$ is a component containing M_ϵ , then $|M_{7n\epsilon}| \geq 1 + |M_\epsilon|$. Proof: Since the disks in M_ϵ are connected, they fit inside of some circle C of radius $< n\epsilon$. From the assumed three large disks, we know that $|M_\epsilon| \leq n - 3$ (i.e., this set is a strict subset of (p, r)). Now we look at the neighbors of our set M_ϵ in G . From 3-connectivity, the set must have at least three neighbors. From Lemma 6.2, at least one of these disks has radius less than $7n\epsilon$, establishing our claim.

Now, let $\epsilon^* := 1/(7n)^n$. If there is any disk with radius less than ϵ^* , then, iteratively applying the previous paragraph, we find that all of disks have radii < 1 , contradicting our assumptions. Finally, an almost maximal planar graph is 3-connected from Lemma 5.3. □

Remark 6.4 If we just know that there is at least one large disk, and we know that it is in the interior of a 3-connected packing, then we can use the same argument to prove that no other disk can be arbitrarily small. (For this we need to argue that no set of $n - 1$ disks, with $n - 2$ of them sufficiently tiny, can wrap around the known large disk.) This then gives us a generalization of Rodin and Sullivan’s “ring-lemma” [27]. (We only assume 3-connectivity, and there may not even be a single “ring” in the packing.)

Lemma 6.5 *Let (\mathbf{p}, \mathbf{r}) be a tridisk-contained, almost-triangulated packing with contact graph G^- . Let ij be a vertex pair that is not an edge in G^- , and is neither e^- nor e^+ . Then disks i and j must remain a bounded distance from each other by a constant depending only on the number of vertices.*

Proof The graph $G' := G^- \cup e_{ij}$ is an almost-maximal planar graph plus one extra edge. But since i and j are not on some common face of G^- , the edge between them does not subdivide any face of G^- and so G' cannot be planar. (Fixing a (topologically unique) embedding of the 3-connected graph G' on the sphere, any curve starting at i and ending at j must enter and then exit one of the faces incident to i . The crossing at the exit point certifies non-planarity.)

From Lemma 2.3, (G^-, \mathbf{p}) is a planar embedding. Let us add in the segment connecting \mathbf{p}_i and \mathbf{p}_j . From non-planarity, this must cross some (contact) edge kl of (G^-, \mathbf{p}) . The radii r_i, r_j, r_k, r_l are all bounded away from zero from Proposition 6.3. This keeps disks k and l bounded away from contact. □

7 Infinitesimal Rigidity

In this section we prove an infinitesimal rigidity result about inversive distances on almost-triangulated packings. This, along with Proposition 7.2, will be used to prove an existence and global rigidity result.

Definition 7.1 Let G be a maximal planar graph on n vertices. Let f be the mapping from \mathbb{R}^{3n} to \mathbb{R}^{3n-6} that measures the inversive distance along the edges of G of a disk configuration. (We do not require that these disks have disjoint interiors, but we do require the radii to be positive.) Explicitly we have

$$f(\mathbf{p}, \mathbf{r})_{ij} := \frac{\|\mathbf{p}_i - \mathbf{p}_j\|^2 - r_i^2 - r_j^2}{2r_i r_j}. \tag{1}$$

Let J (at (\mathbf{p}, \mathbf{r})) be the Jacobian of f . Its row corresponding to an edge ij of G has the following pattern [6]:

$$\left(\dots \quad \mathbf{p}_i \quad \dots \quad \mathbf{p}_j \quad \dots \quad r_i \quad \dots \quad r_j \right. \\ \left. \left(\dots 0 \dots \frac{\mathbf{p}_i - \mathbf{p}_j}{r_i r_j} \dots 0 \dots \frac{\mathbf{p}_j - \mathbf{p}_i}{r_i r_j} \dots 0 \dots \frac{r_j^2 - r_i^2 - \|\mathbf{p}_i - \mathbf{p}_j\|^2}{2r_i^2 r_j} \dots 0 \dots \frac{r_i^2 - r_j^2 - \|\mathbf{p}_i - \mathbf{p}_j\|^2}{2r_j^2 r_i} \right) \right)$$

where the first row above labels the matrix columns. When disks i and j have no interior overlap (i.e., when \mathbf{p}_i is not inside of disk j), we can verify that

$$\frac{r_j^2 - r_i^2 - \|\mathbf{p}_i - \mathbf{p}_j\|^2}{2r_i^2 r_j} \leq 0. \tag{2}$$

When disks i and j are in tangential contact, their inversive distance is 1 and we have $\|\mathbf{p}_i - \mathbf{p}_j\| = r_i + r_j$, giving us

$$\frac{r_j^2 - r_i^2 - \|\mathbf{p}_i - \mathbf{p}_j\|^2}{2r_i^2 r_j} = \frac{-r_i - r_j}{r_i r_j} = \frac{-\|\mathbf{p}_i - \mathbf{p}_j\|}{r_i r_j}. \tag{3}$$

In this section, we will establish the following:

Proposition 7.2 *Let G and H be maximal planar graphs related by flipping an edge e^- . Let (\mathbf{p}, \mathbf{r}) be a (triangulated or almost-triangulated) packing with contact graph of either G , G^- , or H . Then the Jacobian, J , (defined using the edges of G) has rank $3n - 6$.*

Our argument in this section is a variation on ideas explored in [6,13]. The idea is to show that J cannot have any non-zero cokernel vector. The main extra difference here is that when the contact graph is G^- or H , then the edges of G are not in contact. We will argue that, since G^- and H are similar enough to G , this will not destroy the argument.

