Brunnian Spheres

Hugh Nelson Howards

1. INTRODUCTION. Can the link depicted in Figure 1 be built out of round circles in 3-space? Surprisingly, although this link appears to be built out of three round circles, a theorem of Michael Freedman and Richard Skora (Theorem 2.1) proves that this must be an optical illusion! Although each component seems to be a circle lying in a plane, it is only the projection that is composed of circles and at least one of these components is bent in 3-space.

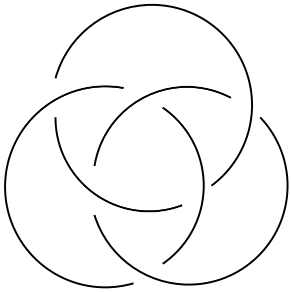

Figure 1. The Borromean Rings are a Brunnian link.

In this article we shed new light on the Freedman and Skora result that shows that no Brunnian link can be constructed of round components. We then extend it to two different traditional generalizations of Brunnian links.

Recall that a "knot" is a subset of $\mathbb{R}^3$ that is homeomorphic to a circle (also called a 1-sphere or S^1). Informally, a knot is said to be an "unknot" if it can be deformed through space to become a perfect (round) circle without ever passing through itself (see Figure 2); otherwise it is knotted (see Figure 3).

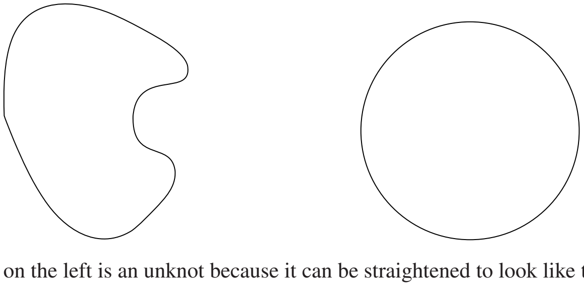

Figure 2. The figure on the left is an unknot because it can be straightened to look like the figure on the right without introducing any self-intersections.

A "link" L is just a collection of disjoint knots. A link L is an "unlink" of n components if it consists of n unknots and if the components can be separated without passing through each other (more rigorous definitions are given in section 3). Figure 2 could be thought of as an unlink of two components. Figure 4 shows *the Hopf link*, the simplest two-component link that is not an unlink.

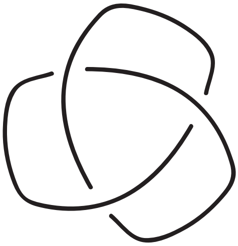

Figure 3. The trefoil is not an unknot because to deform it to a round circle it must pass through itself.

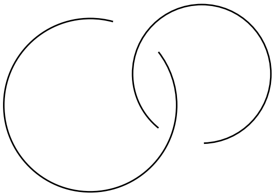

Figure 4. The Hopf link.

A link L is *Brunnian* if L is a link of $n(\geq 3)$ components such that L is not the unlink of n components but such that every proper sublink of L is an unlink. The most famous Brunnian link is called the *Borromean rings* (Figure 1). Note that eliminating any one of the components yields an unlink, as in Figure 5. In Figure 6, we see a nontrivial link that is not Brunnian.

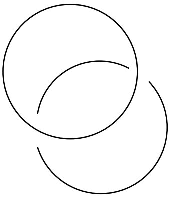

Figure 5. The Borromean Rings become an unlink if any component is deleted (here we have deleted the bottom left component, but any component would have yielded the same result).

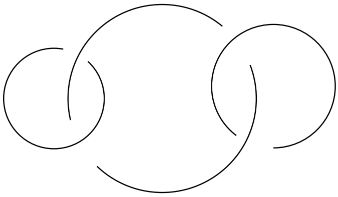

Figure 6. This is not an unlink, but it is not Brunnian either (if we remove the component in the center it becomes an unlink, but if we remove the component on the right, it does not).

2. A LITTLE HISTORICAL CONTEXT.

Brunnian links were introduced over a hundred years ago when knot theory was in its infancy. Hermann Brunn first mentioned them in his 1892 paper "Über Verkettung" (On Linking) [**1**].

In 1961 Debrunner furnished rigorous proofs for some of the claims in [**1**] by using the tools and language of modern knot theory [**3**]. Debrunner also generalized Brunnian links by defining (m, k)-Brunnian links. A link is (m, k)-*Brunnian* if it has m components, if no sublink of k or more components is an unlink, but if every sublink of at most $k - 1$ components is an unlink. Thus an (m, m)-Brunnian link is a traditional Brunnian link. Further work on these links was done by Penney in 1961 [**7**].

