Simplifying and Separating Configurations of Disjoint Unlinked Spheres in Euclidean Space

Allen Hatcher

Consider the space $\mathcal{C} = \mathcal{C}(n, k; d_1, \cdots, d_k)$ of all configurations of k disjoint labeled Euclidean spheres $S_1, \cdots, S_k$ in $\mathbb{R}^n$ of dimensions $d_1, \cdots, d_k$ which are unlinked in the sense that the hemispheres they bound in the upper half-space of $\mathbb{R}^{n+1}$ bounded by $\mathbb{R}^n$ are all disjoint. For circles in $\mathbb{R}^3$ this is equivalent to the usual notion of unlinked. Since the hemispheres are uniquely determined by their boundaries, this means that $\mathcal{C}$ can also be regarded as the configuration space of disjoint labeled hemispheres of specified dimensions in Euclidean half-space with boundaries on the boundary of the half-space. Another viewpoint is as a configuration space of disjoint hyperbolic planes of specified dimensions in hyperbolic n-space.

Configurations in $\mathcal{C}$ can be quite complicated, and the goal of this paper is to provide two ways to simplify them, as a preliminary step to understanding the homotopy type of $\mathcal{C}$. The first simplification method generalizes a technique in [1] which treated the special case of circles in $\mathbb{R}^3$. This involves associating a non-negative real number to each configuration that gives a measure of its complexity, and then proving that $\mathcal{C}$ deformation retracts onto the subspace of configurations of complexity less than a given number $c > 0$. The arguments are a straightforward generalization of those in [1].

The second simplification method involves deforming each configuration in $\mathcal{C}$ so that its spheres can be separated by a collection of codimension one spheres disjoint from each other and from the spheres in the configuration. A space $\mathcal{SC}$ of separated configurations will be defined, and it will be shown that the projection $\mathcal{SC} \to \mathcal{C}$ forgetting the separating spheres is a homotopy equivalence. This answers a question asked by James Griffin who needed such a result for a project to compute the homology of $\mathcal{C}$.

The two simplification methods work equally well for the subspace $\mathcal{U}$ of $\mathcal{C}$ consisting of "untwisted" configurations in which each sphere of dimension d_i lies in a $(d_i + 1)$-plane parallel to the standard $\mathbb{R}^{d_i+1} \subset \mathbb{R}^n$. By sending each sphere to the plane containing it, translated to the origin, we obtain a natural projection from $\mathcal{C}$ to a product of Grassmann manifolds $G(n, d_i + 1)$ of $(d_i + 1)$-planes in $\mathbb{R}^n$, with the fiber of this projection being $\mathcal{U}$:

$$\mathcal{U} \longrightarrow \mathcal{C} \longrightarrow G(n, d_1 + 1) \times \cdots \times G(n, d_k + 1)$$

It is not clear whether this projection is a fibration, but we will show that it is at least a quasifibration. It has a section given by placing the various spheres inside disjoint n-balls in $\mathbb{R}^n$ and taking all possible rotations of the spheres within these balls.

One might hope that the quasifibration is at least homologically a product, but this is not generally true since the inclusion of the fiber does not always induce an injection on homology. For example in the case of circles in $\mathbb{R}^3$ the map $H_1(\mathcal{U}) \to H_1(\mathcal{C})$ has nontrivial kernel when $k \geq 2$.

Reducing Complexity

First we define a notion of "complexity" for configurations in $\mathcal{C}$ generalizing the one used in [BH] for configurations of unlinked circles in $\mathbb{R}^3$. For each sphere S_i in a configuration in $\mathcal{C}$ let its *hull* be the closed n-dimensional ball in $\mathbb{R}^n$ having the same center as S_i and twice the radius. Define the *complexity* of a configuration $(S_1, \cdots, S_k)$ in $\mathcal{C}$ to be the maximum ratio r_i/r_j of the radii of pairs of distinct spheres S_i and S_j whose hulls intersect, with $r_i \leq r_j$. If no two hulls intersect then the complexity is defined to be zero.

