A Porous Aperiodic Decagon Tile

Duane A. Bailey[1] and Feng Zhu[2]

Abstract. We consider the development of a single *universal* aperiodic prototile that tiles the plane without overlap. We describe two aperiodic approaches to constructing tiles that convert overlapping regions of Gummelt's decagon cover to sponge-like point-sets in $\mathbb{R}^2$. One tile has measure zero, the other has positive measure everywhere. Many characteristics of the decagon cover are inherited by these tilings.

1 Introduction

In 1996 Gummelt[6] identified a single decagon that *covered* the plane aperiodically (see Figure 1). The covering process involved overlapping tiles in a small number of ways. The construction, though not a proper tiling, is appealing to physical chemists who believe that the overlap models the overlapping neighborhoods of local influence that seem necessary in the perfect growth of aperiodic physical structures (called *quasicrystals*). The decagon is the direct result of considering a theorem of Conway that suggests that a decagonal *cartwheel* patch of Penrose kites and darts obeys not only a local isomorphism property, but a stronger covering condition.

The study of how single nontraditional tiles (shapes that violate one or more rules of prototile construction) perfectly tile the plane seems likely to shed light on the nature of aperiodicity and its analogy with physical systems[7]. In 1997 Bandt and Gummelt[1] and Gelbrich[3] demonstrated that set of Penrose tiles modified to include fractal edges was sufficient to remove the local matching condition.

Our approach is to develop a single decagon-shaped tile with porous interior (ie. a *sponge*) that allows mutual non-overlapping entanglement of adjacent tiles where overlap would occur in the otherwise analogous decagons of Gummelt. We consider two constructions: a lacy tile that tiles the plane densely everywhere with points almost nowhere (ie. with zero measure), and a hefty tile that tiles densely with positive measure everywhere. In both cases, a perfect tiling fails to cover all points of the plane, but our positive measure construction appears to come as near as possible.

2 The Cartwheel and Decagon Covers

A beautiful structure that regularly appears in Penrose's aperiodic tiling by kites and darts is Conway's *cartwheel* (see Figure 1). While all finite patches appear near isomorphic regions, following observation about the cartwheel is particularly important[2, 5]:

Theorem 1 *Every perfect aperiodic tiling by kites and darts can be covered by cartwheel patches.*

With this result in hand, Gummelt constructed a decagon tile (see Figure 1) whose coverings are in one-to-one correspondence with perfect kite-and-dart tilings. In any cover of the plane by cartwheel patches (or decagon tiles[3]) there are four types of interaction:

1. The patches do not overlap, or
2. The patches meet edge-to-edge, overlapping in an area of measure 0, or
3. The patches overlap in a manner that covers 4 darts and 7 kites (approximately 28.14% of the area of each of the participating tiles), or

[1]bailey@cs.williams.edu, Department of Computer Science, Williams College, Williamstown, Massachusetts 01267.
[2]fzhu@marshall.usc.edu, Marshall School of Business, University of Southern California, Los Angeles, CA 90089-0808.
[3]We will often refer to cartwheel patches and decagon tiles interchangeably.

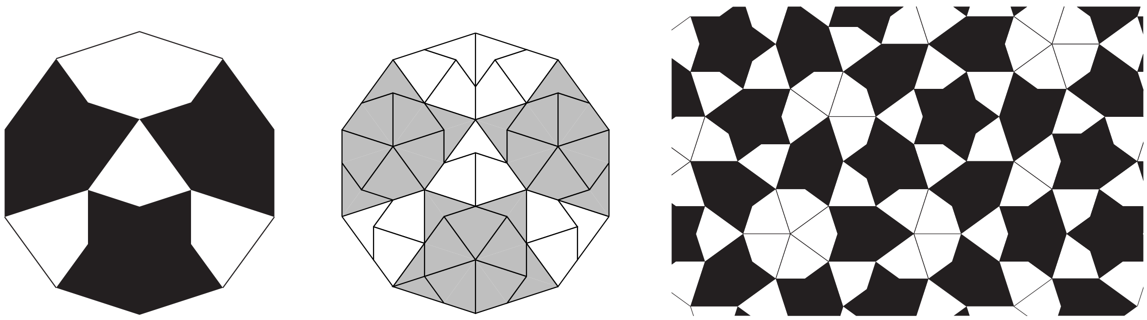

Figure 1: Gummelt's decagon tile (left), the cartwheel patch (middle, shaded to show correspondence with decagon), and a region of a decagon cover (right).

4. The patches overlap in a manner that covers 7 darts and 14 kites (approximately 54.46% of the area of each of the participating tiles).

Gummelt refers to these nontrivial types of overlap as Type A and Type B (see Figure 2). A bit of experimentation with the cartwheel identifies 4 orientations that lead to Type A overlap, and one that leads to Type B overlap (see Figure 3). If two cartwheels meet edge-to-edge then there exists a third cartwheel that (1) includes the edge and (2) consistently overlaps the adjacent patches in a non-trivial Type A or Type B overlap. Thus, we see there are notions of cover minimization that have no analog in tilings.