Takaaki Yanagawa also looked at a different generalization of Brunnian links in 1964. He considered higher dimensional analogues and gave examples of 2-spheres in $\mathbb{R}^4$ that formed Brunnian links [**8**].

In 1987 Freedman and Skora proved the following theorem in their paper "Strange Actions of Groups on Spheres" [**4**]:

Theorem 2.1 (Freedman, Skora). *No Brunnian link can be built out of round circles.*

The theorem went largely unnoticed because it is buried in a technical paper as a tool for investigating conformal and quasiconformal actions on S^3. In the paper Freedman and Skora give an example of a discrete, smooth, uniformly quasiconformal action on S^3 that is not conjugate under any homeomorphism to a conformal action and also an example of an action of the free group F_r of rank $r (\geq 2)$ on S^3 that is not conjugate to a uniformly quasiconformal action, even though every element of F_r individually is conjugate to a conformal transformation.

Because few people knew about Theorem 2.1, the special case of the Borromean rings was reproved in the early '90s by Bernt Lindstrom and Hans-Olov Zetterstrom [**6**], without any reference to the previous result of Freedman and Skora. Their proof was more complicated and less general than Freedman and Skora's. In 1993, Ian Agol independently (unpublished) gave the simplest proof known at this writing for the Borromean rings.

In this paper we generalize Freedman and Skora's result to higher dimensional links by obtaining the following result (Theorem 5.1): *No Brunnian link in* $\mathbb{R}^n$ *can be built out of round spheres.* It is also easy to generalize this further to Corollary 5.3: *No* (m, k)*-Brunnian link in* $\mathbb{R}^n$ *can be built out of round circles (spheres) if* $k > 2$.

Note that the definition of an (m, k)-Brunnian link does not make sense if $k < 2$, and it is not interesting if $k = 2$, since the second condition reduces simply to the requirement that the components of the link be unknots (and thus trivially one-component unlinks). Figure 7 show a (3, 2)-Brunnian link built out of round circles. This time it is not hard to confirm that it is not an optical illusion. The link is a generalization of a Brunnian link, but is not itself a Brunnian link, so we do not have a contradiction to Freedman and Skora's theorem.

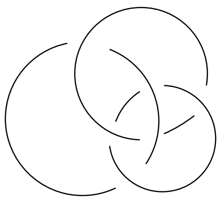

Figure 7. A (3, 2)-Brunnian link made out of three round circles.

A third, partial generalization appears in [**5**]. It replaces the requirement that the components of the link be round circles with the demand that they merely be convex curves (curves that bound convex, planar regions) and concludes:

Theorem 2.2. *The Borromean rings are the unique Brunnian links of three or four components that can be formed out of convex curves.*

Robert Davis has since extended this result and established the following one [**2**]:

Theorem 2.3. *No Brunnian link of five components can be formed out of convex curves.*

We make a few observations about Brunnian links. For each $n (\geq 3)$ there are an infinite number of nonequivalent Brunnian links with n components. One four-component Brunnian link is pictured in Figure 8. We see that if any single component is deleted, the others can be pulled apart one component at a time. Note also the four-fold symmetry. It is easy to use this symmetrical structure to build a Brunnian link of n components.

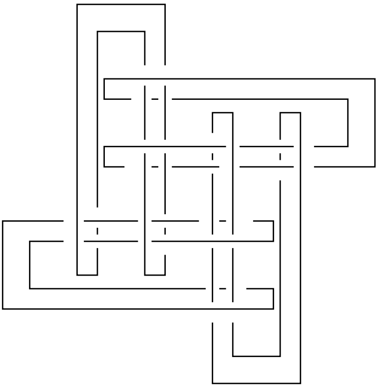

Figure 8. A Brunnian link with four components.

Another, more flexible method of construction is to take the two-component links that can be represented by two-strand braids and take the so-called Bing double of one of the components (see Figure 9). A *Bing double* is obtained by replacing a single knot with two knots. The two new knots must be contained in a small neighborhood of the knot being replaced. The two new components should be clasped together, but not linked as a pair. We see one of the knots in Figure 9 replaced with its Bing double. Notice that the two new components follow the path of the component they replace. Alone they form an unlink of two components, but together with the unaltered component they give a three component link that is not an unlink. This process yields an infinite family of three-component Brunnian links. If we double one of the components of each of the new links we get an infinite family of four-component Brunnian links. We can, of course, continue this process as long as we like.

Section 3 introduces the necessary definitions and notation. Section 4 provides some simple lemmas about intersections of flat balls in dimension three or higher that are then used in section 5 to prove generalized versions of Freedman and Skora's theorem.