Configurations of low complexity thus have the property that whenever two spheres in the configuration are near each other, one sphere is much smaller than the other one. One can think of this in terms of astronomy, with the spheres being regarded as celestial bodies, so that low complexity means that the spheres form systems of planets orbiting a much larger sun, with much smaller moons orbiting the planets, and so on. Thus one has an ordered hierarchy of spheres.

For a constant $c > 0$ let $\mathcal{C}^c$ be the subspace of $\mathcal{C}$ consisting of configurations of complexity less than c. Define the subspace $\mathcal{U}^c$ of $\mathcal{U}$ similarly.

Theorem 1.1. *For each $c > 0$ the inclusions $\mathcal{C}^c \hookrightarrow \mathcal{C}$ and $\mathcal{U}^c \hookrightarrow \mathcal{U}$ are weak homotopy equivalences.*

Proof. The same method will work for both $\mathcal{C}$ and $\mathcal{U}$. The starting point is the idea of deforming a configuration $(S_1, \cdots, S_k)$ in $\mathcal{C}$ by the *canonical shrinking* introduced by Freedman and Skora in [2] which we now describe. By the definition of $\mathcal{C}$ the spheres S_i bound disjoint hemispheres in the upper half of $\mathbb{R}^{n+1}$. The canonical shrinking is obtained by intersecting these hemispheres with the hyperplanes $x_{n+1} = t$ for t increasing from 0 to infinity. Thus each sphere S_i shrinks through concentric spheres until it degenerates to a point and disappears. We do not want it to disappear entirely so we will have to modify the canonical shrinking to achieve this.

Before doing this we should check that the configurations of spheres produced during the canonical shrinking still lie in $\mathcal{C}$, so the hemispheres they bound remain disjoint.

Suppose on the contrary that two of these hemispheres intersect at some interior point. (The spheres themselves remain disjoint since the original hemispheres were disjoint.) Choose two vertical two-dimensional planes in $\mathbb{R}^{n+1}$ containing the intersection point and intersecting the two original hemispheres in semicircles. Putting the center of one of these semicircles at the origin in $\mathbb{R}^2$ we would have a configuration like that in the figure below, with the semicircle the upper half of the circle $x^2 + y^2 = r^2$ for some r.

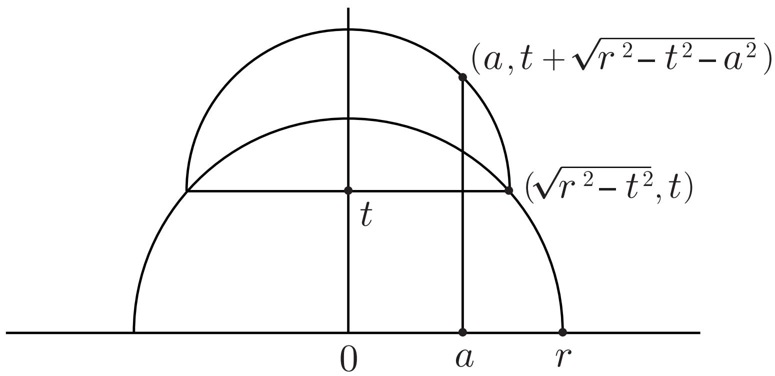

At time t during the canonical shrinking the new hemisphere would intersect the plane in a semicircle with base at height $y = t$, the upper half of the circle $x^2 + (y - t)^2 = r^2 - t^2$. A point on this upper semicircle has coordinates $(a, t + \sqrt{r^2 - t^2 - a^2})$. Similarly, in the other vertical plane a point on the upper semicircle has coordinates $(A, t+\sqrt{R^2 - t^2 - A^2})$. We claim that if $\sqrt{r^2 - t^2 - a^2} < \sqrt{R^2 - t^2 - A^2}$ when $t = 0$ then this inequality holds for all t. But this is obvious just by squaring both sides. Thus if the two upper semicircles are disjoint when $t = 0$ they will be disjoint for all t, so the hemispheres stay disjoint during the canonical shrinking.