Definition 7.3 Given a 3-connected planar graph G with n vertices and m edges, and given a vector \mathbf{v} in \mathbb{R}^m , we assign a sign $\{+, 0, -\}$ to each undirected edge ij using the sign of \mathbf{v}_{ij} . This assignment gives us an associated sign vector \mathbf{s} .

Definition 7.4 Given a sign vector \mathbf{s} , we define the index I_i at vertex i , as the number of times the sign changes as we traverse the edges in order around vertex i , ignoring zeros. (We use the unique face structure of a 3-connected planar graph to define the cyclic ordering of edges at each vertex.) The index I_i is always even.

The following is a variation on Cauchy’s index lemma, which can be proven using Euler’s formula. For a proof, see e.g., [15, Lem. 5.2] or [1, p.87].

Lemma 7.5 *Let G be a planar 3-connected graph with n vertices. Let \mathbf{s} be a sign vector in \mathbb{R}^m . Let n' be the number of vertices that have at least one non-zero signed edge. Then $\sum_i I_i \leq 4n' - 8$.*

Lemma 7.6 *Let G be a maximal planar graph on n vertices. Let (\mathbf{p}, \mathbf{r}) be a configuration of n disks. Let \mathbf{v} be a non-zero cokernel vector of J , giving us an associated sign vector. Let i be the index of a disk that has no interior overlap with its neighbors in G and such that it has at least one non-zero signed edge in \mathbf{v} . Then its associated index I_i is at least 2.*

Proof If vertex i has at least one non-zero sign and no sign changes, then it cannot satisfy equations (2) while having \mathbf{v} annihilates the column of J corresponding to r_i . □

Definition 7.7 Let G be any 3-connected planar graph on n vertices and \mathbf{p} a configuration of n points. Let (G, \mathbf{p}) be its straight line drawing (not necessarily an embedding). We say that a vertex has out of order edges if the drawn (either clockwise or counter-clockwise) cyclic order of its adjacent edges does not match their combinatorial cyclic order in G .

Definition 7.8 Let G be a maximal planar graph on n vertices. Let (\mathbf{p}, \mathbf{r}) be a configuration of n disks. We call a disk with index i , G -like if the edges of vertex i in (G, \mathbf{p}) are in order and if for all of the edges ij in G incident to vertex i , the disks i and j are in tangential contact in (\mathbf{p}, \mathbf{r}) .

Lemma 7.9 *Let G be a maximal planar graph. Let (\mathbf{p}, \mathbf{r}) be a triangulated or almost-triangulated packing with contact graph $G, G^-,$ or H . Then there are at most two disks that are not G -like.*

Proof The contact graph G' of (\mathbf{p}, \mathbf{r}) is one of $G, G^-,$ or H . So Lemma 2.3 tells us that (G', \mathbf{p}) is embedded. G^- is a subgraph of G' so (G^-, \mathbf{p}) is also embedded. It is 3-connected by Lemma 5.3. Thus all the edges of (G^-, \mathbf{p}) are drawn in order and have disks with tangential contact. The graph G includes only one more edge, which can affect only two disks. □

Definition 7.10 Given a graph G with n vertices and m edges, and a configuration \mathbf{p} of n points. A vector $\omega \in R^m$ satisfies the equilibrium condition at vertex i if

$$\sum_j \omega_{ij}(\mathbf{p}_i - \mathbf{p}_j) = 0, \tag{4}$$


$$\sum_j \omega_{ij} \|\mathbf{p}_i - \mathbf{p}_j\| = 0. \tag{5}$$

The sums in (4)–(5) are over neighbors j of i in G . We use the index ij to index one of the m edges of G .

This notion was studied in [6,13]. See also [20].

Lemma 7.11 *Let G be a maximal planar graph on n vertices. Let (\mathbf{p}, \mathbf{r}) be a configuration of n disks. Let \mathbf{v} be in the cokernel of J (defined using the edges of G). Let us define the associated stress vector ω as $\omega_{ij} := r_i r_j \mathbf{v}_{ij}$. Suppose disk i has tangential contact with its neighbors in G . Then ω satisfies the equilibrium condition at vertex i .*

Fig. 7 Illustration of the proof of Lemma 7.12. This vector configuration has only two sign changes (blue are negative signs and red are positive)



Proof This follows from nature of the entries in the associated columns of J and using (3). □

Lemma 7.12 [13] *Let G be a maximal planar graph on n vertices. Let (\mathbf{p}, \mathbf{r}) be a configuration of n disks. Let \mathbf{v} be a cokernel vector of J giving us an associated sign vector. Let i be a G -like disk. Then the index I_i is at least 4.*

Proof In what follows, we ignore any edge with a zero coefficient in \mathbf{v} . From Lemma 7.6, $I_i \geq 2$. So now we will suppose that $I_i = 2$ and arrive at a contradiction.