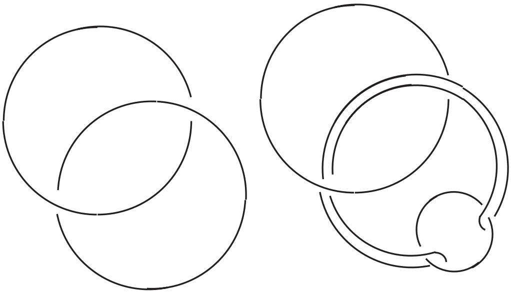

Figure 9. A Bing double of the Hopf link yields a Brunnian link. (The figure on the right is, in fact, a less conventional picture of the Borromean rings.)

Section 6 gives an example of a Brunnian link in $\mathbb{R}^4$ in order to confirm existence of such links. Finally, we close with some open questions.

3. DEFINITIONS. We turn to some standard definitions. Our setting will often be $\mathbb{R}^n$, Euclidean n-space with its traditional notion of distance, call it d. The *unit n-ball* $\hat{B}^n$ in $\mathbb{R}^n$ is all points of distance less than or equal to one from the origin (thus $\hat{B}^1$ is an interval, $\hat{B}^2$ is a disk, and so on). The *unit n-sphere* $\hat{S}^n$ in $\mathbb{R}^{n+1}$ is the boundary of $\hat{B}^{n+1}$ (thus, $\hat{S}^1$ is the unit circle in $\mathbb{R}^2$, $\hat{S}^2$ is the unit sphere in $\mathbb{R}^3$, etc.).

We will follow the topological convention of having very flexible definitions of spheres and balls. Any subset of $\mathbb{R}^n$ that is homeomorphic to the unit k-sphere is called a *k-sphere* (or S^k). Similarly, any subset of $\mathbb{R}^n$ that is homeomorphic to the unit k-ball in $\mathbb{R}^k$ is called a *k-ball* (or B^k). Thus a cube is actually considered a sphere even though it is not round. Note that we differentiate the special case of the unit n-ball (unit n-sphere) from the more general case of a topological n-ball (topological n-sphere) by using the notation $\hat{B}^n$ instead of B^n ($\hat{S}^n$ versus S^n).

Once we accept the topological notion that not all spheres are round, we would like to be able to specify the special case of a sphere that actually is round and have the definition match our intuitive notion of what a round sphere must be. Obviously all points at distance $d(> 0)$ from the origin in $\mathbb{R}^{k+1}$ is a round k-sphere in $\mathbb{R}^{k+1}$. In general, a k-sphere S^k in $\mathbb{R}^n$ is a round k-sphere if there is a bijection f between S^k and the unit k-sphere $\hat{S}^k$ in $\mathbb{R}^{k+1}$ that preserves distance up to multiplication by a constant t (there exists a constant $t > 0$ such that for every pair of points x and y of S^k it is true that $d(x, y) = t\, d(f(x), f(y))$). This indeed matches one's intuitive notion of what a round sphere should be. For instance, the ellipse in $\mathbb{R}^2$ with equation

$$\frac{x_1^2}{4} + \frac{x_2^2}{9} = 1$$

is a 1-sphere (or circle), but it is not a round 1-sphere (round circle).

A subset K of $\mathbb{R}^n$ is a *knot in* $\mathbb{R}^n$ if K is homeomorphic to S^k for some $k(\geq 1)$. By a *link in* $\mathbb{R}^n$ is meant a subset L of $\mathbb{R}^n$ that is homeomorphic to a disjoint union of finitely many knots (possibly of different dimensions). Knot theory is usually restricted to the case where $n = 3$ and each knot is homeomorphic to $\hat{S}^1$.

A link $L = F_1 \cup F_2 \cup \cdots \cup F_m$ in $\mathbb{R}^n$ is *an unlink in* $\mathbb{R}^n$ if for each i the knot F_i bounds a ball B_i (of appropriate dimension) such that $B_i \cap F_j = \emptyset$ $(j \neq i)$. If $n = 3$ and we restrict to circles, we obtain the traditional definition of an unlink, where each component bounds a disk disjoint from the other components. If a link L in $\mathbb{R}^n$ of $m(\geq 3)$ components is not an unlink (in $\mathbb{R}^n$), yet every proper sublink is an unlink (in $\mathbb{R}^n$), we call L a *Brunnian link*. If a link L of m components has no sublink of k or more components that is an unlink in $\mathbb{R}^n$, but every sublink of fewer than k components is an unlink in $\mathbb{R}^n$, we christen L an *(m,k)-Brunnian link*.