Returning now to the overall plan, the main difficulty we face is that if we keep the radius of a sphere S_i a small positive number rather than letting it go all the way to zero and disappear, other shrinking spheres S_j may bump into it and we may have to move it out of the way to avoid such collisions. The challenge is to do this in a way that varies continuously with the original configuration.

Ignoring the problem of disappearing spheres, note that the canonical shrinking decreases complexity monotonically (in the weak sense) since the radii of smaller spheres decrease faster than the radii of larger spheres, and hulls that initially intersect can become disjoint as the shrinking progresses, but they cannot suddenly begin to intersect when they were disjoint before. As a very special case, if all the spheres S_i have the same radius then they shrink until their hulls are disjoint so the complexity decreases to zero before the spheres disappear.

There will be two main steps to the process of modifying the canonical shrinking. In the first step we show how to deal with a single initial configuration, and in the second step we show how to combine these modified shrinkings so that they vary continuously with the initial configuration.

Step 1. Consider a configuration $(S_1, \cdots, S_k)$ in $\mathcal{C}$. Let X_1 be the union of the spheres S_i of largest radius, X_2 the union of the spheres of next largest radius, and so on for $X_3, X_4, \cdots$. Let $Y_i = X_1 \cup \cdots \cup X_i$. The time parameter in the isotopies we construct will be denoted by u with $0 \leq u < \infty$. Let u_i be the time when the canonical shrinking reduces the spheres in X_i to a point, so $u_1 > u_2 > \cdots$.

By induction on i we will construct an isotopy Φ_u^i of Y_i that decreases the complexity of the configuration Y_i so that it approaches zero as u goes to infinity. To begin we let Φ_u^1 be the canonical shrinking applied to the spheres in X_1 for $u \leq u_1' < u_1$ where u_1' is chosen close enough to u_1 so that the hulls of all the spheres in $\Phi_{u_1'}^1(X_1)$ are disjoint. Then for $u \geq u_1'$ we slow down the canonical shrinking of the spheres in X_1 so that their radius $r_1(u)$ stays positive for all u. Having defined Φ_u^1 on X_1 we extend Φ_u^1 to an ambient isotopy $\overline{\Phi}_u^1$ of $\mathbb{R}^n$ in the standard way. Namely, the isotopy gives a tangent vector field to $\Phi^1(X_1 \times [0, \infty))$ in $\mathbb{R}^n \times [0, \infty)$ whose second coordinate is 1 and we extend this to a smooth vector field on $\mathbb{R}^n \times [0, \infty)$ with the same property and whose first coordinate vanishes outside a small neighborhood of $\Phi^1(X_1 \times [0, \infty))$. The flow associated to this extended vector field gives the extended isotopy $\overline{\Phi}_u^1$ on $\mathbb{R}^n$.

Next we extend Φ_u^1 to an isotopy Φ_u^2 defined on the spheres in X_2 as well as those in X_1. For a sphere in X_2 we let Φ_u^2 be the canonical shrinking for $u \leq u_2' < u_2$ for u_2' close enough to u_2 so that the hulls of the spheres in $\Phi_{u_2'}^2(X_2)$ are disjoint. For $u > u_2'$ we will specify the position of a sphere in $\Phi_u^2(X_2)$ by specifying three things: its radius $r_2(u)$, its center, and the $(d+1)$-dimensional plane containing the sphere, where d is its dimension. The center of a sphere in X_2 is constant for $u \leq u_2$ and thereafter we let it move by the ambient isotopy $\overline{\Phi}_u^1$ of $\mathbb{R}^n$. The centers of different spheres in $\Phi_{u_2}^2(X_2)$ are distinct so they remain distinct as they move under this ambient isotopy. The planes containing spheres in X_2 and passing through their centers can be chosen to vary in any smooth way for $u \geq u_2$ as the center moves under the ambient isotopy. For example we could let them move parallel to themselves. This will guarantee that configurations in $\mathcal{U}$ remain in $\mathcal{U}$. Finally, we choose the smooth radius function $r_2(u)$ to be small enough so that the ratio $r_2(u)/r_1(u)$ approaches zero monotonically as u increases and so that the hulls of the spheres in X_2 remain disjoint for all $u \geq u_2$. Having defined Φ_u^2 on Y_2, we then extend Φ_u^2 to an ambient isotopy $\overline{\Phi}_u^2$ of $\mathbb{R}^n$ as before.

