Moduli Spaces of Circle Packings

Allen Hatcher and William Thurston

Historical Note. This paper, written in 2013-2015 by the first author after Bill's untimely death in 2012, is an exposition of parts of an unfinished project that the two of us worked on in 2008, with some valuable input also from Richard Kenyon in the early stages. The project began with a talk by Richard at the Cornell Topology Festival in May 2008 on branched polymers (called circle packings in the present paper). Most of the communication was by email, with over 50 messages from Bill. He also talked about this topic in his course in the Fall semester of 2008, but unfortunately I did not take notes during the course and have only a dim recollection of what he said, so the exposition will be based largely on the email record, apart from details of proofs that were worked out as the paper was written.

The paper presents only a small fraction of the many things discussed in the email, which were often left hanging without a clear resolution. The plan was for me to write up the two theorems included here and then we would move on to try to prove more, in particular the Conjecture discussed in the introduction. To my great regret, I did not get around to doing this before Bill's illness. In particular, I did not learn from Bill all the details of the proof of the first theorem, so only the parts I do understand are given here, although I have been assured by an expert on the relevant tools that there should be no difficulty in completing the proof.

By a *circle packing* we mean a finite collection of circles in $\mathbb{R}^2$ whose interiors are disjoint and whose union is connected. The space of all such configurations of n labeled circles $C_1, \cdots, C_n$ will be denoted $\mathrm{CP}(n)$, the moduli space of packings of n circles. One might ask what can be said about the homotopy type of $\mathrm{CP}(n)$. The question can also be asked for the subspace of 'equal packings', $\mathrm{EP}(n)$, the configurations whose circles all have equal radius, which one can normalize to be 1. Physically, one can think of points in $\mathrm{EP}(n)$ as configurations of n pennies lying flat on a table, with connected union.

There is a map from $\mathrm{CP}(n)$ to the space $\mathrm{C}(n)$ of configurations of n distinct labeled points in $\mathbb{R}^2$, sending a configuration of circles to the centers of these circles. This map is neither surjective nor injective in general. For example, a configuration of four points at the corners of a long thin rectangle cannot be realized as the centerpoints of a circle packing, while a configuration of n equally-spaced points along a line is realized by a whole 1-parameter family of circle packings. The main result of this paper is:

Theorem 1. *The map* $\mathrm{CP}(n) \to \mathrm{C}(n)$ *is a homotopy equivalence.*

As a small preliminary reduction of the problem we can identify configurations that differ by orientation-preserving similarities of the Euclidean plane to obtain a quotient space $\overline{\mathrm{CP}}(n)$ with $\mathrm{CP}(n)$ homeomorphic to $\overline{\mathrm{CP}}(n) \times \mathbb{R}^2 \times (0,\infty) \times S^1$ where the last three factors correspond to translations, rescalings, and rotations of the plane, leaving aside the trivial case $n = 1$. We can view $\overline{\mathrm{CP}}(n)$ as a subspace of $\mathrm{CP}(n)$, the configurations that are normalized to have the first circle centered at the origin and the second circle centered at $(1,0)$. The space $\mathrm{C}(n)$ splits in the same way, so the theorem is equivalent to the assertion that the map $\overline{\mathrm{CP}}(n) \to \overline{\mathrm{C}}(n)$ is a homotopy equivalence.

The idea of the proof is now the following. Associated to any configuration in $\mathrm{CP}(n)$ is its *contact graph*, whose vertices are the centers of the circles in the configuration and whose edges are the line segments connecting the centers of each pair of circles that touch. The subsets of $\mathrm{CP}(n)$ consisting of configurations whose contact graphs vary only by isotopy form the strata of a stratification of $\mathrm{CP}(n)$. The quotient $\overline{\mathrm{CP}}(n)$ has a similarly-defined stratification. It follows from Ahlfors–Bers theory for Kleinian groups that each stratum of $\overline{\mathrm{CP}}(n)$ is homeomorphic to a Euclidean space. The way that the strata fit together is encoded in a finite simplicial complex $\mathrm{G}(n)$ whose vertices correspond to the isotopy classes of connected graphs in the plane with n labeled vertices, and whose simplices correspond to chains of inclusions of such graphs. There is a projection $\overline{\mathrm{CP}}(n) \to \mathrm{G}(n)$, and showing that this is a homotopy equivalence constitutes the first half of the proof. (This is where there are still some details that need to be filled in.) This reduces the problem to the more combinatorial one of understanding the homotopy type of $\mathrm{G}(n)$. Topological techniques then show that $\mathrm{G}(n)$ has the homotopy type of $\overline{\mathrm{C}}(n)$.

Since the strata of $\overline{\mathrm{CP}}(n)$ are homeomorphic to Euclidean spaces, one could hope that they formed the cells of a CW complex structure on $\overline{\mathrm{CP}}(n)$. This is not the case, however. There are only finitely many strata since the strata correspond to the vertices of the finite complex $\mathrm{G}(n)$, so if these strata were the cells of a CW complex, it would be a finite complex, hence $\overline{\mathrm{CP}}(n)$ would be compact. But $\overline{\mathrm{CP}}(n)$ is not compact when $n \geq 2$ since the ratio of the radii of the first two circles gives a continuous map $\overline{\mathrm{CP}}(n) \to (0,\infty)$ which is easily seen to be surjective by considering just configurations of circles with their centers along a line. As we will see, the dimension of $\overline{\mathrm{CP}}(n)$ is $2n-3$ and the dimension of $\mathrm{G}(n)$ is $2n-5$ (when these numbers are positive) so one could hope that $\overline{\mathrm{CP}}(n)$ is homeomorphic to $\mathrm{G}(n) \times \mathbb{R}^2$. This is true at least for $n = 3$.

