To appear, 2014, in The Mathematical Intelligencer, a Springer-Verlag Journal.

Visualizing Paradoxical Sets

Grzegorz Tomkowicz, Bytom, Poland; gtomko@vp.pl
Stan Wagon, Macalester College, St. Paul, Minnesota, USA; wagon@macalester.edu

Two travelers arrive at Hilbert's Grand Hotel. The first asks for a room. Hilbert asks, "Smoking or nonsmoking?" "I need a smoking room," he replies. "You're in luck, sir; we are full, but we can move everyone to an adjacent room, which will free up a smoking room." The second traveler then asks for a nonsmoking room. "Sorry, sir, that's not possible." "But can't you just move people so as to free up a nonsmoking room?" "No, sir, we cannot. Some guests have rooms adjacent to those of family or friends. We can only move the entire hotel so as to preserve all those relationships. That is why all rooms here are the same: we have only smoking rooms. I suggest you try Mycielski's Hotel."

The traveler crosses the street hoping that Mycielski's Hotel, despite its "Always No Vacancy" sign, would have only nonsmoking rooms. At the front desk he overhears a conversation identical to that at Hilbert's place and sees the customer placed in a just-vacated smoking room. So he is in despair, but asks anyway: "Can you find a nonsmoking room for me?" Mycielski answers, "We are full, but let me try … No problem: We will move everyone so as to preserve room type and all room adjacencies, thus freeing up a nonsmoking room for you." "Wonderful. But didn't you just do the same thing for a smoker?" "Yes, indeed. We are a non-Euclidean hotel and can handle any type of request: pet-free rooms for smokers, pet-friendly rooms for nonsmokers. Although we are full, we are able to handle any combination of requests."

1. Introduction

The hyperbolic plane $\mathbb{H}^2$ is a wonderful place. The richness of its isometry group allows one to study in a completely constructive way several counterintuitive ideas of set theory, such as the Banach–Tarski Paradox and the existence of a certain type of Sierpiński set, referred to in our variation of the Hilbert Hotel parable. In 1983 Jan Mycielski wondered whether the Banach–Tarski Paradox could be illustrated in $\mathbb{H}^2$ and in 2013 he asked the same question about Sierpiński sets. We will show here how both phenomena can be given very concrete interpretations in the hyperbolic plane.

We will use the upper-half plane model of hyperbolic geometry, where lines are semicircles perpendicular to the real axis and isometries are given by linear fractional transformations corresponding to matrices in $PSL_2(\mathbb{R})$, the group of 2×2 matrices with determinant 1, modulo the negative identity matrix. We will also use the Poincaré disk model, which is easily obtained from the half-plane by a single linear fractional transformation. We say that two subsets of a metric space E are *congruent* if there is an isometry of E (a distance-preserving bijection) that takes one to the other; we also use the term for subsets of a group where one arises from the other by left multiplication by a group element.

2. The Disappearing Hyperbolic Squares

It is no trick to make a single object vanish while preserving all relations among the others. As with Hilbert's Hotel, start with $\{0, 1, 2, 3, \ldots\}$; translation gives $\{1, 2, 3, \ldots\}$. But what if we want to free up a room chosen from two given rooms? This two-point problem asks for a set E (in some metric space) containing two points p and q so that E is congruent to $E \setminus \{p\}$, and also to $E \setminus \{q\}$? It is not hard to see that no such set exists in $\mathbb{R}^1$ [7, Thm. 6.14]; E. G. Straus proved [4] that even in $\mathbb{R}^2$, one cannot find such a set.

More generally, a *Sierpiński set* in a metric space or group is a set E such that, for any $p \in E$, E is congruent to $E \setminus \{p\}$. In short, the geometry of E is unchanged after deletion of any point. Straus gave a construction of such a set in the free group F_2, using the fact that F_2 has a subgroup that is free on infinitely many generators. By repeated deletion, the set is congruent to the set obtained by removing any finite subset. Let us call a set as in the preceding paragraph (either p or q can be deleted without changing the geometrical essence of the set) a *weak Sierpiński set*.

In 1958, Jan Mycielski [5] showed how one can construct a subset of $\mathbb{R}^3$ that is geometrically unchanged under any set of countably many changes (points can be added or deleted). That gives a blueprint for a remarkably flexible hotel, but only for very small guests: the Axiom of Choice (AC) was used and the set is not discrete, so the guests have to be one-dimensional. Work of K. Satô (§5) can be used to constructively define a Sierpiński set in $\mathbb{R}^3$, without using AC, but it too fails to yield a usable hotel since it is dense in the unit sphere. It seems unlikely that a discrete Sierpiński set exists in 3-space, but that is an open question.

Before turning to the hyperbolic plane, we focus on the abstract free group F_2, with generators σ and τ. Figure 1 shows F_2 as a tree, where words are formed by appending characters on the left; this is not the standard Cayley graph, which works instead on the right. We can easily form a weak Sierpiński set in F_2. Simply let E be the union of S, the words having σ as the rightmost term (red in Fig. 1), and T, the words ending on the right with τ (blue in Fig. 1). Then $\sigma S = S \setminus \{\sigma\}$ while $\sigma T = T$, yielding $\sigma E = E \setminus \{\sigma\}$; similarly $\tau E = E \setminus \{\tau\}$.