The construction of subsequent isotopies Φ_u^i of Y_i is done inductively by the same procedure used to construct Φ_u^2 from Φ_u^1. We choose the radius functions $r_i(u)$ small enough so that the ratios $r_i(u)/r_j(u)$ for $i > j$ go to zero monotonically with increasing u and so that $r_i(u)/r_j(u) \ll r_{i'}(u)/r_{j'}(u)$ for $u \geq u_i$ and for all $j < i$ and $j' < i' < i$.

From the construction of the isotopy it is clear that the complexity becomes small for large u since the hulls of spheres in $\Phi_u^i(X_i)$ are disjoint for $u \geq u_i$ while for spheres in X_i and X_j for $i > j$ the ratio $r_i(u)/r_j(u)$ approaches zero as u increases.

We claim that complexity decreases monotonically during the isotopy. Since the ratios $r_i(u)/r_j(u)$ for $i > j$ decrease monotonically, the only way that complexity could increase would be if the hulls of two spheres intersect for some u value after having been disjoint for slightly smaller u values. This does not happen under the canonical shrinking nor can it happen if both spheres belong to the same X_i by the way we choose $r_i(u)$. Suppose it happens for spheres S_p in X_i and S_q in X_j with $i > j$, so S_p is smaller than S_q. At the time the two hulls begin to intersect, the center of at least one of S_p and S_q must be moving. This motion is caused by the center being 'pushed' by a sphere $S_{q'}$ in $X_{j'}$ as a result of the isotopy extension process. Here $j' < i$ if S_p is being pushed and $j' < j$ if S_q is being pushed. The pushing takes place in a small neighborhood of $S_{q'}$, and thus in the interior of the hull of $S_{q'}$. We cannot have $S_{q'} = S_q$ since if S_q was pushing S_p their hulls would already be intersecting contrary to our assumption that they are coming into contact after being disjoint. If $S_{q'}$ is pushing S_q then their hulls must intersect so since both $r_j(u)/r_{j'}(u)$ and $r_{j'}(u)/r_j(u)$ are greater than $r_i(u)/r_j(u)$, the collision of the hulls of S_p and S_q has no effect on complexity. On the other hand, if $S_{q'}$ is pushing S_p then the hull of $S_{q'}$, which is much larger than S_p, must intersect S_q and hence also the hull of S_q during the collision, so in this case too the complexity is unaffected by the collision. This finishes the argument that complexity decreases monotonically.

Step 2. We now show that the relative homotopy groups $\pi_l(\mathcal{C}, \mathcal{C}^c)$ are zero for each $l \geq 0$ by doing a parametrized version of the shrinking isotopy. Note first that $\mathcal{C}^c$ is an open subset of $\mathcal{C}$ since complexity is upper semicontinuous, meaning that small perturbations of a given configuration in $\mathcal{C}$ cannot produce large increases in the complexity, although they can produce large decreases when two hulls that just touch are perturbed to be disjoint.

An element of $\pi_l(\mathcal{C}, \mathcal{C}^c)$ is represented by a map $f : (D^l, \partial D^l) \to (\mathcal{C}, \mathcal{C}^c)$ of the form $f(t) = (S_1(t), \cdots, S_k(t))$. After a small homotopy of f we can assume that the radii of the spheres $S_i(t)$ are piecewise linear functions of t, so they are linear on the simplices of some triangulation of D^l. Here we are using the fact that $\mathcal{C}^c$ is open in $\mathcal{C}$ to assure that if the homotopy is small enough, the configurations over ∂D^l will still lie in $\mathcal{C}^c$. After a suitable subdivision of the triangulation of D^l the ordering of the spheres $S_i(t)$ by size will be constant on the interior of each simplex of D^l, and when one passes to the boundary of a simplex all that happens to this ordering is that some inequalities among radii become equalities.