Using the assumed G -like property of disk i , we can apply Lemma 7.11. Thus the associated stress vector ω , where $\omega_{ij} := r_i r_j \mathbf{v}_{ij}$, must satisfy the equilibrium condition at vertex i .

With the supposed two sign changes, the edges with one of the signs (say $-$) must be in a cone of angle $2\theta < \pi$ (for example, the dotted red cone in Fig. 7). (Here we use the assumption that these edges are drawn in order in (G, \mathbf{p}) .)

Euclidean transforms have no effect on Eqs. (4) and (5), so, without loss of generality, we may assume that the positive part of the x -axis is the bisector of the cone. The 2D equilibrium condition of (4) must hold after projection along any direction, including onto the x -axis, since (4) is invariant under any affine transformation (see, e.g., [18]).

Let N^+ denote the neighbors of i connected by edges with positive sign and N^- the neighbors connected by negatively signed edges. Let p_i^x be the x -coordinate of the point \mathbf{p}_i . We then get

$$\sum_{j \in N^+} \omega_{ij}(p_i^x - p_j^x) = - \sum_{j \in N^-} \omega_{ij}(p_i^x - p_j^x).$$

But for $j \in N^+$ (outside the cone), we have $p_i^x - p_j^x < \|\mathbf{p}_i - \mathbf{p}_j\| \cos \theta$, while for $j \in N^-$ (inside the cone), $p_i^x - p_j^x > \|\mathbf{p}_i - \mathbf{p}_j\| \cos \theta$. Putting these estimates

together we have

$$\begin{aligned} \sum_{j \in N^+} \omega_{ij} \|p_i - p_j\| \cos \theta &> \sum_{j \in N^+} \omega_{ij} (p_i^x - p_j^x) \\ &= - \sum_{j \in N^-} \omega_{ij} (p_i^x - p_j^x) > - \sum_{j \in N^-} \omega_{ij} \|p_i - p_j\| \cos \theta, \end{aligned}$$

which means that (4) and (5) cannot hold simultaneously. □

See Fig. 7 for an illustration of this argument. And now we can prove our proposition.

Proof of Proposition 7.2 Suppose that J has a non-zero co-kernel vector v . Let n' be the number of disks with at least one incident edge in G with a non-zero v value. From Lemma 7.9, at most two of these n' disks are non G -like, and at least $n' - 2$ are G -like. Summing over the indices of $n' - 2$ disks that are G -like, and using Lemma 7.12, we obtain the value of at least $4n' - 8$. The remaining two of the n' disks do not overlap and so add at least 4 more to the sum using Lemma 7.6, giving us at least $4n' - 4$. But this contradicts Lemma 7.5, which ensures $\sum_i I_i \leq 4n' - 8$. □

7.1 Square Jacobian

The matrix J tells us how differential changes of (p', r') lead to differential changes of f' , in the inversive distances along the edges of a maximal planar graph G . The fact that, for an (almost-)triangulated packing, J has full row rank tells us that all differential changes of inversive distances on the edges of G are achievable. From this row rank and the matrix size, we see that the kernel of J is 6-dimensional, corresponding exactly to the six Möbius degrees of freedom.

Next we wish to mod out the Möbius degrees of freedom. We will do this by marking and pinning:

Definition 7.13 Let (p, r) be a packing with contact graph of either G, G^- , or H . Pick a triangular face of G^- , and mark the corresponding three disks. (Ultimately these three disks will form our canonical tridisk boundary.) We will also refer to the three edges in G , corresponding to this triangle as marked. We then pin the centers (but not radii) of our three marked disks. We call this a center-pinned packing. The effect of center pinning on J is simply to discard the corresponding six columns, giving us a square matrix of size $3n - 6$ we call the center-pinned Jacobian J_c .

Lemma 7.14 Let (p, r) be a center-pinned packing with contact graph of either G, G^- , or H . Then its center-pinned Jacobian (defined by edges in G) is a non-singular matrix.

Proof From Proposition 7.2, J has rank $3n - 6$. From Corollary 3.4 we can fix the five coordinates corresponding to say, p_1 and p_2 , and the x -coordinate of p_3 , and smoothly move the y -coordinate of p_3 using a family of Möbius transformations ϕ_t . Differentiating ϕ_t in t at $t = 0$ gives us a vector (p', r') which, from Theorem 2.7, must be in the kernel of J . By construction, (p', r') has zero entries corresponding

to the five fixed coordinates. Thus the column of J corresponding to the y -coordinate of p_3 must be linearly dependent on the other $3n - 6$ columns. So this column can be removed without changing the rank of the Jacobian. The same reasoning applies symmetrically to all six of these columns. Removing these six columns gives us the center-pinned Jacobian, J_c , which is square and has rank $3n - 6$. \square

Next, we note that J_c has a useful block form.