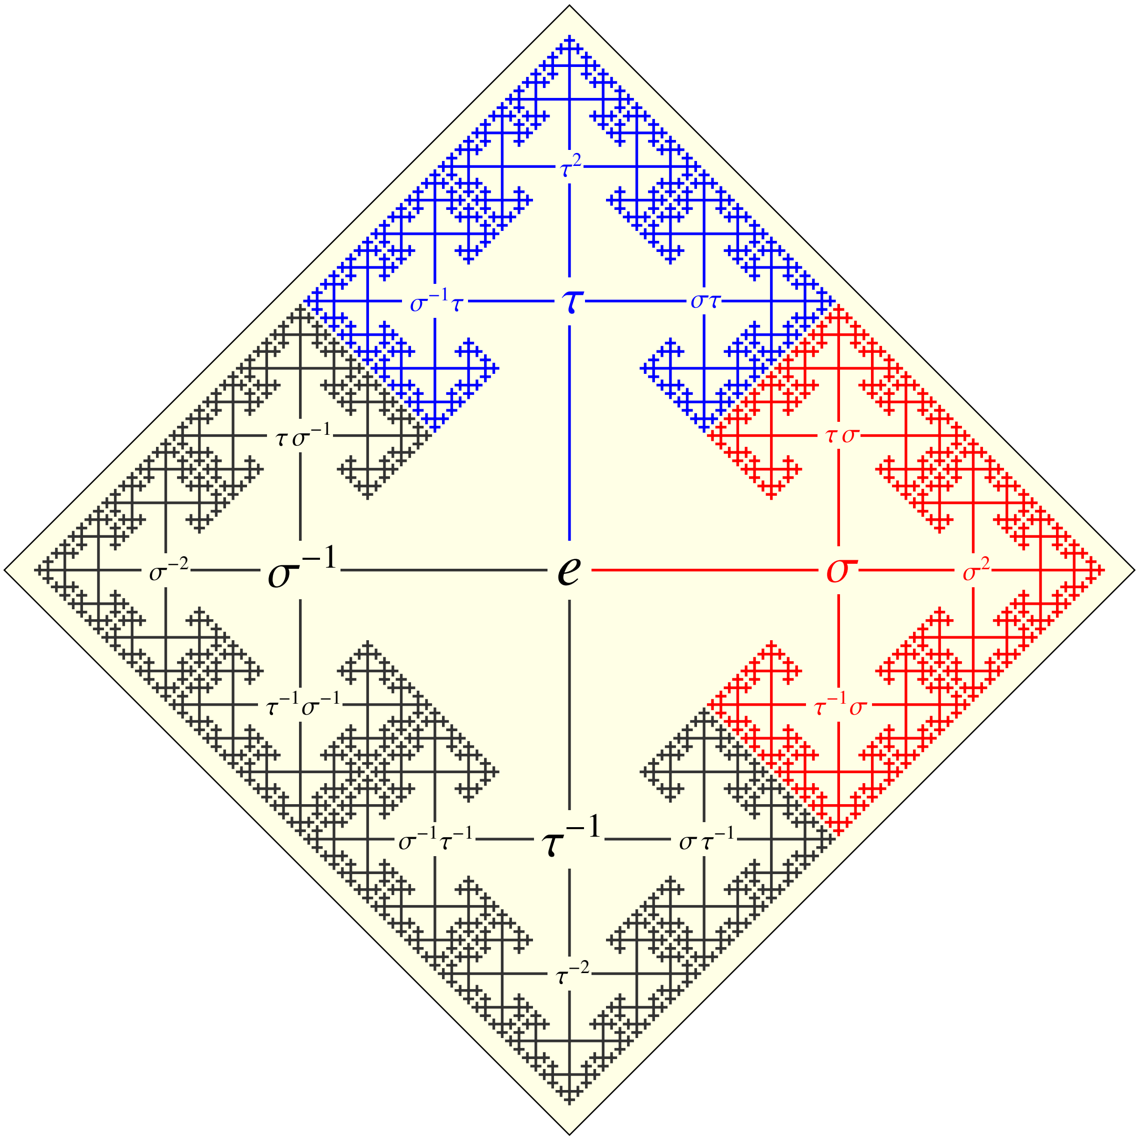

Figure 1. The free group $\langle\sigma, \tau\rangle$ has a subset E (shown in red and blue) that satisfies: $\sigma E = E\setminus\{\sigma\}$ and $\tau E = E\setminus\{\tau\}$.

Now let σ and τ be the hyperbolic isometries defined in the upper half-plane by $\sigma(z) = z/(2z+1)$ and $\tau(z) = z+2$. These generate a group isomorphic to F_2. Recall that composition of linear fractional transformations corresponds to multiplication of the corresponding matrices.

Proposition 1. The matrices $A = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}$ and $B = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$ generate a free group under matrix multiplication.

While a direct algebraic (or geometric) proof of Proposition 1 is not difficult, we will instead derive it from a result we will need in §3. Proposition 2 deals with the independence (except for the finite order of the generators) of the two isometries $z \mapsto -1/z$ and $z \mapsto -1/(z+1)$, where again these formulas are for the upper half-plane.

Proposition 2. The matrices $S = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$ and $T = \begin{pmatrix} 0 & -1 \\ 1 & 1 \end{pmatrix}$ generate the free product $\mathbb{Z}_2 * \mathbb{Z}_3$.

Proof. First note that S and T have order 2 and 3, respectively. Let $R = T^2$. Suppose w is a nonempty string in S, T, R, with no adjacencies of the form SS, TT, TR, or RT, that equals the identity. Conjugating by S, TS, or RS if necessary, we may assume that $w = Sy\ldots ySyS$, where each $y \in \{T, R\}$. We have $ST = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and $SR = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$. Now start with $(-1, 0)$; the rightmost S in w gives $(0, 1)$ and further applications of ST and SR cannot decrease the entries and so never produce $(-1, 0)$, a contradiction. □

And now Proposition 1 follows, since A and B can be expressed in terms of S and T.

Proof of Proposition 1. With S, T, R as in Proposition 2, we have $STST = A$ and $SRSR = B$. Any nontrivial reduced word in $A = SRSR$, $A^{-1} = RSRS$, $B = STST$, and $B^{-1} = TSTS$, viewed as a word in S, T, R, is a nontrivial reduced word in $\mathbb{Z}_2 * \mathbb{Z}_3$. This is because of the eight possible adjacencies AA, BB, AB, BA, AB^{-1}, BA^{-1}, $A^{-1}B$, and $B^{-1}A$, only the last two have any reduction. The first of these is $RSRS\,SRSR = RSTSR$, which still ends in R and so leads to no other cancellation. The other case is the same, with T and R switched. Proposition 2 showed that any such word in S, T, R is not the identity. □

For any discrete group of isometries of $\mathbb{H}^2$, one can select a point in the plane, look at the orbit of the point under the group, and then, for each point X of the orbit, consider the Voronoi region R_X, the set of points closer to X than to any other orbit point. These regions $\{R_X\}$ form a tiling of the $\mathbb{H}^2$; the region containing the initial point is a fundamental domain, and all other regions arise by transforming the fundamental domain by a group element.

Figure 2 shows the tiling for $\langle\sigma, \tau\rangle$; the square labeled "$e$" is the fundamental domain and each tile is obtained from it using the indicated group elements. The labeling here has a tree structure, with e as the root (having four children) and each other node branching outward in three ways: for example, σ branches to σ^2, $\sigma\tau$, and $\sigma\tau^{-1}$. This tree differs in a very important way from the tree in Figure 1: there we wanted words ending in σ (on the right) to be together, so the branching added letters on the left (e.g., σ has the children $\tau\sigma$, $\tau^{-1}\sigma$, and σ^2). But in the hyperbolic tiling, the tree structure moves through words built up by appending letters on the right.