Let us turn now to the subspace $\mathrm{EP}(n) \subset \mathrm{CP}(n)$ consisting of configurations of unit circles. One might guess that $\mathrm{EP}(n)$ is also homotopy equivalent to $\mathrm{C}(n)$. Perhaps surprisingly, this is false once n is sufficiently large. The configuration space $\mathrm{C}(n)$ is a $K(\pi,1)$ for π the pure braid group, but we will show by an explicit elementary construction that $\mathrm{EP}(n)$ is generally not a $K(\pi,1)$. Specifically:

Theorem 2. *$\pi_5\mathrm{EP}(n)$ is nonzero if $n \geq 8$.*

This will be proved by constructing an embedding of S^5 in $\mathrm{EP}(n)$ whose image consists of configurations containing 7 circles arranged in a 'necklace' around a fixed central circle, so there is one more circle in the necklace than the maximum number that can touch the central circle simultaneously. Then we construct a map $\mathrm{EP}(n) \to S^5$ whose restriction to this $S^5 \subset \mathrm{EP}(n)$ has degree 1. This gives a direct summand $\mathbb{Z}$ in $\pi_5\mathrm{EP}(n)$.

The difference between $\mathrm{CP}(n)$ and $\mathrm{EP}(n)$ is reflected in how the strata of $\mathrm{CP}(n)$ intersect $\mathrm{EP}(n)$. Some strata do not intersect $\mathrm{EP}(n)$ at all, for example those corresponding to contact graphs containing a vertex that is an endpoint of more than six edges. It can also happen that the strata of $\mathrm{CP}(n)$ intersect $\mathrm{EP}(n)$ in nontrivial homotopy types. We give some simple examples in section 1 coming from the theory of planar linkages.

It seems reasonable to ask whether the result in the preceding theorem is optimal:

Conjecture. The map $\pi_i\mathrm{EP}(n) \to \pi_i\mathrm{C}(n)$ is an isomorphism if either $n \leq 7$ or $i \leq 4$.

This is easy to check by hand for $n \leq 3$. For $n = 2$ both spaces having the homotopy type of S^1 and when $n = 3$ they have the homotopy type of $S^1 \times (S^1 \vee S^1)$.

A small part of the conjecture is the assertion that $\mathrm{EP}(n)$ is connected for all n. This seems to be harder to prove than one might expect. In the physical model consisting of connected configurations of pennies lying flat on a table, one could attempt to deform an arbitrary configuration to one consisting of a single straight line by taking two parallel rulers above and below a given configuration and squeezing them closer and closer together to force the pennies into a line ("penny pinching"). Then it would be a simple matter to rearrange their order arbitrarily. In general this squeezing process could get stuck in certain configurations, but one would guess that a small perturbation could then be made to get it unstuck. Making this rigorous seems difficult, however. A different sort of argument which seems to work is based on the idea of pulling out a longer and longer string of pennies from a given configuration.

It is not hard to show that the map $\pi_1\mathrm{EP}(n) \to \pi_1\mathrm{C}(n)$ is surjective since all one has to do is construct loops in $\mathrm{EP}(n)$ that project to a set of generators for $\pi_1\mathrm{C}(n)$. It is a little simpler, and it suffices, to do this for the full braid group rather than the pure braid group, so one is talking about configurations of unlabeled points or circles. Then one just has to lift the standard generators σ_i transposing the ith and $(i+1)$st points, which is easy to do starting with a configuration of n circles in a row along a horizontal line. With a little more work one can check that the standard braid relations among the σ_i's also lift, so the map $\pi_1 \mathrm{EP}(n) \to \pi_1 \mathrm{C}(n)$ is in fact split surjective.

One could try to prove the Conjecture by studying the map $\rho : \mathrm{CP}(n) \to (0,\infty)^n$ sending a configuration of circles to the n-tuple of radii of the circles. The map ρ is not a fibration of any sort since if it were, the contractibility of the base $(0,\infty)^n$ would imply that all the fibers, including $\mathrm{EP}(n)$, had the homotopy type of $\mathrm{CP}(n)$ and hence of $\mathrm{C}(n)$, contradicting Theorem 2, at least when $n \geq 8$ (and in fact when $n \geq 4$ by a small generalization of Theorem 2 given in section 1). Thus the projection ρ has certain singular values where the homotopy type of the fibers is not locally constant. The idea would then be to understand these singularities well enough to show for example that the inclusion $\mathrm{EP}(n) \to \mathrm{CP}(n)$ induces an injection on π_1 by taking a contraction in $\mathrm{CP}(n)$ of a loop in $\mathrm{EP}(n)$ representing an element of the kernel of $\pi_1 \mathrm{EP}(n) \to \pi_1 \mathrm{CP}(n)$ and deforming this contraction into the fiber $\mathrm{EP}(n)$. A model for this sort of approach is Morse theory, where one studies how the homotopy type of the fibers of a projection to $\mathbb{R}$ varies.

1. Necklace Configurations

In this section we prove Theorem 2. The argument also works in somewhat greater generality, in particular for subspaces of $\mathrm{CP}(n)$ where all but one of the circles have the same radius. In these cases it is other homotopy groups besides π_5 that are nontrivial. This is explained after the proof of Theorem 2.

The starting point is the observation that for circles of equal radius, exactly six circles fit in a ring or necklace surrounding a central circle, as shown in the first figure below.