Two knots are said to be equivalent if one can be stretched through space to look like the other in a nice continuous manner. The technical term for this stretching is an *ambient isotopy*. People looking into knot theory for the first time often expect homotopy to capture the notion of equivalence, but it turns out this is not a strong enough equivalence. All knotted circles in $\mathbb{R}^3$, for example, are homotopic to each other. The next logical attempt is to define isotopy. Let $h_t : X \to Y$ $(0 \leq t \leq 1)$ be a homotopy with $h_0(X) = X_0$ and $h_1(X) = X_1$. Then $\{h_t\}$ is an *isotopy* from X_0 to X_1 if each h_t is a homeomorphism from X onto its image, in which case X_0 and X_1 are *isotopic*. Surprisingly this is still not quite a strong enough definition. Every knot (including the trefoil pictured in Figure 3) is actually isotopic to the unknot! One can imagine h_1 mapping S^1 to the trefoil, with the knotted portion of the trefoil sitting inside a ball of radius one about the origin and only an unknotted arc escaping the ball. As t drops from 1 to 0 we shrink the knotted portion of the trefoil by tightening the knotted arc so that at time t this portion is contained in a ball of radius t around the origin. At time 0 the result must be an unknot. (We can envision such an isotopy by picturing a small knot tied in a necklace, where we keep pulling the knot tighter and tighter shrinking the knotted portion of the necklace down as t drops from 1 to 0.)

We address this final issue by requiring the homotopy to be a isotopy of both the knot and the space containing it. More formally, let $h_t : Y \to Y$ be a homotopy, let X be a subset of Y, and let $h_0(X) = X_0$ and $h_1(X) = X_1$. Then $\{h_t\}$ is an *ambient isotopy* from X_1 to X_2 if each h_t is a homeomorphism from Y onto itself. In this situation X_0 and X_1 are *ambient isotopic*. When topologists talk about an isotopy between two knots, they are almost always referring to an ambient isotopy. The same will hold true in this paper. If there is an ambient isotopy between knots k_0 and k_1 they are considered to be two different embeddings of some S^p, but the same knot.

4. INTERSECTION OF FLAT BALLS. For the proofs that follow, it will often be easiest to think of the link in $\hat{S}^n$ instead of $\mathbb{R}^n$. It is easy to get from one to the other using a conformal mapping between $\hat{S}^n$ minus a point and $\mathbb{R}^n$ (such as stereographic projection) that preserves all of the crucial properties of the link in which we are interested.

We now establish two lemmas that are necessary for the proof of Theorem 5.1. Assume that $L = F_1 \cup F_2 \cup \cdots \cup F_k$ is a Brunnian link in $\hat{S}^{n-1}$ made of round spheres. We examine what would happen if L contained two round 2-spheres, say F_1 and F_2. The other cases are analogous.

Let X_1 and X_2 be three-dimensional Euclidean planes in $\mathbb{R}^n$ (which in turn contains $\hat{S}^{n-1}$) such that $X_1 \cap \hat{S}^{n-1} = F_1$ and $X_2 \cap \hat{S}^{n-1} = F_2$. Let $D_1 = X_1 \cap \hat{B}^n$ and $D_2 = X_2 \cap \hat{B}^n$. Then D_1 is a 3-ball in $\mathbb{R}^n$ bounded by F_1, and D_2 is a 3-ball bounded by F_2.

Now $X_3 = X_1 \cap X_2$ is one of the following: (1) the null set, (2) a point, (3) a line (a flat copy of $\mathbb{R}^1$), or (4) a two-dimensional plane (a copy of $\mathbb{R}^2$). This implies that $D_3 = D_1 \cap D_2$ is either (1) the null set, (2) a point, (3) an interval, or (4) a disk. The final two options are, of course, impossible, since $D_3 \subset D_1 \cap D_2$ and $\partial D_3 \subset \partial D_1 \cap \partial D_2 = F_1 \cap F_2 = \emptyset$. Thus the intersection of D_1 and D_2 can have no boundary and must be either the null set or a point.