For a fixed $t \in D^l$ we showed in Step 1 how to construct an isotopy Φ_{tu} of the configuration $(S_1(t), \cdots, S_k(t))$ decreasing its complexity monotonically. The construction makes strong use of the ordering of the spheres $S_i(t)$ by size which allowed us to construct Φ_{tu} as the final isotopy in a finite sequence of isotopies Φ^i_{tu}. Over an open simplex σ of the subdivision of D^l where the size ordering is fixed we can arrange that Φ_{tu} depends continuously on $t \in \sigma$. To see this, observe that the numbers u_i are continuous functions of t and we can choose the radius functions $r_i(t, u)$ to vary continuously with t. The isotopy extensions $\overline{\Phi}^i_{tu}$ can be chosen continuously since they just depend on extending certain vector fields.

The construction of Φ_{tu} over σ can be extended continuously over a neighborhood of σ in D^l by using the same partition of the configurations $(S_1(t), \cdots, S_k(t))$ into the sets X_i throughout this neighborhood. Previously the radius functions $r_i(u)$ were the same for all spheres in X_i, but now these spheres can have different sizes so in order for the size ordering of the spheres in $\Phi_{tu}(X_i)$ to be independent of u we will need different radius functions for different spheres. One way to choose these is to start with a radius function $r_i(t, u)$ for $t \in \sigma$ and extend this to a neighborhood of σ via functions of the form $r_i(p(t), u - \rho(t))$ where p projects the neighborhood onto σ and $\rho(t)$ shifts the u-parameter to make $u_i(p(t)) + \rho(t)$ the value of u when the sphere shrinks to a point under the canonical shrinking.

Now we construct inductively a sequence of deformations of the given map $f : D^l \to \mathcal{C}$ to make its image lie in $\mathcal{C}^c$, staying in $\mathcal{C}^c$ over ∂D^l where it already lies in $\mathcal{C}^c$. Suppose inductively that we have already deformed f to take a neighborhood N_{i-1} of the $(i-1)$-skeleton of D^l into $\mathcal{C}^c$. For an i-simplex σ let σ' be a slightly shrunk copy of σ contained in the interior of σ and with $\partial\sigma'$ contained in N_{i-1}. Choose a function $\psi : D^l \to [0, \infty)$ whose support is contained in a neighborhood of σ' where the deformation Φ_{tu} from the preceding paragraph is defined and whose values are large enough on a smaller neighborhood of σ' so that if we perform the deformation Φ_{tu} for $u \le \psi(t)$ then this deforms f to have image in $\mathcal{C}^c$ in the union of N_{i-1} with a neighborhood of σ. After repeating this for the other i-simplices we then proceed inductively to the higher dimensional simplices. ☐

In the proof of Theorem 1.1 we showed how to construct a homotopy of $f : D^l \to \mathcal{C}$ decreasing the complexity to be as small as we like. The construction can be made so that it yields a family with an additional *buffer property*: the hull of each sphere is disjoint from all other spheres in the configuration of equal or larger radius. For a single initial configuration this holds if the radius functions $r_i(u)$ are chosen small enough. In the parametrized setting we used damping functions $\psi(t)$ to fit together the various shrinking isotopies, and the point to observe is that if one starts with a configuration satisfying the buffer property, then this continues to hold during the shrinking isotopy provided that the radius functions $r_i(t,u)$ are chosen small enough.

Let us now show how to obtain the quasifibration described in the introduction:

$$\mathcal{U} \to \mathcal{C} \to G(n, d_1+1) \times \cdots \times G(n, d_k+1)$$

Proposition 1.2. *The projection* $\mathcal{C} \to G(n, d_1+1) \times \cdots \times G(n, d_k+1)$ *is a quasifibration.*

Proof. To show the quasifibration property it suffices to verify the homotopy lifting property for homotopies $f_t : D^l \to G(n, d_1+1) \times \cdots \times G(n, d_k+1)$ that are stationary for $t \in [0, \varepsilon]$ for some $\varepsilon > 0$, with a given lifting $\widetilde{f}_0$ for $t = 0$. The idea is to perform the complexity-reducing deformation of $\widetilde{f}_0$ in the proof of Theorem 1.1, compressed into the t interval $[0, \varepsilon]$, to obtain configurations whose spheres can then be rotated according to the homotopy f_t.