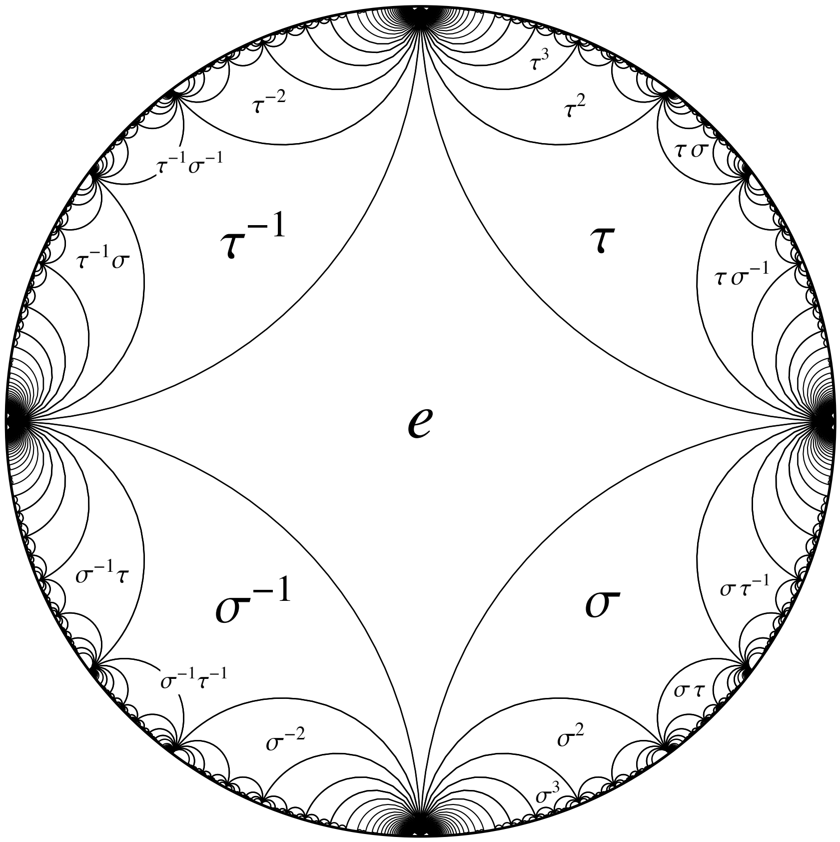

Figure 2. The hyperbolic tiling corresponding to the free group generated by σ and τ.

It is now a simple matter to identify the tiles corresponding to words in the weak Sierpiński set in F_2 shown in Figure 1. One could use points (from the orbit of the center of the disk under the group) rather than regions, but visualization is improved — and the hotel analogy respected — if we identify each such point with the tile it lies in. Working this way, the collection of all colored tiles in Figure 3 represents a weak Sierpiński set E. If E^{-} is the result of removing the large light red square in the fourth quadrant, then $\sigma(E) = E^{-}$; the same holds for the light blue square and τ. For example, in the red-deletion case, the large blue square in the third quadrant gets taken by σ to the large blue square in the first quadrant, since this corresponds to $\sigma\left(\sigma^{-1}\tau\right) = \tau$. Note that the left-right distinction of the previous paragraph is the reason the hotel consists of partially disconnected wings, rather than the contiguous shapes of Figure 1.

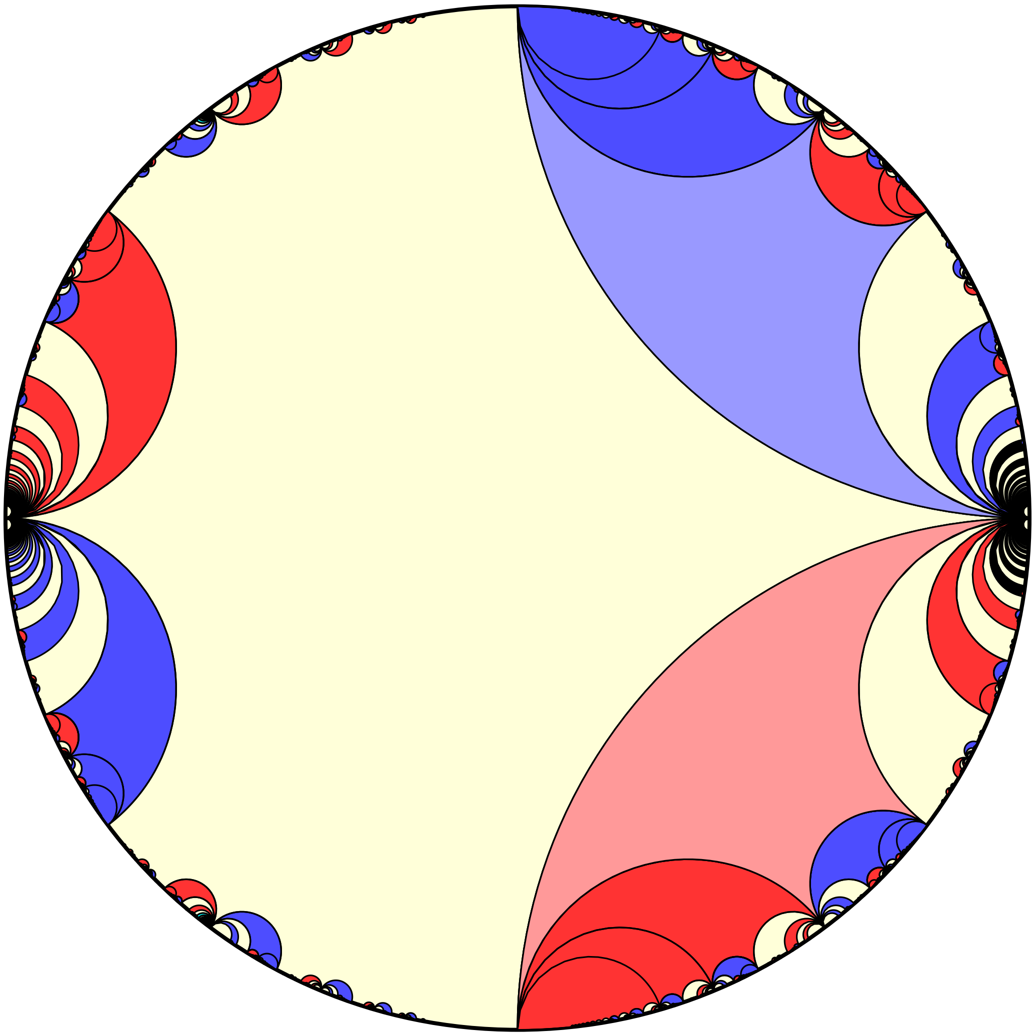

Figure 3. The weak Sierpiński set viewed as the set of colored tiles. Removing the light red tile leads to a set congruent to E via σ, while removing the light blue tile leaves a set congruent to E by τ.

For a Mycielski Hotel with rooms of two types, the red squares would be the smoking rooms and the blue squares the nonsmoking rooms. Then σ is an isometry whose application sends rooms to other rooms, preserves all adjacency relationships, and frees up the one smoking room labeled σ; similarly a transformation by τ frees up the nonsmoking room labeled τ. An animation of the disappearance of either room is available at [8].

Now, if a large family should arrive requiring a linear sequence of rooms such as (smoking, non-smoking, nonsmoking), then they can be accommodated in the Mycielski Hotel as follows. Using the transformation $\sigma\tau^2$, we have

$$\tau(E) = E \setminus \{\tau\}$$
$$\tau^2(E) = \tau(E \setminus \{\tau\}) = \tau(E) \setminus \tau^2 = E \setminus \{\tau, \tau^2\}$$
$$\sigma\tau^2(E) = \sigma\left(E \setminus \{\tau, \tau^2\}\right) = \sigma(E) \setminus \{\sigma\tau, \sigma\tau^2\} = E \setminus \{\sigma, \sigma\tau, \sigma\tau^2\}$$

So this transformation releases the three desired rooms σ, $\sigma\tau$, and $\sigma\tau^2$. This idea works for any linear sequence, but the hotel is not truly universal: room adjacencies form a tree, so one cannot get a triangular array of three rooms, each adjacent to the other two rooms.