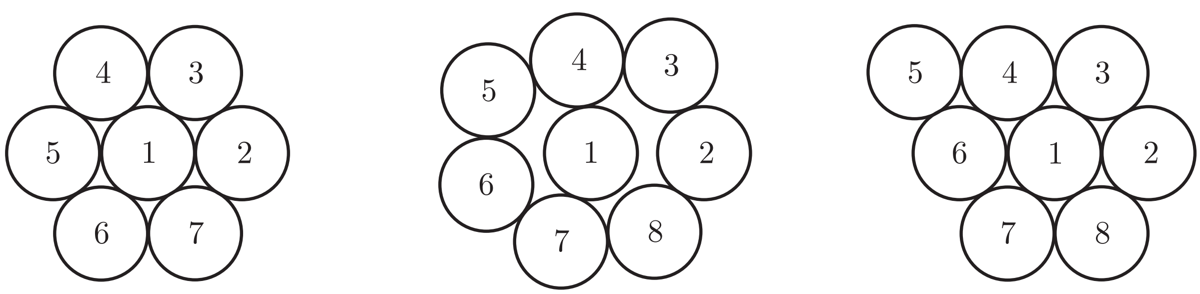

The idea is to see what happens when the necklace is enlarged to seven circles surrounding a central circle, as in the second figure. We do not require that the seven circles have their centers exactly in a circle, just that each circle touches its two neighbors. The central circle must also touch at least one of the seven circles in the necklace in order to get a configuration in $\mathrm{EP}(8)$. We allow the limiting case that one of the seven circles moves so far away from the central circle that its two neighbors touch each other, as in the third figure.

Let $N \subset \mathrm{EP}(8)$ be the set of all such necklace configurations with the central circle C_1 surrounded by a ring of seven circles $C_2, \cdots, C_8$ in that order, normalized so that C_1 is centered at the origin and C_2 has its center on the positive x-axis. We claim that N is homeomorphic to the sphere S^5. Because of the normalizations, a configuration in N is uniquely determined by the distances δ_i from C_i to C_1 for $i = 2, \cdots, 8$. Namely, δ_2 determines the position of C_2 on the positive x-axis, then δ_3 determines the position of C_3 touching C_2, then δ_4 determines the position of C_4 touching C_3, and so on. Here we are assuming the δ_i's are not too large, so δ_i ranges only from 0 to the value in the limiting configuration with C_i outside the other six circles that all touch C_1. At the end of the process of placing the circles $C_2, \cdots, C_8$ around C_1 there is also the condition that these circles close up to form a necklace, so C_8 just touches C_2 without a gap or an overlap. This imposes a certain constraint on the δ_i's, but we will not need to determine the exact form of this constraint.

Let us show that for a given nonzero vector $(\delta_2, \cdots, \delta_8)$ with each $\delta_i \geq 0$ there is a unique positive scalar multiple $\lambda(\delta_2, \cdots, \delta_8)$ that makes the necklace exactly close up. To see this, observe that when the δ_i's are all sufficiently small the last circle C_8 definitely overlaps C_2, as the sum of the angles between the rays from the origin to the centers of successive C_i's is more than 2π. Then as one takes larger and larger scalars λ the angles decrease strictly monotonically (an exercise), so there is a unique λ for which the angle sum is 2π. Thus the possible positions for $C_2, \cdots, C_8$ forming a necklace are parametrized by Δ^6, the projectivization of $[0, \infty)^7$. We also have the condition that some δ_i must be 0 so that C_1 touches at least one other C_i, so this means that we are in $\partial\Delta^6$, a 5-sphere as claimed.

To show that this sphere $N = S^5$ is nontrivial in $\pi_5\mathrm{EP}(8)$ we define a map $\mathrm{EP}(8) \to S^5$ that restricts to a homeomorphism on N. The distance functions δ_i from C_i to C_1 for $i = 2, \cdots, 8$ are defined for arbitrary configurations in $\mathrm{EP}(8)$. They cannot all be simultaneously 0 since at most six C_i's can touch C_1, so we can projectivize to get a map $\mathrm{EP}(8) \to \Delta^6$. This must have image in $\partial\Delta^6$ since for each configuration some δ_i must be 0. This gives the desired map $\mathrm{EP}(8) \to S^5$, and it obviously restricts to a homeomorphism on N. Thus S^5 is a retract of $\mathrm{EP}(8)$ (not a deformation retract), giving a $\mathbb{Z}$ summand of $\pi_5\mathrm{EP}(8)$.

Now we extend this construction to $\mathrm{EP}(n)$ for $n > 8$. We can enlarge the necklace configurations in N by placing new circles $C_9, \cdots, C_n$ to the right of C_2 with their centers along the positive x-axis, with C_9 touching C_2, then C_{10} touching C_9, and so on. This gives an embedding $N = S^5 \subset \mathrm{EP}(n)$.

The distance functions δ_i from C_i to C_1 for $i \geq 2$ give a map $\mathrm{EP}(n) \to [0, \infty)^{n-1}$. The image of this map misses the origin since at most six other C_i's can touch C_1, so we can projectivize to get a map $f : \mathrm{EP}(n) \to \Delta^{n-2}$. This has image in $\partial\Delta^{n-2}$ since at least one C_i must touch C_1. We cannot have seven or more δ_i's equal to 0 at once, so the image of f must lie in the complement of the codimension 7 skeleton of Δ^{n-2}. In $\partial\Delta^{n-2}$ this is the complement of the codimension 6 skeleton of $\partial\Delta^{n-2}$. Call this complement $W \subset \partial\Delta^{n-2}$, so f can be viewed as a map $\mathrm{EP}(n) \to W$. It is an elementary fact that W has the homotopy type of a wedge of 5-spheres, namely the 5-skeleton of the $(n-2)$-simplex dual to Δ^{n-2} , but we will not need to use this fact and instead we make the following direct argument.