Since F_1 and F_2 are 2-spheres bounding 3-balls D_1 and D_2, we may assume that if $D_3 = D_1 \cap D_2$ is a point, then $n = 6$. To understand why this is true, we observe that this is analogous to the way two straight lines (which are one-dimensional sets) can intersect in $\mathbb{R}^n$ (an n-dimensional space). Obviously in $\mathbb{R}^1$ two ostensibly different lines actually coincide. In $\mathbb{R}^2$ if two nominally different lines overlap in every point, then it is easy to create an isotopy of the second line that pushes it off of the first line, leaving them disjoint, but moving each point on the second line an arbitrarily small distance. A specific example of this is that, given two copies of the line $y = 0$, we can push one copy up to the line $y = \epsilon/2$, leaving the two lines disjoint but not moving any point on the new line more than distance ϵ. On the other hand, if the two lines in $\mathbb{R}^2$ intersect in only one point, as the x- and y-axes do, any isotopy that makes the lines disjoint will have to move points arbitrarily large distances. In $\mathbb{R}^n$ with $n \geq 3$, however, every pair of lines can be made disjoint by means of an isotopy that displaces no point more than ϵ, for any prescribed ϵ (> 0). For example, if our first line were the x-axis and our second line the y-axis, we could take an isotopy of the y-axis that moves each point $(0, y, 0)$ to $(\epsilon/2, y, \epsilon/2)$. The lines end up disjoint, yet no point of either line is moved more than ϵ. We see that for two lines to intersect in a point in $\mathbb{R}^n$ in such a way that the intersection cannot be eliminated via a small isotopy n must equal 2. We are looking at $\mathbb{R}^1$ and $\mathbb{R}^1$ in $\mathbb{R}^2$. The intersection is interesting (neither trivial nor the entire line) in this case because the dimensions of the spaces involved add up perfectly: $1 + 1 = 2$. A plane will intersect a line in one point that cannot be removed by small isotopies in $\mathbb{R}^3$, but not $\mathbb{R}^2, \mathbb{R}^4, \ldots$ since $2 + 1 = 3$.

We will be intersecting two 3-balls and examining when they might intersect in a point that cannot be removed via a small isotopy. Since 3-balls are three-dimensional, the principle just discussed implies that this can happen only in $\mathbb{R}^6$ and not in $\mathbb{R}^n$ for $n \neq 6$. If two surfaces intersect minimally with respect to isotopy, we say that they are in *general position*. Thus the x-axis and y-axis are in general position in $\mathbb{R}^2$, but not in $\mathbb{R}^3$.

Lemma 4.1. *If F_1 and F_2 bound 3-balls D_1 and D_2 that are in general position in $\mathbb{R}^6$ and if $X_3 = D_1 \cap D_2$ is a point, then $F_1 \cup F_2$ is not an unlink.*

The easiest proof we could find was suggested by Genevieve Walsh. It shows that in this context there is an isotopy of $\hat{S}^5$ that does not introduce any intersections and takes F_1 and F_2 to linked great spheres.

Proof. We want to reduce to the case where X_1 and X_2 intersect at the origin. Assume that they do not. Let S^5_t be the round 5-sphere of radius t centered at the origin in $\mathbb{R}^6$, let B^6_t be the ball bounded by S^5_t in $\mathbb{R}^6$, let $F_i(t) = X_i \cap S^5_t$, and let $D_i(t) = X_i \cap B^6_t$. We observe the changes in the configuration as t grows from 1 to ∞.

If we rescale the B^6_t to unit balls (and thus the S^5_t to unit spheres and $F_1(t) \cup F_2(t)$ to a new link in the unit sphere), this yields an isotopy of $F_1 \cup F_2$ through S^5. As the isotopy progresses the intersection of the planes containing the spheres moves toward the origin. In the limit it reaches the origin (this is a direct result of the rescaling: since B^6_t is shrunk by a factor of t, all distances, including the distance from $D_1(t) \cap D_2(t)$ to the origin, are also shrunk by a factor of t). This means that, as t approaches ∞, the link components converge to great spheres bounding flat balls in B^6 that intersect at the origin. In this case the two spheres constitute a simple generalization of the Hopf link, so are indeed linked. ∎

The foregoing proof works for links composed of 2-spheres, but the argument for other dimensions is completely analogous. Thus, we arrive at the following result:

Lemma 4.2. *If F_i and F_j are not pairwise linked, then $D_i \cap D_j = \emptyset$ (the flat balls that F_i and F_j bound are disjoint).*

5. BRUNNIAN SPHERES. With the lemmas from section 4 in hand and with the aid of the Freedman and Skora techniques, Theorem 2.1 becomes fairly easy to generalize to higher dimensions (the proof of Theorem 2.2, by contrast, requires completely different methods).