These ideas can also be used, in theory, to visualize a hyperbolic hotel that is invariant under the removal of any finite set of rooms [4; 7, Thm. 6.16]. That uses a set of generators that is countably infinite; they exist in F_2 via $\rho_n = \sigma^n \tau^n$. The words involved grow too quickly in length to yield a useful picture. But one can easily see how the generators lead to a weak Sierpiński set with 2 replaced by ∞, which is the concept underlying the Mycielski Hotel having infinitely many room types. Let E consist of all words ending in some ρ_n on the right, and let such a word be given type n. Then for any n, $\rho_n(E) = E \setminus \{\rho_n\}$ and so a customer requesting a room of type n can be accommodated.

In an attempt to avoid the widely separated lobes of Figure 3, one could let E consist of all the cells in Figure 2 lying right of the y-axis. Then $\sigma(E)$ consists of all words beginning on the left with $\sigma\tau$ or σ^2, and so is a subset of E. But this process frees up infinitely many rooms and so when one of them becomes reoccupied, infinitely many are left vacant.

Another way to look at the construction in $\mathbb{H}^2$ is to observe that the set of inverses of words defining E is the set of words having σ^{-1} or τ^{-1} on the left. This set is quite simple: it is represented by the tiles (Fig. 2) that are left of the y-axis, and so is a union of two half-planes. So E is the result of starting with those two disjoint hyperbolic half-planes and applying group inversion (not a distance-preserving map) to the corresponding words.

3. The Hausdorff Paradox

In 1914, Felix Hausdorff showed how to get a paradoxical set in the free product $G = \mathbb{Z}_2 * \mathbb{Z}_3 = \langle \sigma, \tau : \sigma^2 = \tau^3 = e \rangle$. He also found two free generators that are rotations of the sphere. These two results allowed him to show that there is no finitely additive, rotation-invariant measure on the sphere $\mathbb{S}^2$ in 3-space having total measure 1. Moreover, they were the critical steps in the Banach–Tarski Paradox, discovered ten years later. Also, Banach proved that such measures do exist in $\mathbb{R}^1$ and $\mathbb{R}^2$.

For the paradox, Hausdorff found a way to partition G into three sets A, B, C, so that

$$(*) \qquad \tau(A) = B, \quad \tau^2(A) = C, \quad \sigma(A) = B \cup C.$$

In short, A is simultaneously a half and a third of the group. This is a paradoxical situation: One can use $G = A \cup B \cup C$ and $\sigma A = B \cup C$ to partition G into $A_1 \cup A_2 \cup B \cup C$, where A_1, A_2 are G-congruent (denoted $\sim$) to B, C, respectively. Then $A_1 \sim B \sim A \sim B \cup C \sim C \cup A$ (the last by τ) so A_1 and B yield $A \cup B \cup C = G$, and the same holds for A_2 and C.

Hausdorff constructed the sets inductively in a way that gives preference to B when there is a choice. To our surprise, if one instead gives preference to C, the result is very different. So we focus on that method; here are the placement rules.

Start with $A = \{e\}$, $B = C = \emptyset$ and consider τ^2 as an atom when dealing with word length (so τ^2 is a word of length 1). Working inductively by length, any unassigned word w has the form σu, τu, or $\tau^2 u$, where u has been placed.

• If w is τu or $\tau^2 u$, place w as forced by $(*)$: if $u \in A$, place τu into B and $\tau^2 u$ into C, and similarly moving cyclically if u is in B or C.

• If $w = \sigma u$, place w into A if $u \in B \cup C$ (forced) and into C if $u \in A$.

The C-preference occurs in the only unforced move, the very last clause; Hausdorff used B at this stage.

There is a much more direct way to describe the placement of words. For $L = \sigma$, τ, or τ^2, let W_L be the set of words with L at the left end (recall that τ^2 is an atom, and so is in W_{τ^2}, not W_τ). One first tries simply $A = W_\sigma$, $B = W_\tau$, and $C = W_{\tau^2}$, which works except that the identity cannot be assigned so as to preserve $(*)$. To address this, we absorb e into A using powers of $\tau\sigma$ as follows, where $j = 0, 1, 2, \ldots$:

$$A = \{\text{all } (\tau\sigma)^j \text{ and all of } W_\sigma \text{ except } \tau^2(\tau\sigma)^j)\};$$
$$B = \{\text{all } \tau(\tau\sigma)^j \text{ and all of } W_\tau \text{ except } (\tau\sigma)^j\};$$
$$C = \{\text{all } \tau^2(\tau\sigma)^j \text{ and all of } W_{\tau^2} \text{ except } \tau(\tau\sigma)^j\}.$$

In short, powers of $\tau\sigma$ and their translates by τ or τ^2 are assigned directly, with all other words assigned according to their leftmost term. It is not hard to show that the sets defined by this direct algorithm, shown in Table 1, satisfy $(*)$ and, in fact, coincide with Hausdorff's C-preferred sets. In the table, entries from the direct assignment of powers are shown in gray.