Definition 7.15 The three rows of J_c corresponding to the marked edges are only supported on the three columns corresponding to the three radii corresponding to the marked disks (the positional columns for these disks have already been removed). Let us now also pin these three radii, resulting in a pinned packing. By removing these three rows and columns, we obtain a square matrix of size $3n - 9$. We call this the pinned Jacobian J_p . Its rows correspond to the unmarked edges of G and its columns correspond to the positions and radii of the unpinned vertices.

Lemma 7.16 Let (p, r) be a pinned packing with contact graph of either G , G^- , or H . Then its pinned Jacobian (defined by the edges of G) is a non-singular matrix.

Proof This follows from Lemma 7.14 and its block form just described. \square

8 Packing Manifold

Definition 8.1 Let G and H be maximal planar graphs related by flipping an edge e^- , and G^- their common, almost-maximal planar graph. Pick a triangular face T of G^- , and mark (in G , G^- , and H) the corresponding three vertices, in some order, with the labels 1, 2, and 3. We also refer to the three edges of T as marked. Given any packing P with contact graph G , G^- , or H , we mark three of its disks using these vertex markings.

Define the packing set S_{G^-} , to be the subset of configurations of n disks that are packings and where the contact graph is exactly the graph G^- (no extra contacts), and that are tridisk-contained under the above marking. There are three common pinned disks in S_{G^-} . Let us consider the configuration space of the unpinned $n - 3$ disks as \mathbb{R}^{3n-9} .

This set satisfies a set of strict radius inequalities $r_i > 0$, strict inequalities over the non-edges, and equalities over the contact edges. Given these constraints, the tridisk containment can then be enforced by the strict inequality requirement that the unmarked-disk centers are in the open interior of the tricusp region.

Definition 8.2 The frontier points of S_{G^-} are defined as $\overline{S_{G^-}} - S_{G^-}$.

Now we establish some geometric properties of S_{G^-} .

Lemma 8.3 S_{G^-} is a 1-dimensional smooth manifold.

Proof Let P be a point in S_{G^-} . In a sufficiently small neighborhood of P in configuration space, all of the radii must remain positive, all non-edge connected disk pairs

must remain disjoint, and all of the unmarked-disk centers must be inside of the triscusp. (These are open conditions.) Thus locally, we need to only consider the $3n - 10$ inversive distances equalities of (1).

The partials of these constraints are described in the $3n - 10$ rows of J_p corresponding to the unmarked edges of G^- . By Lemma 7.16, these rows are linearly independent and thus, from the implicit function theorem, S_{G^-} , restricted to some neighborhood of P , is a smooth manifold of dimension 1. □

See [17] for some related and more general results.

Lemma 8.4 *S_{G^-} is bounded, and so its closure is compact.*

Proof The tridisk constraint bounds the coordinates of \mathbf{p} and \mathbf{r} . □

Lemma 8.5 *A frontier point of S_{G^-} can only be a tridisk-contained maximal triangulated packing with contact graph G or H (either ac or bd in contact.)*

Proof Let us consider the possible frontier points of S_{G^-} . From Proposition 6.3, for a packing in S_{G^-} , every radius must be bounded above zero. From Lemma 6.5, for a packing in S_{G^-} , every disk pair corresponding to any non-edge, except for ac or bd , must be bounded away from contact. Meanwhile, any disk configuration with positive radii, where some edge of G^- is not in tangential contact, cannot be approached while satisfying the (closed) equality conditions of S_{G^-} . Finally, the above conditions also keep all of the disk centers bounded away from the triscusp boundary. The lemma follows by this process of elimination. □

Ultimately, we will see below as a result of our flow argument, that each connected component of S_{G^-} has one frontier point that is a tridisk-contained packing with contact graph G and another frontier point that is a tridisk-contained packing with the flipped contact graph H . Moreover, a packing P with contact graph G (or resp. H) is in the closure of exactly one connected component of S_{G^-} .

9 Flow

In this section we will describe how to flow continuously from a tridisk-contained marked packing $P(0)$ with contact graph G to a tridisk-contained packing $P(1)$ with contact graph H (and the same marking). As above, the graph H will be obtained from G by removing an edge e^- (without loss of generality, its last edge), resulting in the graph G^- , and then adding in its cross edge e^+ . For $t \in (0, 1)$, the packing $P(t)$ will have contact graph G^- . Throughout the flow, the outer three disks will remain completely pinned, and thus we will work with $(\mathbf{p}, \mathbf{r}) \in \mathbb{R}^{3n-9}$.