In the shrinking process of Theorem 1.1 we used isotopy extension to determine where the centers of spheres move, and once the spheres were small we could specify their radii and we could choose deformations of their spanning planes. In the present situation we choose these deformations so that the spanning planes are fixed or move parallel to themselves during the shrinking process until the shrinking has been completed for all the spheres for all parameter values in D^l, say for $u \le u_0$. We compress the time interval $[0, u_0]$ to $[0, \varepsilon]$. Then for $t \ge \varepsilon$ we continue the shrinking process beyond $u = u_0$ and use the rotations of the spanning planes specified by the homotopy f_t. □

Separating Spheres

There is another way of avoiding complicated configurations of spheres that is expressed in terms of separation by codimension-one spheres. Let us say that a system of codimension-one spheres $E_1, \cdots, E_l$ in $\mathbb{R}^n$ *separates* a configuration $(S_1, \cdots, S_k)$ in $\mathcal{C}$ if the spheres E_j are disjoint from each other and from each S_i and there is at most one S_i in each component of $\mathbb{R}^n - \cup_j E_j$, which is equivalent to saying that each pair of spheres S_i is separated by some E_j. We also require that for each E_j there is at least one sphere S_i in each component of $\mathbb{R}^n - E_j$. Define a space $\mathcal{SC}$ of separated sphere configurations to consist of all configurations $(S_1, \cdots, S_k)$ in $\mathcal{C}$ together with unordered systems $E_1, \cdots, E_l$ of separating spheres for $(S_1, \cdots, S_k)$ and weights $w_1, \cdots, w_l$ in $(0, 1]$ associated to the spheres $E_1, \cdots, E_l$. We allow the number of separating spheres to vary by allowing spheres E_j to be deleted when their weights go to zero, provided that the remaining E-spheres with nonzero weights still separate $(S_1, \cdots, S_k)$. To define the topology on $\mathcal{SC}$ we start with the configuration spaces of systems with a fixed number l of separating spheres E_j and corresponding weights w_j in $[0,1]$, with the requirement that the E-spheres with nonzero weights separate the S-spheres. For a fixed number m we then take the quotient space of the disjoint union of these configurations spaces with $l \le m$ separating spheres, with identifications corresponding to deleting spheres of weight zero. Then we take the direct limit as the number m goes to infinity.

In the definition of $\mathcal{SC}$ we required each sphere E_j to have at least one sphere S_i on each side of it. Omitting this condition replaces $\mathcal{SC}$ by a larger space that deformation retracts to it by letting the weights of the spheres E_j not satisfying this condition go to zero. The condition is thus fairly harmless, and has the virtue of excluding "separating spheres" that do not actually separate the spheres S_i.

As a slight generalization, we can choose some subset A of the indices $1, \cdots, k$ and only require that the spheres E_j separate the spheres S_i for $i \in A$. For simplicity we use the same notation $\mathcal{SC}$ for this generalization.

Replacing $\mathcal{C}$ by the subspace $\mathcal{U}$ of untwisted sphere configurations we can then have an analogous space $\mathcal{SU}$ of separated untwisted configurations, which is just $\mathcal{U} \cap \mathcal{SC}$.

Theorem 1.3. *The projections $\mathcal{SC} \to \mathcal{C}$ and $\mathcal{SU} \to \mathcal{U}$ forgetting the separating spheres and their weights are weak homotopy equivalences.*

Proof. To show that the projection $p\colon \mathcal{SC} \to \mathcal{C}$ induces isomorphisms on all homotopy groups π_l it suffices to show the following: Given a map $F\colon \partial D^l \to \mathcal{SC}$ and a map $f\colon D^l \to \mathcal{C}$ which equals pF on ∂D^l then there exist homotopies $F_t\colon \partial D^l \to \mathcal{SC}$ of $F = F_0$ and $f_t\colon D^l \to \mathcal{C}$ of $f = f_0$ with $pF_t = f_t$ for all t, such that F_1 extends to a map $F_1\colon D^l \to \mathcal{SC}$ with $pF_1 = f_1$ on D^l. There will be four steps to the construction. The first two steps produce homotopies f_t and F_t and the second two arrange the desired compatibility over ∂D^l.