A	B	C
e	τ	$\tau\tau$
$\sigma\tau$	$\tau\sigma\tau$	$\tau\tau\sigma\tau$
$\tau\sigma$	$\tau\tau\sigma$	σ
$\sigma\tau\tau$	$\tau\sigma\tau\tau$	$\tau\tau\sigma\tau\tau$
$\sigma\tau\sigma\tau$	$\tau\sigma\tau\sigma\tau$	$\tau\tau\sigma\tau\sigma\tau$
$\sigma\tau\tau\sigma$	$\tau\sigma\tau\tau\sigma$	$\tau\tau\sigma\tau\tau\sigma$
$\tau\sigma\tau\sigma$	$\tau\tau\sigma\tau\sigma$	$\sigma\tau\sigma$
$\sigma\tau\sigma\tau\tau$	$\tau\sigma\tau\sigma\tau\tau$	$\sigma\tau\sigma\tau\sigma\tau\sigma$
$\sigma\tau\tau\sigma\tau$	$\tau\sigma\tau\tau\sigma\tau$	$\tau\tau\sigma\tau\sigma\tau\tau$
$\sigma\tau\sigma\tau\sigma\tau$	$\tau\sigma\tau\sigma\tau\sigma\tau$	$\tau\tau\sigma\tau\tau\sigma\tau$
$\sigma\tau\sigma\tau\tau\sigma$	$\tau\sigma\tau\sigma\tau\tau\sigma$	$\tau\tau\sigma\tau\sigma\tau\sigma\tau$
$\sigma\tau\tau\sigma\tau\sigma$	$\tau\sigma\tau\tau\sigma\tau\sigma$	$\tau\tau\sigma\tau\sigma\tau\tau\sigma$
$\sigma\tau\tau\sigma\tau\tau$	$\tau\sigma\tau\tau\sigma\tau\tau$	$\tau\tau\sigma\tau\tau\sigma\tau\sigma$
$\tau\sigma\tau\sigma\tau\sigma$	$\tau\tau\sigma\tau\sigma\tau\sigma$	$\sigma\tau\sigma\tau\sigma$
$\sigma\tau\sigma\tau\sigma\tau\tau$	$\tau\sigma\tau\sigma\tau\sigma\tau\tau$	$\tau\tau\sigma\tau\tau\sigma\tau\tau$
$\sigma\tau\sigma\tau\tau\sigma\tau$	$\tau\sigma\tau\sigma\tau\tau\sigma\tau$	$\tau\tau\sigma\tau\sigma\tau\sigma\tau\tau$
$\sigma\tau\tau\sigma\tau\sigma\tau$	$\tau\sigma\tau\tau\sigma\tau\sigma\tau$	$\tau\tau\sigma\tau\sigma\tau\tau\sigma\tau$
$\sigma\tau\tau\sigma\tau\tau\sigma$	$\tau\sigma\tau\tau\sigma\tau\tau\sigma$	$\tau\tau\sigma\tau\tau\sigma\tau\sigma\tau$
$\sigma\tau\sigma\tau\sigma\tau\sigma\tau$	$\tau\sigma\tau\sigma\tau\sigma\tau\sigma\tau$	$\tau\tau\sigma\tau\tau\sigma\tau\tau\sigma$
$\sigma\tau\sigma\tau\sigma\tau\tau\sigma$	$\tau\sigma\tau\sigma\tau\sigma\tau\tau\sigma$	$\tau\tau\sigma\tau\sigma\tau\sigma\tau\sigma\tau$
$\sigma\tau\sigma\tau\tau\sigma\tau\sigma$	$\tau\sigma\tau\sigma\tau\tau\sigma\tau\sigma$	$\tau\tau\sigma\tau\sigma\tau\sigma\tau\tau\sigma$
$\sigma\tau\sigma\tau\tau\sigma\tau\tau$	$\tau\sigma\tau\sigma\tau\tau\sigma\tau\tau$	$\tau\tau\sigma\tau\sigma\tau\tau\sigma\tau\sigma$
$\sigma\tau\tau\sigma\tau\sigma\tau\sigma$	$\tau\sigma\tau\tau\sigma\tau\sigma\tau\sigma$	$\tau\tau\sigma\tau\sigma\tau\tau\sigma\tau\tau$
$\sigma\tau\tau\sigma\tau\sigma\tau\tau$	$\tau\sigma\tau\tau\sigma\tau\sigma\tau\tau$	$\tau\tau\sigma\tau\tau\sigma\tau\sigma\tau\sigma$
$\sigma\tau\tau\sigma\tau\tau\sigma\tau$	$\tau\sigma\tau\tau\sigma\tau\tau\sigma\tau$	$\tau\tau\sigma\tau\tau\sigma\tau\sigma\tau\tau$
$\tau\sigma\tau\sigma\tau\sigma\tau\sigma$	$\tau\tau\sigma\tau\sigma\tau\sigma\tau\sigma$	$\sigma\tau\sigma\tau\sigma\tau\sigma\tau\sigma$

Table 1. The sets *A*, *B*, *C*, using powers of $\tau\sigma$ to absorb the identity into *A*.

We saw in Proposition 2 that the isometries $\sigma(z) = -1/z$ and $\tau(z) = -1/(z+1)$ generate the free product we need for the algebraic paradox. So we can now visualize the paradox by using the Klein–Fricke tessellation of the hyperbolic plane via the isometries σ and τ (Fig. 4); we simply assign each triangle to whichever of *A*, *B*, *C* contains the word defining it in the tiling. This was done by Mycielski and Wagon in 1984 for Hausdorff's *B*-preferred method [7, Chap. 5]; see Figure 8. The results for the *C*-preferred method used here turn out to be quite different; Figure 5 shows the three sets, where the powers of $\tau\sigma$ appear as the connected set in a lighter shade of red.

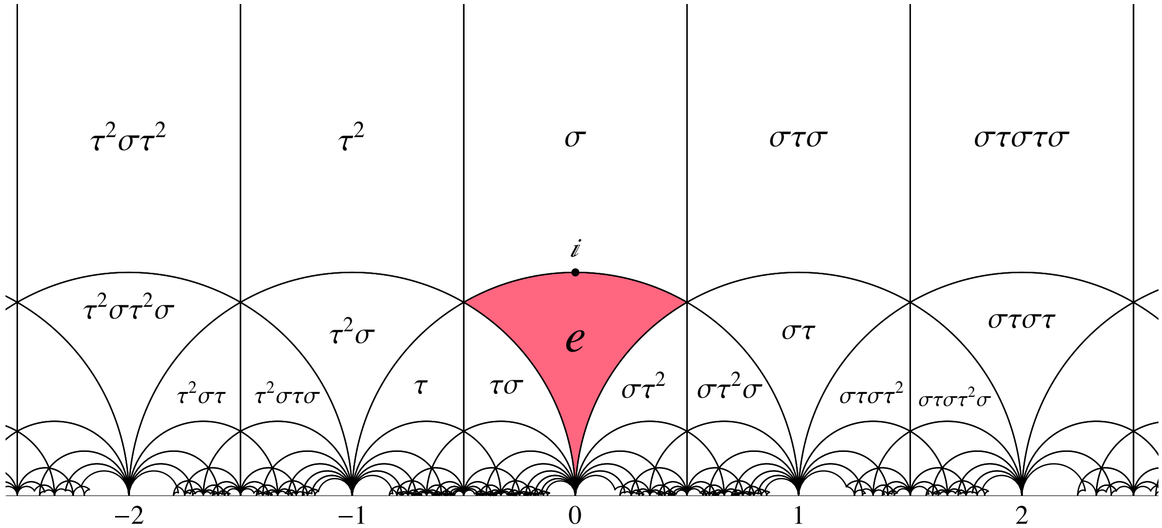

Figure 4. The Klein–Fricke tessellation of the hyperbolic plane using $\mathbb{Z}_2 * \mathbb{Z}_3$.