We will find this flow as the solution to an ordinary differential equation that we set up now. Given any pinned disk configuration (\mathbf{p}, \mathbf{r}) , we have a pinned Jacobian matrix J_p . Where J_p is non-singular, we define the velocity field: $(\mathbf{p}', \mathbf{r}') := J_p^{-1}[0, 0, \dots, 0, 1]^t$. Differentially, this leaves all of the inversive distances fixed except for increasing the inversive distance corresponding to e^- . The matrix J_p is

non-singular for some (\mathbf{p}, \mathbf{r}) (Lemma 7.16), thus it is non-singular over a Zariski-open subset U of \mathbb{R}^{3n-9} . This defines a smooth (never zero) velocity field over U that includes $P(0)$, setting up a system of ODEs (ordinary differential equations). Given this system of ODEs, we start a trajectory at $P(0)$ and integrate forward in time. As a smooth ODE, this defines a unique maximal trajectory $(\mathbf{p}(t), \mathbf{r}(t))$ for some (possibly bi-infinite) open time interval.

The following is standard. See [29, Sect. 4.1] for a readable discussion.

Theorem 9.1 *Going forward in time, a maximal trajectory of a C^1 system of ODEs defined over an open set U must leave any compact set contained in U , or go on for infinite time.*

Lemma 9.2 $\overline{S_{G^-}} \in U$. So too is $P(0)$.

Proof From Lemma 8.5, all disk configurations in $\overline{S_{G^-}}$ are tridisk-contained almost-triangulated packings with packing graph G^- , or triangulated packings with contact graph G or H . Lemma 7.16 then guarantees they are in U . In particular, it guarantees that $P(0) \in U$. \square

Lemma 9.3 $(\mathbf{p}(t), \mathbf{r}(t)) \in S_{G^-}$ for sufficiently small and positive t .

Proof From Lemma 9.2, the ODE is well defined at $P(0)$. The ODE is constructed to maintain the equality constraints along the edges of G^- . Also by construction, the velocity of our trajectory at $t = 0$ must destroy the contact along the edge e^- for $t > 0$. For small enough t , no other inequality constraints can become violated. \square

Remark 9.4 This lemma also tells us that every packing P with contact graph G must be the frontier point of at least one connected component of S_{G^-} . By symmetry, the same is true of every packing Q with contact graph H .

Lemma 9.5 *For any trajectory starting in S_{G^-} , if the trajectory leaves S_{G^-} , it must first hit a frontier point of S_{G^-} .*

Proof The set S_{G^-} is defined by the intersection of some strict inequality (open) conditions and some equality constraints. Our ODE is constructed to maintain these equality constraints. To leave the set it must violate an inequality, and so, by continuity, the trajectory must first violate this inequality with an equality. Such a point is in the closure of S_{G^-} making it a frontier point. \square

Lemma 9.6 *Any trajectory starting in S_{G^-} must leave S_{G^-} at some positive t .*

Proof From Theorem 9.1, (using Lemmas 8.4 and 9.2) if the trajectory is only defined for a finite time, then it must leave S_{G^-} . But if the trajectory goes on for infinite time, it also must leave S_{G^-} by the following argument. Our ODE increases the inversive distance along e^- at a constant rate, and since the radii are bounded away from zero (Proposition 6.3), their centers must be moving apart at a lower bounded rate. Meanwhile S_{G^-} is bounded (Lemma 8.4). \square

Proposition 9.7 *Any trajectory starting in S_{G^-} must hit a positive t where the edge e^+ is in contact.*

Proof From Lemma 9.6, the trajectory must leave S_{G^-} . From Lemma 9.5, it must leave at a frontier point. From Lemma 8.5 this must either be contact along e^- or e^+ . But by construction, the inversive distance along e^- is increasing away from 1, so we cannot achieve contact along e^- . \square

Remark 9.8 This proposition also tells us that every connected component of S_{G^-} must have a frontier point Q with contact graph H . From the uniqueness of ODE solutions, two trajectories cannot cross, and thus only one component of S_{G^-} can have Q as a frontier point. By symmetry, every component also must have a frontier point P with contact graph G , and no other component can hit P .

In summary, each connected component of S_{G^-} must connect one tridisk-contained packing with contact graph G to one tridisk-contained packing with contact graph H . And each tridisk-contained packing with contact graph G (resp. H) is an endpoint of exactly one component of S_{G^-} .

We will stop our flow when the edge e^+ is in contact. After a time rescaling, we will consider this $t = 1$. We denote this tridisk-contained packing with contact graph H as $P(1)$. We refer to this trajectory between $t = 0$ and $t = 1$ as our flow.

With Lemma 3.7, Proposition 9.7, and our defined flow we have just proven the following result:

Theorem 9.9 *Let P be a triangulated packing, with contact graph G . Let e^- be a flippable edge of G . Let $G^- := G \setminus e^-$ and let H be the graph resulting from flipping e^- . Then there is a continuous path of packings $P(t)$, with $t \in [0, 1]$ such that:*

- $P(0) = P$.
- For t in the interval $(0, 1)$, the contact graph of $P(t)$ is G^- .
- The contact graph of $P(1)$ is H .

Now let us consider the inversive distance along e^+ during this flow.