Step 1. Deforming the given $f\colon D^l \to \mathcal{C}$ as in the proof of Theorem 1.1 gives a homotopy f_t of $f = f_0$ whose final family f_1 has small complexity and satisfies the buffer property. For a fixed configuration $f_1(x)$ in the family f_1 consider the part of the hull of a sphere S_i of $f_1(x)$ consisting of codimension one spheres concentric with the boundary sphere of the hull and lying between S_i and the boundary sphere. We call these spheres *shell spheres*. If the complexity is small compared to the number of spheres in $f_1(x)$ then among the shell spheres for S_i there will always be some that are disjoint from the hulls of all smaller spheres S_j. These shell spheres will also be disjoint from all the spheres S_j of equal or larger radius by the buffer property. If we choose at least one of these shell spheres for each S_i with $i \in A$ then we claim that this gives a separating system of E-spheres for the given configuration. To see this, observe first that these shell spheres are disjoint from the spheres S_i by construction, and they are disjoint from each other since for each pair of spheres $S_i \neq S_j$, if these two spheres have the same size then their hulls are disjoint, and if one is larger than the other, say S_i is larger than S_j, then the shell sphere for S_i was chosen disjoint from the hull of S_j and hence also from all shell spheres for S_j. Finally, observe that the chosen shell spheres separate all the spheres S_i with $i \in A$ since the shell spheres for S_i separate S_i from all S_j's of equal or larger size.

This was for a fixed configuration $f_1(x)$ in the family f_1. For small perturbations of x we can vary the shell spheres continuously giving new separating systems of shell spheres for the perturbed configurations. Thus we can choose functions $\phi_x : D^l \to [0,1]$ supported in a neighborhood of x in which the shell spheres vary continuously, and these ϕ_x's can be taken to be PL in some triangulation of D^l. By compactness a finite number of the interiors of their supports cover D^l. Relabeling these ϕ_x's as ϕ_j, we can use them as weights for the corresponding shell spheres. As the configurations $f_1(x)$ vary with x, some of the chosen shell spheres for a given S_i might merge, but this can be avoided by specifying that the radii of the chosen shell spheres are always rational multiples of the radii of the associated spheres S_i. This guarantees that these ratios are locally constant since the weight functions ϕ_j were chosen to be PL, with locally connected support sets. Thus we have a map $F_1 : D^l \to \mathcal{SC}$ with $pF_1 = f_1$.

Step 2. First we deform the given map $F : \partial D^l \to \mathcal{SC}$ so that the weights of the E-spheres of $F(x)$ are PL functions of x in the following way. For each $x \in \partial D^l$ choose a PL function $\omega_x : \partial D^l \to [0,1]$ that is positive at x and vanishes outside a neighborhood of x in which all the E-spheres in $F(x)$ have positive weights. We regard ω_x as assigning a new weight to the E-spheres in $F(x)$. As x varies over ∂D^l the interiors of the supports of the functions ω_x cover ∂D^l, and by compactness this cover has a finite subcover. Relabeling the resulting finite set of ω_x's as ω_j, we obtain a new system of weighted E-spheres by taking the spheres corresponding to each ω_j and summing the weights ω_j when the same sphere occurs for more than one ω_j. For each $x \in \partial D^l$ the new E-spheres with positive weights give a subsystem of the original E-spheres which is still a separating system. Letting the old weights go to zero while the new weights go from zero to $\sum_j \omega_j$ gives a homotopy of F to a new map with PL weight functions without changing the map $f = pF$.

After this preliminary deformation of F we now want to apply the shrinking process in the preceding theorem to the union of the spheres in the systems $f(x)$ and the E-spheres in $F(x)$. This certainly works for a fixed value of x, but as x varies over ∂D^l some E spheres appear and disappear as their weights go to zero, so we need to modify the shrinking process to take this into account. Since the weight functions are PL we can triangulate ∂D^l so that over each open simplex of the triangulation there is a fixed set of E-spheres varying only by isotopy. Some of these can disappear at the boundary of the simplex if their weights become zero.