Let the simplex Δ^{n-2} have vertices $v_2, \cdots, v_n$ corresponding to the circles $C_2, \cdots, C_n$. The sphere $f(N)$ is homotopic in W to a small linking sphere of the simplex $\langle v_9, \cdots, v_n \rangle$ by letting the values of $\delta_9, \cdots, \delta_n$ become large. We can include W into $\partial\Delta^n - \partial\langle v_8, \cdots, v_n \rangle$ which is homotopy equivalent to a 5-sphere, the link of $\langle v_9, \cdots, v_n \rangle$. Thus we have a composition $BP(n) \xrightarrow{f} W \to S^5$ whose restriction to $N = S^5$ has degree 1. It follows that N generates a $\mathbb{Z}$ summand of $\pi_5 \mathrm{EP}(n)$.

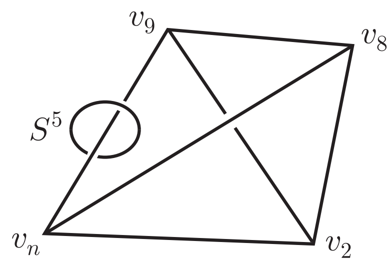

This finishes the proof of Theorem 2.

Varying the Size of the Central Circle

Consider the subspace $\mathrm{CP}_r(n)$ of $\mathrm{CP}(n)$ consisting of configurations where C_1 has radius r and the other circles all have radius 1. For small values of r and $n = 4$ one can form configurations as in the figure below, with C_1 inside the curvilinear triangle formed by arcs of C_2, C_3, C_4.

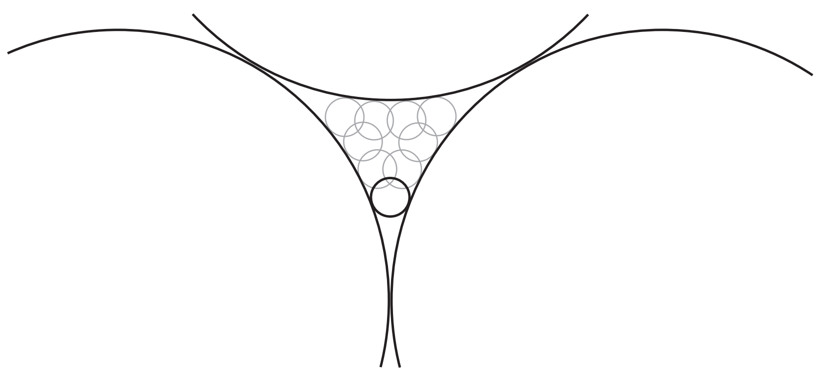

Modulo translations and rotations of the plane, such configurations obviously form a circle $S^1 \subset \mathrm{CP}_r(4)$, with C_1 moving inside the triangle so as to stay in contact with at least one of the other circles. This circle represents an element of the kernel of $\pi_1 \mathrm{CP}_r(4) \to \pi_1 \mathrm{C}(4)$, and it is nontrivial in $\pi_1 \mathrm{CP}_r(4)$ by the same sort of argument as before. This works in $\mathrm{CP}_r(n)$ for $n > 4$ as well.

As r increases nothing changes until r reaches the value $r = r_3$ where C_1 is exactly large enough to touch all three of C_2, C_3, C_4 simultaneously. The S^1 in $\mathrm{CP}_r(4)$ then degenerates to a point, so we have an explicit contraction in $\mathrm{CP}(4)$ of the S^1 in $\mathrm{CP}_r(4)$.

Now for $r = r_3$ we add another circle C_5 to form a necklace of four circles around

C_1. Arguing as before, the set of such configurations is an $S^2 \subset \mathrm{CP}_r(5)$ giving a direct summand $\mathbb{Z} \subset \pi_2\mathrm{CP}_r(5)$ for $r_3 \leq r < r_4$ where r_4 is the value of r for which four unit circles exactly fit around C_1. Letting r continue to increase, we obtain in this way an infinite sequence of values $r_3 < r_4 < \cdots$ with subgroups $\mathbb{Z} \subset \pi_{k-2}\mathrm{CP}_r(k+1)$ for $r_{k-1} \leq r < r_k$. These subgroups survive in $\pi_{k-2}\mathrm{CP}_r(n)$ for $n > k+1$ as well.

Noncontractible Strata

From the theory of planar linkages of rigid rods connected by pivots at their ends it is not hard to produce noncontractible strata in the quotient $\overline{\mathrm{EP}}(n)$ of $\mathrm{EP}(n)$ by translations and rotations of $\mathbb{R}^2$. Consider for example configurations whose contact graphs have the form in the figure below, consisting of a rigid supporting base with three rigid trusses that are hinged to rotate where they touch each other and the base.