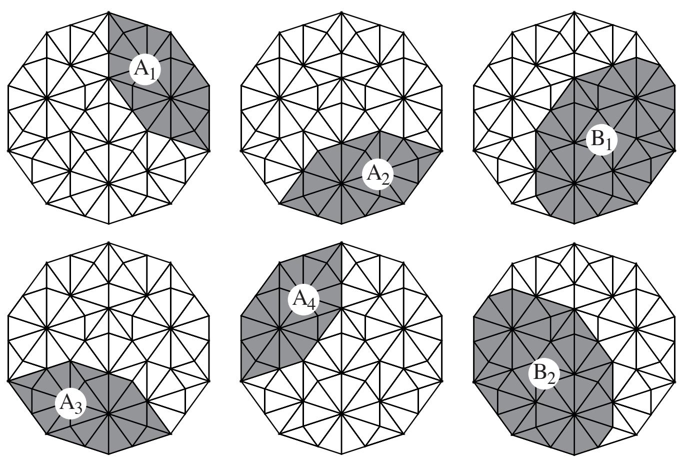

Figure 2: The six portions of the cartwheel that may overlap other cartwheel patches. Kites and darts have been divided into Robinson's tiles.

Gummelt ingeniously shaded the decagon to force similar overlapping rules: two decagons are allowed to overlap if they are consistently shaded (both black, or both white) on the set of overlapping points. This notion of overlap bends the traditional notions of a tile's edge. There are, essentially, two types of edges: an external edge that defines the boundary of the decagon, and an internal edge that is the potential image of an external edge of an overlapping decagon. When the tiles are used to cover the plane in a manner that respects these rules, the cover is aperiodic (see Figure 1). The proof of aperiodicity is fully developed by Gummelt[6].

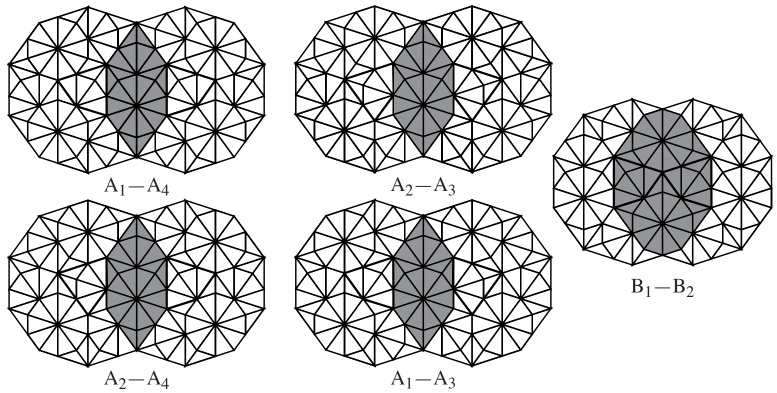

Figure 3: Legal overlaps of two cartwheels.

3 The Decagon Sponge Tiles

We seek to construct a decagon tile that tiles the plane, but without significant overlap. Our approach is similar to that of Gummelt: we begin with the cartwheel patch and develop a single tile that has similar tiling behavior. The construction, however, violates another traditional tile concept: the tile is a decagon-bounded collection of porous, Cantor-like sets (*sponges*).

We begin by considering the potential overlap of a pair of cartwheel tiles. Figure 4 (left) depicts the regions of the cartwheel that act similarly during any of the five different Type A or Type B overlaps. In many cases two or more regions of one decagon tile are covered by one region of another. If two points appear in different regions, they are covered by distinct regions for some overlap. For each of the five types of overlap Table 5 describes the alignment of regions in the participating tiles. If two regions are mapped to one another, we hope our tile's overlap, if any, is restricted to the edges. Given this anti-dependence between regions, it is possible to 5-color the cartwheel. We suggest a coloring on the right of Figure 4.

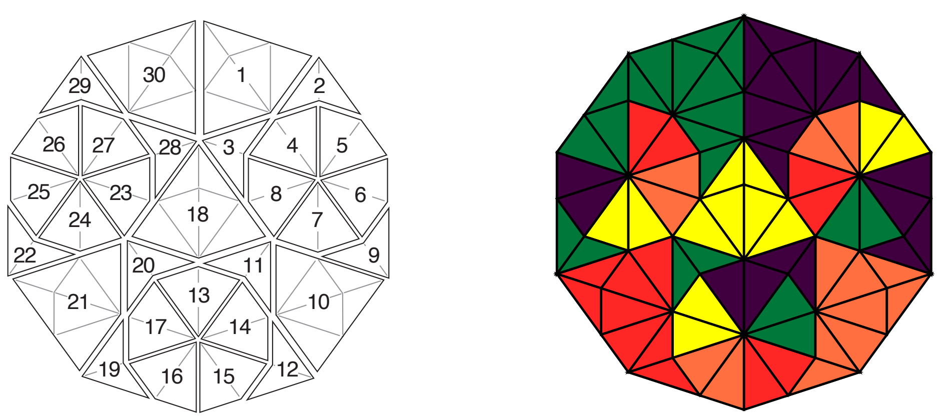

Figure 4: Left, regions of the cartwheel patch distinguished by overlap behavior and right, coloring of the cartwheel that ensures only overlaps of different colors.