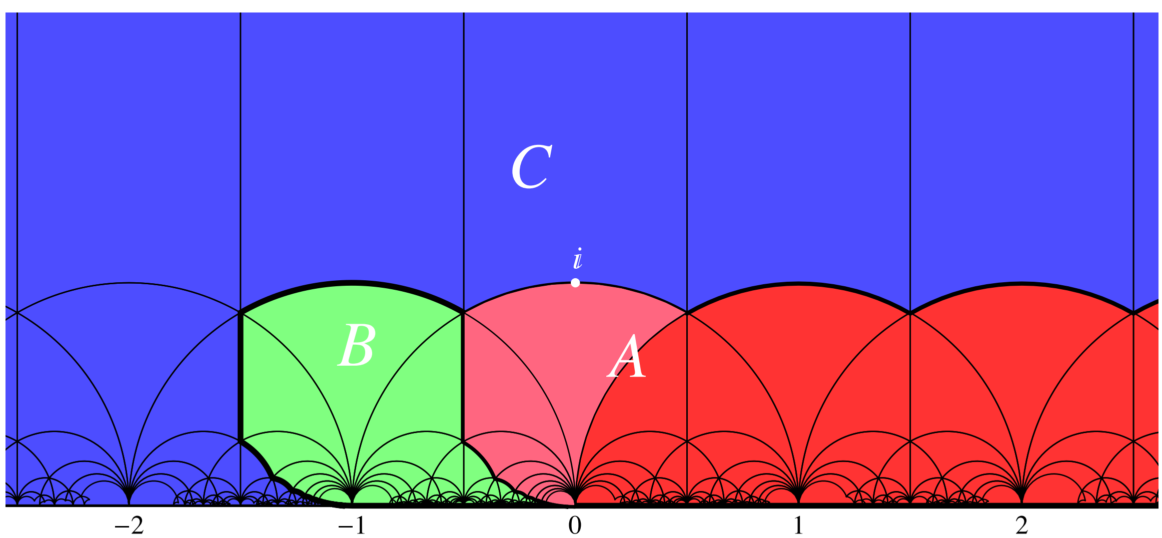

Figure 5. The sets *A*, *B*, *C* of a Hausdorff paradox viewed in the tessellation.

When we wrap the half-plane into a Poincaré disk, we get the pleasing symmetry of Figure 6. The projection in Figure 6(a) is not the standard one, but rather one that places the triple point $-\frac{1}{2}+\frac{\sqrt{3}}{2}i$ at the origin. That turns τ into a clockwise Euclidean rotation of 120°, thus making evident the congruence of *A* to *B* to *C*. The projection (b) is the standard one (it uses the Cayley transformation $(z-i)/(z+i)$ and places i at the origin); in this view, σ becomes a Euclidean half-turn and it is evident that *A* is congruent to $B \cup C$. Note that in (b) nothing is changed except the viewpoint: It is as if one flies over the plane to look straight down at a different region.

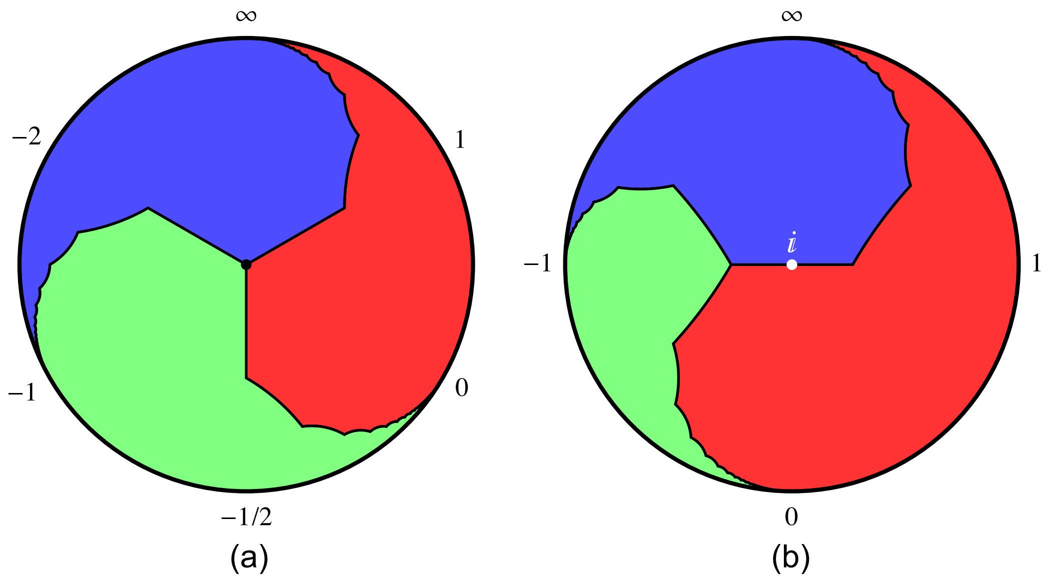

Figure 6. Two views of a Hausdorff paradox in the hyperbolic plane. (a) The three sets are congruent. (b) Changing the viewpoint shows that the red set is congruent to the combined green and blue sets. The labels refer to corresponding points in the upper half-plane model.

The sets are connected, and have a striking resemblance to a triple yin-yang motif (Fig. 7), which one might call a tian-di-ren figure, a Chinese phrase that refers to the heaven-earth-human trichotomy. Indeed, the knots on the boundaries in Figure 6 lie *exactly* on the semicircles in the tian-di-ren figure. But when the focal point is changed (Fig. 7(b)) as was done in Figure 6(b), the red region's area is about 99.7% of half the disk: because the origin is in the exterior of the red region, when red is reflected, a small lens around the origin is in neither red nor its reflection. The curved boundary is the problem; in the true paradox the boundary near the origin is perfectly straight.

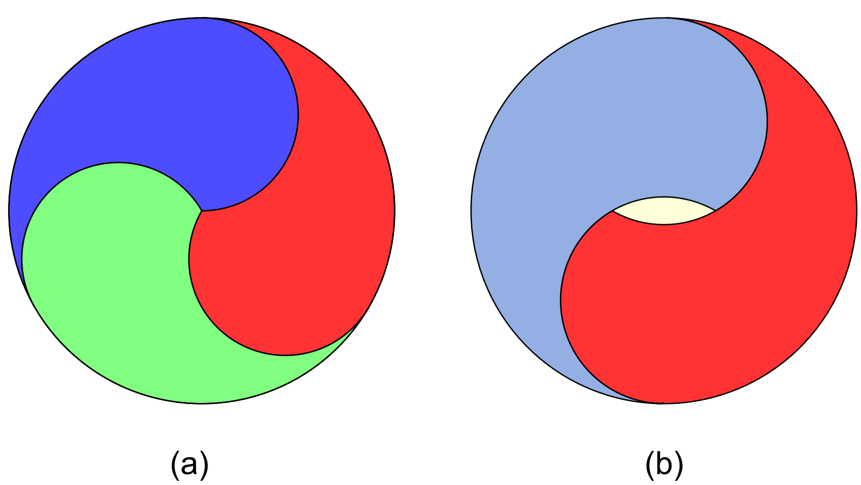

Figure 7(a). A tian-di-ren motif made from three semicircles. (b). The hyperbolic transformation of the red region as in Figure 6(b); reflecting red in the origin picks up less than its complement.

One can look at one of the tian-di-ren regions, the red one, say, in the upper half-plane to see exactly how it compares to the red region of the Hausdorff paradox. Figure 8 shows the two regions: they differ by two infinite families of circular caps. The leftmost horizontal cap corresponds to the lower half of the yellow lens in Figure 7(b).