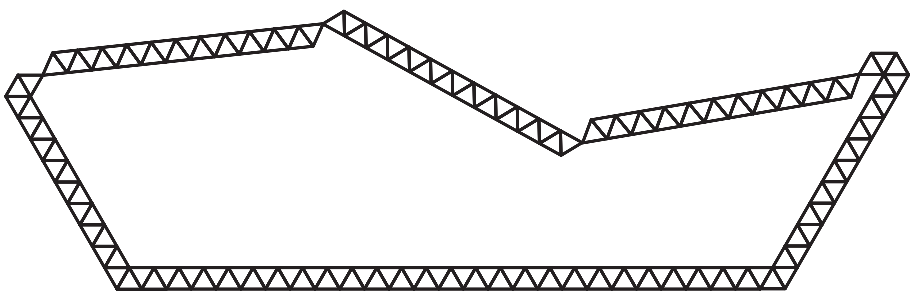

When the base is held fixed there is one degree of freedom in these configurations, and they form a stratum of $\overline{\mathrm{EP}}(n)$ homeomorphic to a circle. More generally, with k trusses whose total length is slightly longer that the distance between the support points at the ends of the base, the stratum is homeomorphic to a sphere S^{k-2}. This can be seen by induction on k, each increase in k producing a suspension of the previous space of configurations.

One of the basic results about planar linkages is that every closed connected smooth manifold can be realized as a component of the space of configurations of some linkage, and one could ask whether the same is true for strata of the spaces $\overline{\mathrm{EP}}(n)$.

2. Circle Packings and Contact Graphs

In order to show that the map $\overline{\mathrm{CP}}(n) \to \overline{\mathrm{C}}(n)$ is a homotopy equivalence we will define two spaces of graphs $\mathrm{G}(n)$ and $\overline{\mathrm{GC}}(n)$ together with maps as in the following diagram:

$$\begin{tikzcd}
\overline{\mathrm{CP}}(n) \arrow[rr] \arrow[d, "(1)"'] \arrow[drr] & & \overline{\mathrm{C}}(n) \\
\mathrm{G}(n) & & \overline{\mathrm{GC}}(n) \arrow[ll, "(2)"] \arrow[u, "(3)"']
\end{tikzcd}$$

Consider finite connected graphs embedded in $\mathbb{R}^2$ with n vertices, labeled $1, \cdots, n$, and edges that are smooth arcs. We mean the term 'graph' here in the strict sense of a 1-dimensional simplicial complex, so there are no edges whose two endpoints coincide, nor are there pairs of edges with the same two endpoints. In addition, positive weights are assigned to the edges. The set of all isotopy classes of such weighted graphs is $\mathrm{G}(n)$, topologized so that the weight on an edge is allowed to go to 0 and that edge is deleted, provided this still yields a connected graph. If we only allow isotopies that fix the vertices, we obtain a space $\mathrm{GC}(n)$ which is something of a hybrid of $\mathrm{G}(n)$ and $\mathrm{C}(n)$, with natural maps to both of these spaces, where the map to $\mathrm{G}(n)$ is the quotient map and the map to $\mathrm{C}(n)$ assigns to a graph its set of vertices. If we further factor out translations, rotations, and rescalings we obtain $\overline{\mathrm{GC}}(n)$ with maps to $\mathrm{G}(n)$ and $\overline{\mathrm{C}}(n)$ as shown in the diagram.

The idea for defining a map $\overline{\mathrm{CP}}(n) \to \overline{\mathrm{GC}}(n)$ is to send a circle packing to its contact graph. However, weights must be assigned in order to get a continuous map, and we do this in the following way. Instead of joining two vertices at the centers of circles in a circle packing by an edge only when the circles touch, insert the edge whenever the distance between the two circles is sufficiently small, say at most one-tenth the minimum radius of any of the circles in the packing. Then assign a weight to this edge which decreases from 1 when they touch to 0 when they are at distance one-tenth the minimum radius. With these modifications to contact graphs we obtain a map $\overline{\mathrm{CP}}(n) \to \overline{\mathrm{GC}}(n)$. Composing this map with the quotient map $\overline{\mathrm{GC}}(n) \to \mathrm{G}(n)$ we obtain a map $\overline{\mathrm{CP}}(n) \to \mathrm{G}(n)$.

Now we show that each of the maps (1)–(3) in the diagram is a homotopy equivalence in turn. Commutativity of the diagram will then imply that $\overline{\mathrm{CP}}(n) \to \overline{\mathrm{C}}(n)$ is a homotopy equivalence.

(1) To each configuration in $\mathrm{CP}(n)$ we associate the Kleinian group generated by the hyperbolic reflections across the circles in the configuration. Here we are thinking of the upper half-space model for hyperbolic space $\mathbb{H}^3$, and by a reflection across a circle in $\mathbb{R}^2$ we mean the hyperbolic reflection of $\mathbb{H}^3$ across the plane in $\mathbb{H}^3$ bounded by the circle. As one varies the configuration of circles within a given stratum of $\mathrm{CP}(n)$ consisting of configurations with isotopic contact graphs, the corresponding Kleinian group varies accordingly. It follows from Ahlfors-Bers theory that, after factoring out orientation-preserving similarities of $\mathbb{R}^2$, the space of such variations is contractible, a product of Teichmüller spaces, one for each component of the complement of the union of the disks bounded by the circles C_i in the configuration. This complement is a fundamental domain for the restriction of the action to the domain of discontinuity. For the unbounded complementary component there is also the choice of the point at infinity, so one is dealing with a Teichmüller space for a surface with a marked point in this case, but it is still contractible, homeomorphic to a Euclidean space.