Lemma 9.10 *The inversive distance along e^+ monotonically decreases, with derivative bounded away from zero during the flow from $P(0)$ to $P(1)$.*

Proof Let f^+ be the inversive distance along e^+ . If f^+ were to change sign, we must have a time during our flow when $df^+/dt = 0$. As $(\mathbf{p}', \mathbf{r}') \neq 0$ (the velocity field is never zero), this would have to be a time where the Jacobian of the "dual flip", defined using the edge of H , goes singular. But Lemma 7.16 can be applied just as well to this dual flip, so this cannot happen. Thus during our flow, f^+ changes monotonically, with non-zero derivative. Meanwhile, we know that f^+ must overall decrease from its starting value (non-contact) down to 1 (contact). Since we are working over a compact time interval $[0, 1]$, this derivative is bounded away from zero. \square

Corollary 9.11 *If we start with the graph H and remove the edge e^+ , and set up a dual flow starting at $P(1)$, increasing the inversive distance of this edge, then this dual flow must drive e^- to contact and give us $P(0)$. This dual flow must travel along the reverse path of the original forward flow, defined using the graph G and the edge e^- .*

Proof We can set up a dual flow starting at $P(1)$ and increasing the inversive distance on e^+ . From Theorem 9.1, Lemma 9.3, and Proposition 9.7 (applied to this dual

differential equation) there is also a *unique* trajectory satisfying this dual differential equation. But from Lemma 9.10, the reverse path from $P(1)$ to $P(0)$ satisfies this dual differential equation (up to a smooth regular parameterization of the time variable). \square

10 Proving KAT

With these pieces all in place, we can now establish all of the elements of the KAT circle packing theorem.

Theorem 10.1 *Let M be a maximal planar graph on n vertices. Then there exists a packing of n disks with contact graph M .*

Proof Let P be a trilaterated packing with n disks and with contact graph C . From Theorem 5.6 we can combinatorially transform C to M using a finite sequence of edge flips. Starting with P , we wish to flow the packing across each of the edge flips in our sequence.

To this end, let G and H represent two maximal planar graphs related by flipping some edge e^- . Let P be a packing with contact graph G . We mark three disks in P corresponding to a triangle in G^- . Using Lemma 3.7, we apply a Möbius transformation to P so that the resulting marked packing is tridisk-contained. Now we can apply Proposition 9.7 to flow this packing to a new tridisk-contained marked packing Q with contact graph H .

We can do this for each edge flip in our sequence, arriving at a packing with contact graph M . \square

Theorem 10.2 *Let M be a maximal planar graph on n vertices. Then its packing is unique up to a generalized Möbius transformation.*

Proof We will follow the same sequence of edge flips and flows as in the proof of Theorem 10.1. From Lemma 4.2, the trilaterated packing is unique up to generalized Möbius transformations.

Let us look at a single edge flip, flowing from a tridisk-contained marked packing P with contact graph G to the tridisk-contained packing Q with the same marking and with contact graph H . Let us assume (as an inductive invariant) that P was the only packing with contact graph G up to a generalized Möbius transformation. Then from Lemma 3.8, it is the only tridisk-contained packing with this contact graph and marking.

Now let Q' be any tridisk-contained packing with this marking and contact graph H . Then, from Proposition 9.7, the dual differential equation starting from Q' must flow to *some* tridisk-contained marked packing P' with contact graph G . Setting up the double dual ODE (which is just our original ODE), starting from P' and using Corollary 9.11, the double dual flow must lead us back to the marked packing Q' . But from our assumed marked uniqueness, $P = P'$. And since our double dual ODE is just our original ODE, then $Q = Q'$ (as marked packings). Thus from Lemma 3.8, Q is the only packing with contact graph H up to a generalized Möbius transformation. (See also Remark 9.8 for the geometric point of view.)

In summary, the number of distinct packings, up to generalized Möbius transformations, cannot increase across an edge flip, maintaining our inductive invariant, and proving the theorem. \square

Remark 10.3 If we fix G , but extend our configuration space to allow non-edge disk pairs to overlap, and thus allow (G, \mathbf{p}) to be a non-embedding, then all bets (such as Lemma 7.16) are off. Indeed, global uniqueness no longer holds. (For example, smallest disk in Fig. 3 can be changed to be identical to the disk labeled 1 in the figure, while maintaining all of the necessary external tangential contact.) This leads to especially interesting situations when using general inversive distance constraints on the sphere [23].

Theorem 10.4 *Let G be a planar graph on n vertices. Then there exists a packing of n disks with contact graph G .*

Proof We can always add edges to G to obtain a maximal planar graph M . From Theorem 10.1, there must be a tridisk-contained (using Lemma 3.7) packing P with contact graph M . To remove the extraneous contacts, all we need to do is create a short lived flow starting at P that increases the inversive distances along the added edges, while maintaining contact along the edges of G .

Note that in this setting, we may also need to separate the three outer disks of P (as there may be no triangles at all in G). To do this we will work in the $(3n - 6)$ -dimensional configuration space, where the three outer radii, but not their centers, are variable.

From Lemma 7.14, the center-pinned Jacobian J_c , defined using the edges of M , is non-singular at P . From upper semi-continuity of rank, it remains non-singular in a sufficiently small neighborhood U of P .