Theorem 5.1. *If L is a Brunnian link in $\mathbb{R}^n$, then L cannot be constructed exclusively of round components.*

Theorem 5.1 is a direct consequence of the following result:

Theorem 5.2. *Any link in $\hat{S^n}$ (or $\mathbb{R}^n$) of round spheres with at least three components and all components pairwise unlinked is the unlink.*

Proof. Let $L = F_1 \cup F_2 \cup \cdots \cup F_m$ be any link in $\hat{S^n}$ with all components pairwise unlinked and with F_i represented by a round v_i-sphere in $\hat{S^n}$. Consider S^n as the boundary of $\hat{B^{n+1}}$, the unit $(n+1)$-ball in $\mathbb{R}^{n+1}$, as in Lemma 4.2. As in that case, F_i bounds a ball $D_i = X_i \cap B^{n+1}$ (the ball has dimension one higher than F_i). By Lemma 4.2 we may assume that $D_i \cap D_j = \emptyset$ when $i \neq j$.

As earlier, let S^n_t signify the round sphere of radius t centered at the origin in $\mathbb{R}^{n+1}$, let B^{n+1}_t be the ball bounded by S^n_t in $\mathbb{R}^{n+1}$, and let $F_i(t) = X_i \cap S^n_t$. We can assume that none of the D_i contains the origin (if not, we perturb the corresponding knot slightly so that this is true).

This time we take the limit as t goes from 1 to 0. As t goes to 0, each $F_i(t)$ shrinks to a point and then disappears when D_i becomes tangent to S^n_t. Rescaling the S^n_t to unit spheres yields an isotopy of the generalized link in S^n that keeps the components disjoint but shrinks each component to a point, showing we have a generalized unlink. (We focus on S^n_t, but rescale at each t so that S^n_t is a unit sphere, so it is as if S^n is fixed and the knots in $\mathbb{R}^{n+1}$ move instead of the other way around.) ■

Corollary 5.3. *No (m, k)-Brunnian link in $\mathbb{R}^n$ can be constructed from round spheres if $k > 2$.*

Proof. Let L be an (m, k)-Brunnian link in $\mathbb{R}^n$, and consider a k-component sublink L' of L. We know that L' is not an unlink by the definition of an (m, k)-Brunnian link in $\mathbb{R}^n$. Also by the definition of (m, k)-Brunnian links, every sublink of L' is an unlink in $\mathbb{R}^n$. Thus L' is a Brunnian link in $\mathbb{R}^n$ and cannot be built out of round spheres. Since L' is contained in L, L clearly cannot be constructed from round spheres either. Note that, as with all results in this paper, if we set $n = 3$ Corollary 5.3 holds true in the traditional setting of Debrunner's original generalization of Brunnian links. ■

6. A BRUNNIAN LINK IN $\mathbb{R}^4$. In this section we give an example to show that Brunnian links in $\mathbb{R}^n$ $(n > 3)$ do exist. Let Π_ω project $\mathbb{R}^4$ onto $\mathbb{R}^3$ via the map $\Pi_\omega((x, y, z, w)) = (x, y, z)$. Take $L = F_1 \cup F_2 \cup k$, where F_1 and F_2 are two-spheres in $\mathbb{R}^4$ lying in the three-plane X defined by the equation $w = 0$ and k is a knot containing the points $p_1, p_2, \ldots p_8$ with the property that p_i lies in the three-plane given by $w = (-1)^i$ (Figure 10 shows $\Pi_\omega(L)$). Note that

$$\Pi_\omega(k) \cap \Pi_\omega(F_1) = \Pi_\omega\{p_1, p_2, p_3, p_4\}$$

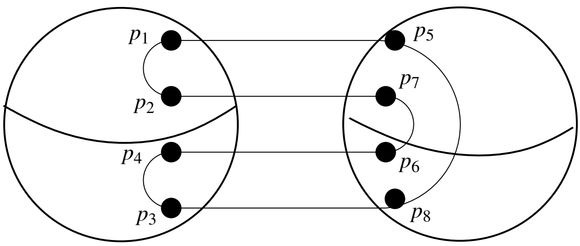

Figure 10. The projection of a Brunnian link in $\mathbb{R}^4$ into $\mathbb{R}^3$ by Π_ω

and

$$\Pi_\omega(k) \cap \Pi_\omega(F_2) = \Pi_\omega\{p_5, p_6, p_7, p_8\}.$$

Let s_{ij} be the component of $k - \{p_i, p_j\}$ that intersects $F_1 \cup F_2$ the fewest times when both are projected into $\mathbb{R}^3$ (for example, $\Pi_\omega(s_{12})$ is contained in the ball bounded by $\Pi_\omega(F_1)$).