Consider first the case of a bounded complementary component. This can be viewed as an ideal polygon in the hyperbolic plane. The moduli space of such polygons with s

sides, modulo Möbius transformations, is homeomorphic to $\mathbb{R}^{s-3}$. As parameters one can choose the lengths of $s-3$ disjoint common perpendiculars to pairs of nonadjacent sides of a polygon. These vary independently over the interval $(0,\infty)$ and determine the polygon completely. The choice of a set of $s-3$ disjoint common perpendiculars is not unique, but there is a finite simplicial complex P_s whose vertices correspond to common perpendiculars and whose simplices correspond to sets of disjoint perpendiculars. The top-dimensional simplices of P_s then correspond to the different parametrizations of the moduli space. It is well known that P_s is a triangulation of the sphere S^{s-4}. This sphere forms a natural boundary for the moduli space, compactifying it to a closed ball. Going to a point in the boundary of this ball is achieved by letting the lengths of a set of disjoint common perpendiculars go to zero. Assuming that the circles in the given configuration that form the boundary of the ideal polygon in question are all distinct, then the pairs of circles corresponding to the shrinking common perpendiculars come together and touch as the length of each of these perpendiculars goes to zero.

For example when $s=4$ we have $P_4 = S^0$, two points corresponding to the two choices for the perpendicular joining opposite edges of a 4-gon. As the length of one perpendicular goes to 0 the length of the other goes to ∞. The four sides of the 4-gon belong to four circles of the configuration. Letting the length of a perpendicular go to 0 means that two nonadjacent circles are coming together to touch. The contact graph then changes by adding one of the two diagonals of a quadrilateral in the graph.

When $s=5$ the polyhedron P_5 is S^1 subdivided as a pentagon with vertices the five common perpendiculars and edges the pairs of disjoint common perpendiculars. Shrinking one common perpendicular turns the ideal 5-gon into an ideal triangle and an ideal 4-gon. Shrinking a second perpendicular disjoint from the first then splits the 4-gon further into two triangles. In terms of contact graphs we are adding two diagonals to a pentagon in the graph.

It can happen that two sides of a polygon belong to the same circle of a configuration. The simplest example occurs for a 5-gon as in the figure below.

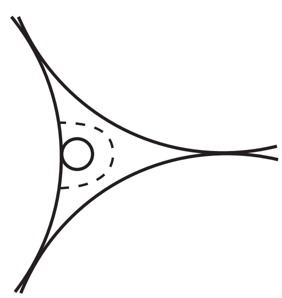

If this happens, there is a common perpendicular joining the two sides that lie in the same circle. Shrinking the length of this perpendicular to zero would have the effect of shrinking the circle or circles cut off by the perpendicular to points, which is not allowed. This means that the corresponding face of the compactified moduli space is not part of $\mathrm{CP}(n)$. For the corresponding contact graphs, shrinking this perpendicular to zero would add an edge to the graph whose two endpoints coincide, and such graphs cannot occur as contact graphs.

Contact graphs are also not allowed to have two edges joining the same pair of vertices. The situation that might lead to this is when one would attempt to shrink a perpendicular joining two sides that belong to circles that already touch, as in the first figure below, or when one would shrink two perpendiculars joining sides that belong to the same two circles as in the second figure. In both situations shrinking the perpendiculars to zero would cause one or more circles in the configuration to shrink to points.

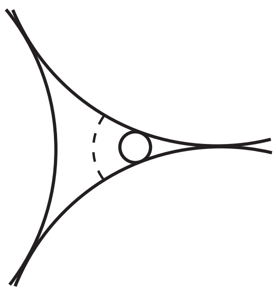

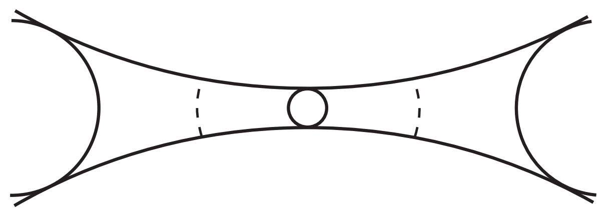

Summarizing, faces of the compactified moduli spaces are omitted when they would create contact graphs that are not simplicial graphs.

In the case of the unbounded complementary component, if this has s sides then the dimension of the moduli space is $s-1$, with the extra two dimensions coming from the location of the point at infinity within the ideal polygon, thinking of this polygon as lying in S^2 rather than $\mathbb{R}^2$. Parameters are the lengths of disjoint common perpendiculars between nonadjacent sides and also common perpendiculars between adjacent sides or from a side to itself.

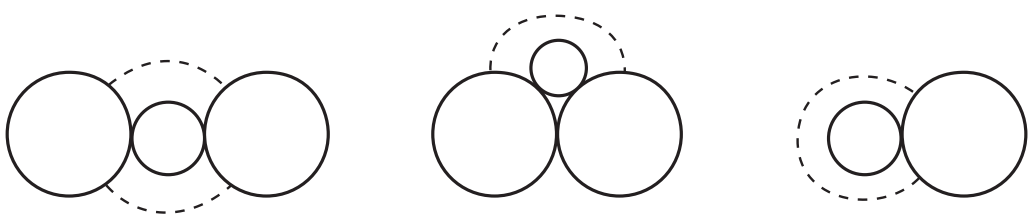

There is again a polyhedron whose top simplices correspond to these parametrizations. This polyhedron is a sphere S^{s-2} compactifying the moduli space to form a closed $(s-1)$-ball. There are always faces of this ball that cannot be realized by deformations of the circle configurations, deformations that would produce inadmissible contact graphs with edges forming loops or edge pairs sharing both endpoints.

.... A lot needs to be added here to complete the argument for part (1) ....