We define our velocity field as $(\mathbf{p}', \mathbf{r}') := J_c^{-1} \mathbf{f}'$, where $f'_{ij} > 0$ for the added edges, and $f'_{ij} = 0$ for the edges of G . Finally, we integrate this flow from $t = 0$ to $t = \epsilon$ for a sufficiently small positive ϵ . \square

The KAT circle packing theorem is simply the union of Theorems 10.1, 10.2, and 10.4.

11 Generalizations

11.1 Other Powerful Tools

This proof of the KAT circle packing theorem is based on establishing the infinitesimal rigidity of the inversive distances along the edges of a maximal planar graph G , for any (\mathbf{p}, \mathbf{r}) that is an almost-triangulated packing with contact graph G^- . We established this in Sect. 7 using the equilibrium condition. In particular, we use the fact that all but one of the edges of G represent tangential contact.

Another set of tools for proving such infinitesimal rigidity results comes from the work of Guo [16]. Although he works in the intrinsic setting of polyhedral metrics, his result can also be applied to any set of disks (\mathbf{p}, \mathbf{r}) , with any inversive distances in

$[0, \infty]$ along the edges of G , whenever the extrinsic \mathbb{R}^2 drawing, (G, \mathbf{p}) , is a planar embedding.

Now in our setting, we know that (G^-, \mathbf{p}) is a planar embedding, but we still have one pesky edge e^- , which might become out of order during the flow. And at some time, one of the triangles with this edge can become degenerate. Luckily, it seems that one can always find a Möbius transformation (that does not invert any disks) under which (G, \mathbf{p}) does become an embedding. The *existence* of such an embedded drawing should be enough to argue the non-singularity of J_p for the packing under consideration.

11.2 More General Inversive Distance Targets

The full KAT theorem covers configurations of disks that are allowed some amount of overlap [7,33]. One fixes each of the inversive distances along the edges ij of a maximal planar graph G to some target value f_{ij} in the interval $[0, 1]$ (instead of equal to 1). This allows the disks across an edge to intersect with some fixed angle $\leq \pi/2$. Disks not connected by an edge are still not allowed to overlap. In this setting, for existence, one must also add in an assumption about the sums of the overlapping angles in 3-cycles that are not triangles in G , and also on the sums over 4-cycles. It would be interesting to see if we can use our flow-based approach to start with a packing with tangential contact along the edges of G , and then flow the disks together to overlap by the desired amount. For this, it might be useful to apply the tools of Sect. 11.1.

Work starting with Rivin [26] has investigated the situation with $f_{ij} \in (-1, 1]$ (allowing disk intersections up to π). Existence and uniqueness can be shown under some additional assumptions. Again, it would be interesting to see if we can apply our flow-based approach to this setting.

If one has a set of disks (\mathbf{p}, \mathbf{r}) , with some set of inversive distances $f_{ij} \in [0, \infty]$ along the edges of a maximal planar graph G , the story gets even more complicated. If we work in the category of disks where (G, \mathbf{p}) is a planar embedding, then infinitesimal rigidity follows from [16,35] and global uniqueness follows from [22,35] (though see [5,23] for non-uniqueness in the spherical embedded setting). But when we place such general inversive distance targets on the edges of G , then existence conditions are not well understood.

11.3 Higher Genus

The KAT theorem also generalizes to higher genus settings [33]. We conjecture that given a graph G that triangulates the torus, and a packing on a flat torus with this contact graph, we can flip an edge, and smoothly flow the packing across this flip (this necessarily requires us to also flow the shape of the torus). For example, in his book, "Regular Figures" [14] László Fejes Tóth takes the periodic packing on the left in Fig. 8, which is triangulated, and continuously deforms the packing. This is a flip-and-flow motion as we have described in this paper, applied to the edge between the blue and yellow regions, but now on a torus. This flow changes the radii as well as the




Fig. 8 This shows a continuous flip of an edge of a triangulated packing of a torus as used by László Fejes Tóth. The packing on the left is the start of the continuous motion of the packing. The middle packing was used because the relative ratio of the packing radii were closer to one, and the packing on the right is the standard triangulated packing with all the radii the same. Three fundamental regions are shown, each in a different color

density of the packing. If one continues the motion, the packing eventually moves to the standard packing with all radii the same as in the right of the figure.

We would also be interested in seeing if our method extends to surfaces of higher genus in the hyperbolic setting.

Acknowledgements The authors would like to thank Louis Theran for many helpful discussions along the way.