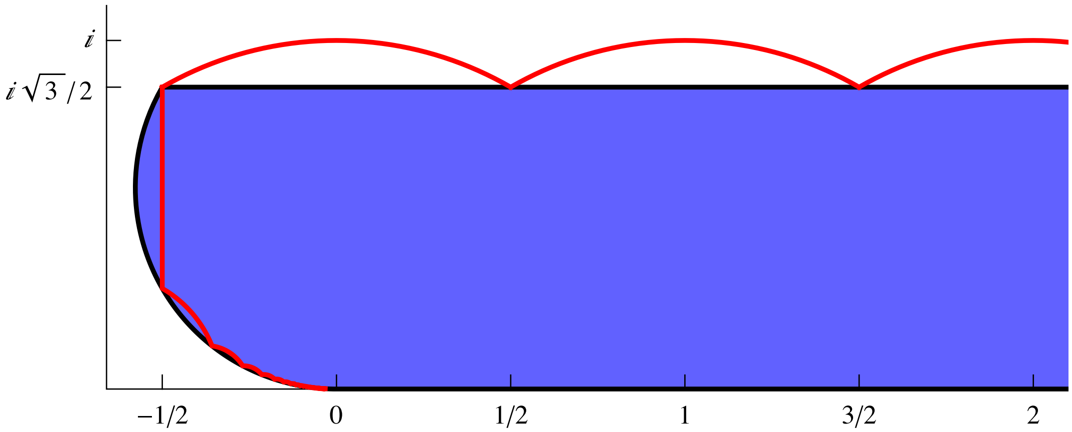

Figure 8. The blue region, bounded by a circular arc and straight lines, defines the tian-di-ren red region in the upper half-plane; the red curve bounds the red Hausdorff region from Figure 5.

The original *B*-preferred Hausdorff paradox also has a direct algebraic interpretation, where one uses powers of $\tau^2\sigma$ (as opposed to $\tau\sigma$) to absorb the identity. This leads to the visualization in Figure 9, where *A* is red and the sequence of powers is the tail of red triangles at upper left.

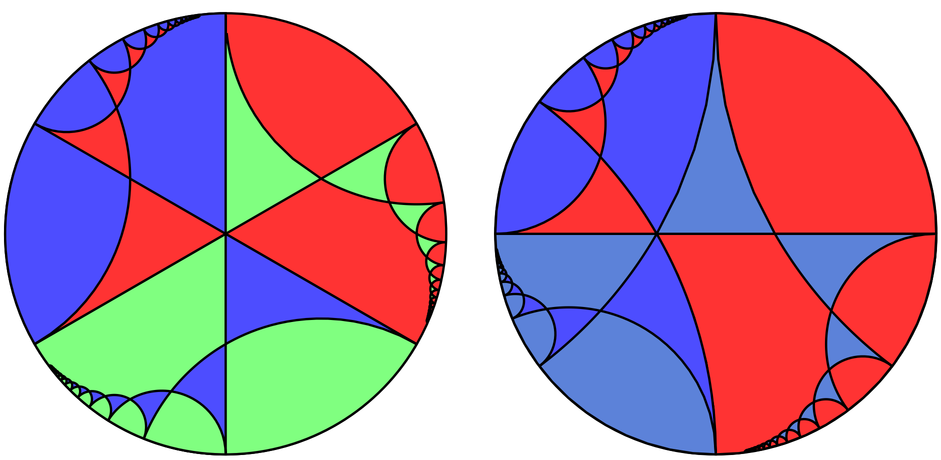

Figure 9. A view of Hausdorff's original paradox, using the same viewpoint change as in Figure 6. This *B*-preferred approach leads to disconnected sets.

These constructions have a measure-theoretic consequence: There is no isometry-invariant, finitely additive measure on the Borel subsets of $\mathbb{H}^2$ having total measure 1. The paradox using the Klein-Fricke tiling is built from open triangles and so, because the boundaries are unassigned, does not strictly prove this nonexistence. But it is not hard [7, chap. 5] to use the free group and tiling of §1 to build a paradox of the same type, but where every single point is taken into account. Similar measures on the Borel subsets of $\mathbb{R}^n$ do exist in every dimension [7, Thm. 11.15].

In 2000, Curtis Bennett [1] showed how to apply these ideas to color the angels and devils of M. C. Escher's famous woodcut using red, green, and blue so as to illustrate Hausdorff's paradox. That required finding a paradoxical partition in the orientation-preserving triangle group (also called a von Dyck group): $D(3, 4, 4) = \langle \sigma, \tau : \sigma^3 = \tau^4 = (\sigma\tau)^4 = e \rangle$. Figure 10 shows such a coloring: the left image shows that the red set is a third of the plane, while the viewpoint change at right shows the red set as one half of the plane. We thank Bennett for sharing the devil shape used in this image.