(2) For this part of the proof we will replace $\mathrm{G}(n)$ and $\overline{\mathrm{GC}}(n)$ by their *spines*, which are subspaces to which they deformation retract. In the case of $\mathrm{G}(n)$ the spine is a finite simplicial complex, the geometric realization of the partially ordered set consisting of isotopy classes of connected graphs in $\mathbb{R}^2$ with n vertices labeled $1, \cdots, n$, under the partial ordering given by inclusion. The deformation retraction to this spine can be achieved in two steps. First, a collection of weights on edges can be rescaled so that it sums to 1. The normalized weights can then be viewed as barycentric coordinates. Each graph thus determines a simplex, with certain faces omitted whenever deleting the corresponding set of edges from the graph produces a disconnected subgraph. What remains of the simplex deformation retracts onto the subcomplex of the barycentric subdivision corresponding to the partially ordered set of connected subgraphs.

A similar construction works for $\overline{\mathrm{GC}}(n)$, producing a spine which is a (semi-)simplicial space rather than a simplicial complex. For notational convenience we use the same symbols $\mathrm{G}(n)$ and $\overline{\mathrm{GC}}(n)$ to denote these spines.

The map $\overline{\mathrm{GC}}(n) \to \mathrm{G}(n)$ is a simplicial map from a simplicial space to a simplicial complex, so it will be a homotopy equivalence if its fibers are contractible. There is one fiber F for each isotopy class of graphs (weights are fixed in a fiber and can be ignored), where the equivalence relation on graphs in F is isotopy fixing the vertices. The graph can be assumed to be a maximal tree since specifying the isotopy classes of any edges not in a maximal tree results in a homeomorphic space F.

The projection $F \to \overline{\mathrm{C}}(n)$ sending an embedded tree to its vertices is a covering space. We will show this is the universal cover, and this will imply that F is contractible since $\overline{\mathrm{C}}(n)$ is a $K(\pi, 1)$.

The group $\pi_1 \mathrm{C}(n) = \pi_1 \overline{\mathrm{C}}(n) \times \mathbb{Z}$ is the pure braid group P_n. We can also view P_n as the mapping class group of a disk with n punctures, where the punctures are not allowed to be permuted and diffeomorphisms restrict to the identity on the boundary of the disk. Dropping this last condition gives the quotient group $\pi_1 \overline{\mathrm{C}}(n)$. This 'relaxed' mapping class group acts on F by taking trees to their images under diffeomorphisms fixing the punctures. This is a faithful action since if a diffeomorphism fixes a tree, it is isotopic to the identity fixing the punctures but not fixing the boundary of the disk. The action of $\pi_1 \overline{\mathrm{C}}(n)$ on a fiber of the covering space by lifting loops is the same as the action by the mapping class group. Since this is a faithful action, elementary covering space theory implies that the covering space is the universal cover.

(3) The projection $\overline{\mathrm{GC}}(n) \to \overline{\mathrm{C}}(n)$ is a fiber bundle whose fiber consists of isotopy classes of weighted graphs with a fixed vertex set. Let us denote this fiber by $\mathrm{X}(n)$. It will suffice to prove that $\mathrm{X}(n)$ is contractible. We will show this by enlarging $\mathrm{X}(n)$ to a space $\mathrm{Y}(n)$ whose contractibility is more easily seen, then we will show that the relative groups $\pi_i(\mathrm{Y}(n), \mathrm{X}(n))$ are zero, hence $\pi_i \mathrm{X}(n) \approx \pi_i \mathrm{Y}(n) = 0$ for all i, which implies that $\mathrm{X}(n)$ is

contractible.

The space $\mathrm{Y}(n)$ has the same definition as $\mathrm{X}(n)$ except that graphs are allowed to have more than one edge joining the same pair of endpoints, provided that the disk bounded by any two such edges contains at least one vertex in its interior. We use the term *multigraph* for such a graph.

To show $\mathrm{Y}(n)$ is contractible we use the surgery process from [H] which gives a path in $\mathrm{Y}(n)$ from an arbitrary graph to a fixed maximal graph. First choose representatives of the isotopy classes of these two graphs that intersect transversely in the minimum number of points (apart from vertices). These intersection points are then eliminated one by one by cutting each edge of the arbitrary graph wherever it meets the fixed graph and pushing the newly created pair of endpoints to the vertices along the fixed edges, as in the figure below. Note that the two new edges can be isotoped to be disjoint from the old one and all the other edges in the given graph.

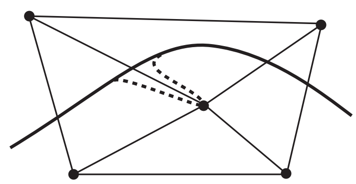

Such a surgery is realized by a continuous path in $\mathrm{Y}(n)$ by gradually shifting the weight on the edge being surgered to each edge in the new pair of edges. Note that each surgery preserves the connectedness of the graph. During the surgery process pairs of isotopic edges may be created, and these are to be replaced by a single edge, adding the two weights. Edges connecting a vertex to itself may also be created, but these can simply be discarded without affecting connectedness of the graph. To make the surgery process well defined, choose orientations for the edges of the fixed graph and surger in the direction indicated by the orientation. Also, choose an ordering for the edges of the fixed graph, and first surger to eliminate all intersections with the first edge, then all intersections with the second edge, and so on. When the surgery process is finished the arbitrary given graph is replaced by a graph disjoint from the fixed graph, and hence isotopic to a subgraph of it since the fixed graph was chosen to be maximal. To complete the deformation, weights are deformed linearly from those on the subgraph to those on the fixed graph. This surgery process in fact depends continuously on the arbitrary given graph; see [H] for details.