References

- Alexandrov, A.D.: *Convex Polyhedra*. Springer Monographs in Mathematics. Springer, Berlin (2005)
- Andreev, E.M.: On convex polyhedra of finite volume in Lobachevskii space. *Math. USSR-Sbornik* **12**(2), 255–259 (1970)
- Bobenko, A.I., Springborn, B.A.: Variational principles for circle patterns and Koebe’s theorem. *Trans. Am. Math. Soc.* **356**(2), 659–689 (2004)
- Bose, P., Hurtado, F.: Flips in planar graphs. *Comput. Geom.* **42**(1), 60–80 (2009)
- Bowers, J.C., Bowers, P.L.: Ma–Schlenker c -octahedra in the 2-sphere. *Discrete Comput. Geom.* **60**(1), 9–26 (2018)
- Bowers, J.C., Bowers, P.L., Pratt, K.: Almost all circle polyhedra are rigid. *Geom. Dedicata* **203**, 337–346 (2019)
- Bowers, P., Stephenson, K.: A branched Andreev–Thurston theorem for circle packings of the sphere. *Proc. Lond. Math. Soc.* **73**(1), 185–215 (1996)
- Chow, B., Luo, F.: Combinatorial Ricci flows on surfaces. *J. Differ. Geom.* **63**(1), 97–129 (2003)
- Colin de Verdière, Y.: Un principe variationnel pour les empilements de cercles. *Invent. Math.* **104**, 655–669 (1991)
- Collins, Ch.R., Stephenson, K.: A circle packing algorithm. *Comput. Geom.* **25**(3), 233–256 (2003)
- Connelly, R.: Rigidity and energy. *Invent. Math.* **66**(1), 11–33 (1982)
- Connelly, R.: What is . . . a tensegrity? *Not. Am. Math. Soc.* **60**(1), 78–80 (2013)
- Connelly, R., Gortler, S.J., Theran, L.: Rigidity for sticky discs. *Proc. R. Soc. A* **475**(2222), # 20180773 (2019)
- Fejes Tóth, L.: *Regular Figures*. A Pergamon Press Book, Macmillan, New York (1964)
- Gluck, H.: Almost all simply connected closed surfaces are rigid. In: *Geometric Topology* (Park City 1974). *Lecture Notes in Math.*, vol. 438, pp. 225–239. Springer, Berlin (1975)
- Guo, R.: Local rigidity of inversive distance circle packing. *Trans. Am. Math. Soc.* **363**(9), 4757–4776 (2011)
- He, Z., Liu, J.: On the Teichmüller theory of circle patterns. *Trans. Am. Math. Soc.* **365**(12), 6517–6541 (2013)

18. Izmistiev, I.: Projective background of the infinitesimal rigidity of frameworks. *Geom. Dedicata* **140**, 183–203 (2009)
19. Koebe, P.: Kontaktprobleme der konformen Abbildung. *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig, Mathematisch-Physische Kl.* **88**, 141–164 (1936)
20. Lam, W.Y.: Minimal surfaces from infinitesimal deformations of circle packings. *Adv. Math.* **362**, # 106939 (2020)
21. Laumond, J.-P.: Connectivity of plane triangulations. *Inform. Process. Lett.* **34**(2), 87–96 (1990)
22. Luo, F.: Rigidity of polyhedral surfaces, III. *Geom. Topol.* **15**(4), 2299–2319 (2011)
23. Ma, J., Schlenker, J.-M.: Non-rigidity of spherical inversive distance circle packings. *Discrete Comput. Geom.* **47**(3), 610–617 (2012)
24. Marden, A., Rodin, B.: On Thurston’s formulation and proof of Andreev’s theorem. In: *Computational Methods and Function Theory (Valparaiso 1989)*. Lecture Notes in Math., vol. 1435, pp. 103–115. Springer, Berlin (1990)
25. Mohar, B.: A polynomial time circle packing algorithm. *Discrete Math.* **117**(1–3), 257–263 (1993)
26. Rivin, I.: Euclidean structures on simplicial surfaces and hyperbolic volume. *Ann. Math.* **139**(3), 553–580 (1994)
27. Rodin, B., Sullivan, D.: The convergence of circle packings to the Riemann mapping. *J. Differ. Geom.* **26**(2), 349–360 (1987)
28. Roeder, R.K.W., Hubbard, J.H., Dunbar, W.D.: Andreev’s theorem on hyperbolic polyhedra. *Ann. Inst. Fourier (Grenoble)* **57**(3), 825–882 (2007)
29. Schaeffer, D.G., Cain, J.W.: *Ordinary Differential Equations: Basics and Beyond*. Texts in Applied Mathematics, vol. 65. Springer, New York (2016)
30. Schramm, O.: Existence and uniqueness of packings with specified combinatorics. *Israel J. Math.* **73**(3), 321–341 (1991)
31. Sleator, D.D., Tarjan, R.E., Thurston, W.P.: Short encodings of evolving structures. *SIAM J. Discrete Math.* **5**(3), 428–450 (1992)
32. Stephenson, K.: *Introduction to Circle Packing*. Cambridge University Press, Cambridge (2005)
33. Thurston, W.P.: *The Geometry and Topology of Three-Manifolds*. Princeton University, Princeton (1980). (unpublished notes)
34. Wagner, K.: Bemerkungen zum Vierfarbenproblem. *Jahresbericht der Deutschen Mathematiker-Vereinigung* **46**, 26–32 (1936)
35. Xu, X.: Rigidity of inversive distance circle packings revisited. *Adv. Math.* **332**, 476–509 (2018)

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.