We assert the following:

Theorem 6.1. *L is a Brunnian link.*

Proof. Note that if k is removed, then F_1 and F_2 are not linked. To see that if F_1 is deleted, F_2 and k are not linked we want to show that they can be deformed through space and pulled apart. Because it is hard to picture $\mathbb{R}^4$ it may be easier to perform the deformation in R^4 but then watch what happens after projecting to $\mathbb{R}^3$ with Π_ω. If we can figure out how to deform them through $\mathbb{R}^4$ so that the projection becomes separated, but throughout the deformation any time $\Pi_\omega(k)$ intersects $\Pi_\omega(F_2)$ the preimage on k has a different w value from the preimage on F_2, then we know that the deformation has kept the components disjoint. This is true because, if the two components intersect, each has a point with identical x-, y-, z-, and w-coordinates. The intersections in the projection are the only points with the same x-, y-, and z-coordinates and thus are the only potential intersections back in $\mathbb{R}^4$. The preimages do intersect if and only if the w-coordinates are also the same. If we tried to pull s_{67}, for example, directly out of F_2 we would run into trouble when the point on the arc with $w = 0$ hit F_2. If there were no point on s_{67} with $w = 0$, then it would be easy to pull it out without introducing an intersection in $\mathbb{R}^4$.

On the other hand, since F_1 was deleted, k can be deformed through $\mathbb{R}^4$ in such a way that every point of s_{57} has $w = -1$, $\Pi_\omega(k)$ never changes, and the rest of k is kept fixed. Thus the trajectory of k is kept disjoint from F_2 throughout the process. This is easy to do because both endpoints of s_{57} have $w = -1$ and $\Pi_\omega(s_{57})$ intersects F_2 only in its endpoints (the w-component is homeomorphic to a subset of $\mathbb{R}$, and since $\mathbb{R}$ is simply connected, it is easy to alter just the w-component until it is constant).

Similarly, s_{68} can be deformed through space to have $w = 1$. Now these two arcs no longer have points with $w = 0$ on them and thus can be pulled to the inside of F_2, so that $\Pi_\omega(k)$ is entirely contained inside of $\Pi_\omega(F_2)$. Now leaving the x-, y-, and z-components fixed (and thus $\Pi_\omega(k)$ disjoint from $\Pi_\omega(F_2)$), we can deform k through space so that all the points on k have $w = 1$, again while leaving the x-, y-, and z-components of each point fixed. By the previous observations, once every point on k has a nonzero w-coordinate, k can clearly be deformed through space so that $\Pi_\omega(k)$ is entirely outside of $\Pi_\omega(F_2)$, showing that this is indeed an unlink. The same type of argument can be made for F_1 and k.

Now that we have shown that every sublink of L is an unlink we need demonstrate only that L is not the unlink and we will know that it is indeed Brunnian. To prove this we look at the simpler link $M = k_1 \cup k_2 \cup k'$ in $\mathbb{R}^3$ pictured in Figure 11, for which our intuition is better but the argument is analogous. Note that M is just another presentation of the Borromean rings, so we are actually proving that the Borromean Rings are not an unlink and thus are indeed a Brunnian link, a fact that we have previously asserted but not proved.

Note that, as with L, when we delete k, k_1, or k_2 from M we get an unlink. In general, a circle that is a component of an unlink must bound a disk in the link complement. It therefore must represent a trivial element in the fundamental group of the

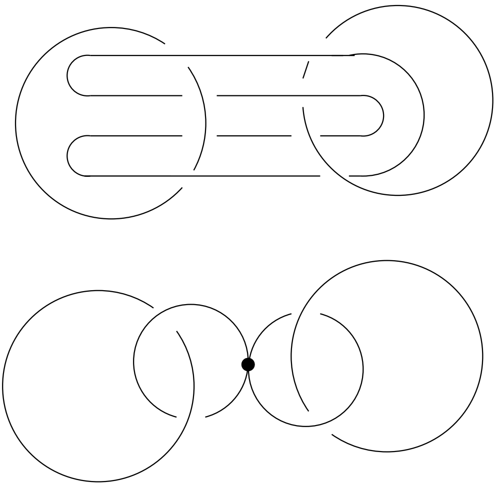

Figure 11. The Borromean Rings, represented as $M = k_1 \cup k_2 \cup k'$ are pictured on top. The second picture shows k_1, k_2, and the generators of $\pi_1(\mathbb{R}^3 - (k_1 \cup k_2))$.