If the surgery process stayed within $\mathrm{X}(n)$ we would be done, but unfortunately it does not. Thus we need a procedure for eliminating pairs of edges joining the same two vertices. Simply discarding all such pairs of edges does not work since this might produce a disconnected graph. Discarding only one edge of a pair cannot be done in a way that depends continuously on the given graph. This is illustrated in the following figure, where the three graphs across the top show a path in $\mathrm{Y}(n)$ with ends in $\mathrm{X}(n)$, in which a continuous choice of which edge to discard cannot be made. What we will do instead is deform this path to one where discarding both edges of the pair does not disconnect the graph, as shown in the lower part of the figure.

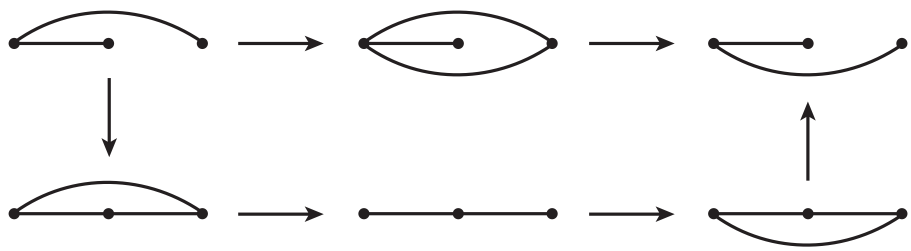

A pair of edges joining two vertices p and q bounds a disk D. Let $\mathrm{X}'(k)$ be the space of isotopy classes of weighted graphs in D with k fixed vertices in the interior of D in addition to the two vertices p and q in ∂D, such that the graphs have at most two components, the components containing p and q. We further specify that no edges join p directly to q. The subspace of $\mathrm{X}'(k)$ consisting of graphs with only one component is denoted $\mathrm{X}_1'(k)$. Allowing multigraphs instead of just graphs results in a larger space $\mathrm{Y}'(k) \supset \mathrm{X}'(k)$.

Lemma. *Every map $K \to \mathrm{X}'(k)$ from a finite polyhedron K to $\mathrm{X}'(k)$ is homotopic to a map to the subspace $\mathrm{X}_1'(k)$.*

Proof. For the given family of graphs in $\mathrm{X}'(k)$ parametrized by K we perform the surgery process, deforming this family to a constant family consisting of a fixed maximal graph in $\mathrm{X}_1'(k)$. This deformation takes place in the space $\mathrm{Y}'(k)$ provided that we orient the edges of the fixed target graph toward vertices in the interior of the disk D, so that no edges joining the two boundary vertices are created. It will suffice to deform the resulting map $K \times I \to \mathrm{Y}'(k)$ into $\mathrm{X}'(k)$ staying fixed over $K \times \partial I$ where it already maps to $\mathrm{X}'(k)$. We will show that this is possible more generally for any map of a finite polyhedron L to $\mathrm{Y}'(k)$.

We may assume the given map $f : L \to \mathrm{Y}'(k)$ is PL in the sense that the weights on edges are PL functions on L after we extend these functions to be 0 where the edges are not part of the multigraph. (One way to do this would be to replace $\mathrm{Y}'(k)$ by its spine, a simplicial complex, and take a PL map to this simplicial complex.) If the multigraphs in the family defined by f do not consist entirely of graphs in $\mathrm{X}'(k)$, choose a pair of edges a and b joining the same two vertices such that the number m of vertices in the interior of the disk bounded by $a \cup b$ is minimal for the family. These two edges live over an open set in the domain L with closure a subpolyhedron A containing a subpolyhedron $A_0 \subset A$ where at least one of the weights on a and b is zero. Thus we have a map $g : A \to \mathrm{X}'(m)$ by restricting just to vertices and edges in the disk bounded by $a \cup b$. By induction on k

there is a homotopy g_t of g to a map with image in $\mathrm{X}_1'(m)$. The induction can start with the case $k=1$ where the spine of $\mathrm{X}'(1)$ is homeomorphic to a closed interval with $\mathrm{X}_1'(1)$ as its midpoint.

Now we perform three successive homotopies of f:

(1) By a homotopy supported in a neighborhood of A_0 in A we can reduce the weight on the edge a to 0 near the part of A_0 where this weight is 0 and similarly for b. Note that this does not destroy connectedness of graphs since any edges with nonzero weights over A_0 have nonzero weights over a neighborhood of A_0.

(2) If we damp down the homotopy g_t in a smaller neighborhood of A_0 by taking shorter and shorter initial segments as we approach A_0, we can extend this damped down homotopy to a homotopy of f which is constant outside A.

(3) Now outside this smaller neighborhood of A_0 in A we deform the f produced in the previous step by letting the weights on both a and b go to zero. This preserves connectedness of graphs since g_1 has image in $\mathrm{X}_1'(m)$. The result is a new f in which a and b are never simultaneously edges over A.

These steps give a deformation of f which eliminates the pair of edges a and b from the finitely many pairs of edges joining the same endpoints, without introducing any new such pairs. After finitely many repetitions of this process we eventually obtain a homotopy of f into $\mathrm{X}'(k)$. This homotopy is fixed where f already maps to $X'(k)$. This finishes the proof of the lemma. $\square$

The same arguments applied to a map $(D^i, \partial D^i) \to (\mathrm{Y}(n), \mathrm{X}(n))$ give a homotopy to a map with image in $\mathrm{X}(n)$, staying fixed over ∂D^i. This shows that $\pi_i(\mathrm{Y}(n), \mathrm{X}(n)) = 0$, as desired.

References

[H] Allen Hatcher, On triangulations of surfaces. Topology Appl. 40 (1991) 189-194. For an updated version see the author's webpage.