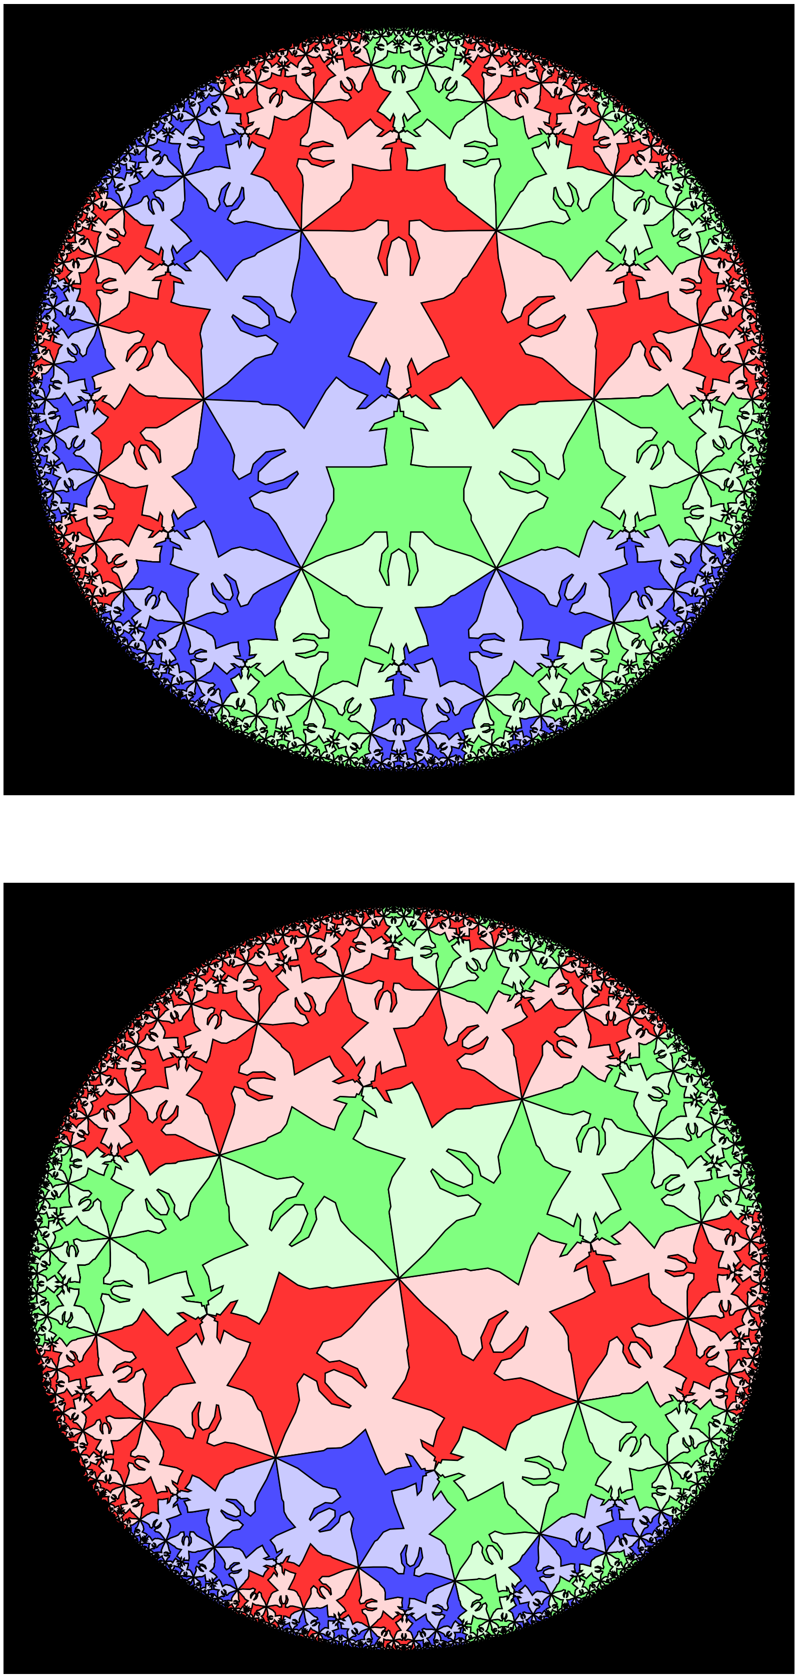

Figure 10. A paradox that combines ideas of Escher and Hausdorff.

4. Rationalizing the Paradox

Very often the Axiom of Choice is criticized for allowing such strange constructions as the classic Banach–Tarski Paradox of the Euclidean sphere or ball. The paradox in $\mathbb{H}^2$ does not require AC, but it involves the entire space, so leads only to an $\infty = 2\,\infty$ equation from the view of classic measure theory. Yet there are two other noteworthy ways to eliminate AC from such paradoxes.

K. Satô [3] found that the two 3-dimensional rotations given by

$$\frac{1}{7}\begin{pmatrix} 6 & 2 & 3 \\ 2 & 3 & -6 \\ -3 & 6 & 2 \end{pmatrix} \quad \text{and} \quad \frac{1}{7}\begin{pmatrix} 2 & -6 & 3 \\ 6 & 3 & 2 \\ -3 & 2 & 6 \end{pmatrix},$$

generate a free group G that acts on $\mathbb{S}^2_{\mathbb{Q}}$, the set of rational points on the unit sphere, and has no nontrivial fixed points on the rational sphere. A proof can be given along the lines of the proof of Proposition 2, but working modulo 7. Therefore G acts without fixed points on $\mathbb{S}^2_{\mathbb{Q}}$ (which is dense in $\mathbb{S}^2$; stereographic projection to $\mathbb{R}^2$ preserves rational points), and therefore the rational sphere admits a paradoxical decomposition. In fact, the sets of the paradox can be constructed explicitly. But visualization is difficult since each piece is dense in the sphere. Moreover, Satô's action, restricted to a single orbit, can be used to construct a Sierpiński set on the sphere. It too will be dense in the sphere.

Another constructive approach comes from the important work of Dougherty and Foreman [2], who found that the Banach–Tarski paradox is possible using pieces having the property of Baire (sets that are the union of a Borel set with a meager set). Their construction uses AC, but they can avoid AC to get results such as: There are finitely many disjoint open subsets of the unit ball in $\mathbb{R}^3$ which can be rearranged by isometries to form a set dense in a much larger ball.

A very famous theorem of Alfred Tarski states that a group has a paradoxical decomposition (with respect to its action on itself) iff it is not amenable. A group G is *amenable* if there is a measure on $\mathcal{P}(G)$ that is finitely additive, G-invariant, and of total measure 1. But the connection with free groups is much more complex. The famous group $B_{2,665} = \langle x, y : w^{665} = e \rangle$, which was proved by S. Adian to be infinite, and therefore a counterexample to Burnside's Conjecture, is now known (again, by Adian) to be nonamenable, and therefore paradoxical. Of course it contains no free subgroup even of rank 1.

The connection between Sierpiński sets and free groups is much clearer, as we have the following theorem of Straus from 1958 [7, Thm 6.19].

Theorem 3 (E. G. Straus). A group has a Sierpiński set iff it has F_2 as a subgroup.

But Straus also showed that even the action of an Abelian group can lead to a weak Sierpiński set: Let G be generated by the shears $\begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}$; then G acts on $\mathbb{R}^2$ and $E = \{(0, 2^n), (2^n, 0) : n = 0, 1, 2, \ldots\}$ satisfies $\sigma E = E \setminus \{(1, 0)\}$ and $\tau E = E \setminus \{(0, 1)\}$.

The proof of Theorem 3 is not difficult, but it uses the full hypothesis and we do not know if the conclusion holds assuming only that either one of two specific points may be deleted.

Conjecture. If a group has a weak Sierpiński set, then it has a free non-Abelian subgroup.

We conclude with the following conjecture, which is certainly believable since an old theorem of Bieberbach tells us that the isometry group of $\mathbb{R}^3$ does not have a discrete free non-Abelian subgroup.

Conjecture. There is no discrete weak Sierpiński set in $\mathbb{R}^3$.

References

1. C. Bennett, A paradoxical decomposition of Escher's Angels and Devils (Circle Limit IV), *The Mathematical Intelligencer*, **22:3** (2000) 39–46.

2. R. Dougherty and M. Foreman, Banach–Tarski decompositions using sets with the property of Baire, *Journal of the American Mathematical Society,* **7** (1994) 75–124.

3. K. Satô, A free group acting without fixed points on the rational unit sphere, *Fundamenta Mathematicae,* **148** (1995) 63–69.

4. E. G. Straus, On a problem of W. Sierpiński on the congruence of sets, *Fundamenta Mathematicae,* **44** (1957) 75–81.

5. J. Mycielski, About sets invariant with respect to denumerable changes, *Fundamenta Mathematicae,* **45** (1958) 296–305.

6. J. Mycielski and S. Wagon, Large free groups of isometries and their geometrical uses, *L'Enseignement Mathématique* **30** (1984) 247–267.

7. S. Wagon, *The Banach–Tarski Paradox,* New York: Cambridge University Press, 1985.

8. S. Wagon, The disappearing hyperbolic squares, Wolfram Demonstrations Project, 2013, http://demonstrations.wolfram.com/TheDisappearingHyperbolicSquares