A handle structure on ∂D^l can be associated to the triangulation of ∂D^l in the usual way, where 0-handles are neighborhoods of vertices of the triangulation and i-handles are attached inductively to a neighborhood of the $(i-1)$-skeleton to form a neighborhood of the i-skeleton. We will construct the deformation F_t as a sequence of deformations over neighborhoods of the i-handles by downward induction on i.

The first step is to deform F by applying the shrinking process over a neighborhood of the $(l-1)$-handles, damping this deformation down to zero near the boundary of the neighborhood, with the full deformation performed over the $(l-1)$-handles themselves. The $(l-1)$-handles are contained in the interiors of $(l-1)$-simplices, so the E-spheres are varying only by isotopy and none are appearing or disappearing.

For the induction step of extending the shrinking over an i-handle corresponding to an i-simplex σ, we first deform F (as modified by the previous steps in the induction) in a neighborhood of the i-handle by decreasing the weights of all E-spheres except those that live over σ until these weights are zero over a smaller neighborhood of the i-handle, damping the deformation down to zero at the boundary of the larger neighborhood. The remaining E-spheres with nonzero weights still form a separating system since this was true over σ, and eliminating E-spheres in this way does not increase complexity or destroy the buffer property. After doing this deformation of weights, we perform the shrinking isotopy over a small neighborhood of the i-handle, damping it down to zero outside the handle. This finishes the induction step in the construction of the deformation F_t.

Step 3. By assumption we have $pF_0 = f_0$ and we need to arrange that $pF_t = f_t$ for all t. We will do this by inserting a collar on ∂D^l in D^l and placing a family of shrinkings in this collar connecting pF_t and f_t. Let us parametrize points in this collar by pairs $(x, s) \in \partial D^l \times I$ where for $s = 0$ we have the shrinkings pF_t and for $s = 1$ we have the shrinkings f_t. To begin, we use these same shrinkings for s near 0 and 1, respectively, then we damp these shrinkings down as we move away from $s = 0$ and $s = 1$. After performing these damped shrinkings we have a map $\partial D^l \times I \to \mathcal{C}$ and we can do the shrinking process for this family of configurations. Damping this down at $s = 0$ and $s = 1$, we then have the desired family of shrinkings on the collar interpolating between pF_t and f_t.

Step 4. For the new f_t produced in Step 3 we can apply the procedure in Step 1 to construct a lifting F_1 of f_1 with $pF_1 = f_1$. This yields shell spheres that serve as E-spheres for the sphere configurations $f_1(x)$. It remains to construct one more homotopy that changes these shell spheres to the E-spheres in F_1 without changing f_1. Recall that the deformation F_t in Step 2 involved shrinking the given E-spheres as well as the given spheres S_i. This means that if we apply the method in Step 2 to choose shell spheres for the spheres S_i as well as the E-spheres of the family F_1, these will be disjoint from the E-spheres in the family F_1. Forgetting the shell spheres for the E-spheres of F_1, we can then construct a homotopy of F_1 letting the weights of the given E-spheres go to zero while increasing the weights of the new shell spheres from zero to their specified values.

The shell E-spheres we have just constructed for the family f_1 over ∂D^l may not agree with the shell E-spheres constructed earlier over D^l, but we can just shift weights from one system of shell spheres to the other to get the final homotopy needed to finish the proof for the case of $\mathcal{C}$. The same proof works for $\mathcal{U}$ since the only time the spheres S_i were moved was when the complexity-reduction process was applied, and this works in $\mathcal{U}$ as well as $\mathcal{C}$. □

References

[BH] Tara Brendle and Allen Hatcher, Configuration spaces of rings and wickets. *Commentarii Math. Helv.* 88 (2013), 131–162.

[FS] Michael Freedman and Richard Skora, Strange actions of groups on spheres. *J. Diff. Geom.* 25 (1987), 75–98.