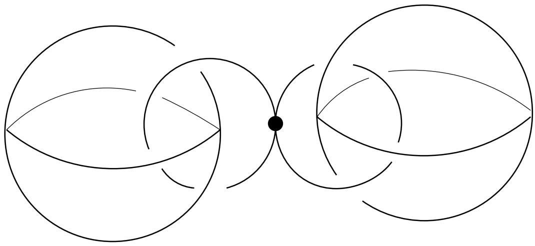

Figure 12. The images under Π_ω of F_1, F_2, and the generators of $\pi_1(\mathbb{R}^4 - (F_1 \cup F_2))$.

complement of the other link components. For M we will show that k' is not trivial in the fundamental group of $\mathbb{R}^3 - (k_1 \cup k_2)$, $\pi_1(\mathbb{R}^3 - (k_1 \cup k_2))$. This is a free group on two generators. The generators, which we may call a and b, are represented by the two circles that make up the figure eight in Figure 11. Note that the fundamental group of $\mathbb{R}^3 - k_i$ (or, for that matter, of any simply connected subset of $\mathbb{R}^3$ minus any unknotted circle in its interior) is a free group on one generator. The Seifert/Van Kampen theorem ensures that if a space can be decomposed into two subspaces that intersect in a connected, simply connected subset, then the fundamental group of the original space is the free product of the fundamental groups of the smaller spaces. Now k' is homotopic in $\mathbb{R}^3 - (k_1 \cup k_2)$ to $aba^{-1}b^{-1}$, the commutator of two generators of the fundamental group. Thus, k' is not homotopic to the trivial element in the fundamental group of $\mathbb{R}^3 - (k_1 \cup k_2)$, and M is not an unlink. We may make a totally analagous argument for L. The fundamental group of $\mathbb{R}^4 - (F_1 \cup F_2)$ is again the free group on two generators, generated by two circles forming a figure eight. We can see the images under Π_ω of the generators of the fundamental group of $\mathbb{R}^4 - (F_1 \cup F_2)$ in Figure 12. Each of the two loops intersects a respective sphere in two points, but in the preimage of the figure eight, for each loop one of these two points has positive w-coordinate and the other has negative w-coordinate. Again k represents the commutator element of the fundamental group, a free group on two generators. As such, it is not trivial, proving that L is not the unlink. ■

7. OPEN QUESTIONS.

We end with a few open questions.

Question 7.1. *Is there a Brunnian link other than the Borromean rings that can be formed out of convex curves?*

The combinatorics in such problems become much more complicated as the number of components goes up. Bob Davis has proved that the answer is no for $n = 5$, but a new strategy seems necessary for large n.

Question 7.2. *If $n \geq 6$, how many n-component Brunnian links can be formed out of planar curves?*

As we observed in the introduction, for each n the answer is at least one, but perhaps all such examples arise from the relatively small family of Brunnian links that are formed by iterated doubles of the Hopf link.

Question 7.3. *Is there a generalized Brunnian link other than the Borromean rings that can be formed out of convex embeddings of spheres in $\mathbb{R}^n$?*

This question appears to be wide open: perhaps there is a good example of such a link, or perhaps no such link exists.

ACKNOWLEDGMENT. The author would like to thank Ian Agol and Seungsang Oh for helpful comments.

REFERENCES

1. H. Brunn, Über Verkettung, Sitzungsber. Bayerische Akad. Wiss., *Math. Phys. Klasse* **22** (1892) 77–99.
2. R. M. Davis, *Brunnian Links of Five Components*, Master's thesis, Wake Forest University, Winston-Salem, NC. 2005.
3. H. E. Debrunner, Links of Brunnian type, *Duke Math. J.* **28** (1961) 17–23.
4. M. H. Freedman and R. K. Skora, Strange actions of groups on spheres, *J. Differential Geom.* **25** (1987) 75–98.
5. H. N. Howards, Convex Brunnian links, *J. Knot Theory Ramifications* **15** (2006) 1131–1140.
6. B. Lindstrom and H. O. Zetterstrom, Borromean circles are impossible, this MONTHLY **98** (1991) 340–341
7. D. E. Penney, Generalized Brunnian links, *Duke Math J.* **36** (1969) 31–32.
8. T. Yanagawa, Brunnian systems of 2-spheres in 4-space, *Osaka J. Math.* **1** (1964) 127–132.

HUGH HOWARDS received his B.A. at Williams College and received his Ph.D. (1997) at the University of California, San Diego, where his advisor was Mike Freedman. He is a Sterge Faculty Fellow and associate professor at Wake Forest University. Since coming to Wake Forest in 1997, he has won both the Reid-Doyle prize for the top young teacher on campus and the Student Government award for the top teacher on campus.