It remains to develop tile shapes that (1) cover the indicated regions when no overlap occurs in that region and (2) allow the joint covering of the indicated regions when overlap does occur. It is difficult to see how a single shape can accomplish both tasks. Our approach is to develop independent sets of points that are dense in the specified regions. Each "color" maps to a two-dimensional sponge that has sufficiently

$A_1 \leftrightarrow A_3$	$A_1 \leftrightarrow A_4$	$A_2 \leftrightarrow A_3$	$A_2 \leftrightarrow A_4$	$B_1 \leftrightarrow B_2$	
1 $\leftrightarrow$ 21	1 $\leftrightarrow$ 30	10 $\leftrightarrow$ 21	10 $\leftrightarrow$ 30	3 $\leftrightarrow$ 22	12 $\leftrightarrow$ 11
2 $\leftrightarrow$ 20	2 $\leftrightarrow$ 28	11 $\leftrightarrow$ 19	11 $\leftrightarrow$ 29	4 $\leftrightarrow$ 25	13 $\leftrightarrow$ 17
3 $\leftrightarrow$ 21	3 $\leftrightarrow$ 29	12 $\leftrightarrow$ 20	12 $\leftrightarrow$ 28	5 $\leftrightarrow$ 26	14 $\leftrightarrow$ 13
4 $\leftrightarrow$ 17	4 $\leftrightarrow$ 27	13 $\leftrightarrow$ 16	13 $\leftrightarrow$ 26	6 $\leftrightarrow$ 27	15 $\leftrightarrow$ 14
5 $\leftrightarrow$ 13	5 $\leftrightarrow$ 23	14 $\leftrightarrow$ 17	14 $\leftrightarrow$ 27	7 $\leftrightarrow$ 23	16 $\leftrightarrow$ 15
6 $\leftrightarrow$ 14	6 $\leftrightarrow$ 24	15 $\leftrightarrow$ 13	15 $\leftrightarrow$ 23	8 $\leftrightarrow$ 24	17 $\leftrightarrow$ 16
7 $\leftrightarrow$ 15	7 $\leftrightarrow$ 25	16 $\leftrightarrow$ 14	16 $\leftrightarrow$ 24	9 $\leftrightarrow$ 28	18 $\leftrightarrow$ 21
8 $\leftrightarrow$ 16	8 $\leftrightarrow$ 26	17 $\leftrightarrow$ 15	17 $\leftrightarrow$ 25	10 $\leftrightarrow$ 18	20 $\leftrightarrow$ 19
9 $\leftrightarrow$ 12	9 $\leftrightarrow$ 22	19 $\leftrightarrow$ 12	19 $\leftrightarrow$ 22	11 $\leftrightarrow$ 20	

Figure 5: The association of regions under each of the 5 different overlaps.

small measure (or more evocatively, sufficiently high *porousness*) to allow the non-overlapping integration of differently colored sponges.

In the next sections we describe several techniques for constructing porous tiles that are closely related to the cartwheel and decagon covers.

3.1 A Lacy Fixed Point Decagon Tile

One method for constructing a tiling is to use *tile decomposition*. For each tile, a single step decomposes the tile into smaller tiles of similar shape. Given a coloring of prototiles, some points may be guaranteed to be colored the same after some fixed number of decompositions. This "fixed point" feature may be used to generate equivalence classes of points that may be associated with colors. The appropriately colored equivalence class of points may be used in each tile of a colored region to construct a single "tile" with tiling properties that are analogous to the decagon cover.

To see this we consider the expansions of the Robinson tiles (that is, tiles that correspond to half kite and half dart tiles). We now develop a dense point-coloring process. First, we color regions of the Robinson S and L tiles (see Figure 6). For our purposes, we need five equivalence classes, thus we partition the two tiles into five colors, conveniently picking division lines along tile boundaries in the decomposition. (Our partitioning decision is somewhat arbitrary. More colors or alternate partitioning could be used without impact.) After three or more decomposition steps, the image of each colored region contains at least one similarly shaped subtile that is similarly colored; within these regions we will find colored fixed points. Since the existence of these points is independent of the context of the decomposition, each tile in any decomposed patch of tiles provides additional fixed points of each of the five colors. Because each fixed point is introduced at some finite stage of decomposition, it is clear that each of the fixed points is uniquely colored at all but a finite number of steps of decomposition. In fact, fixed points of each color are dense in the plane—they reside within each open ball. We construct a porous decagon tile by selecting, for each colored region, the equivalence class of fixed points of the appropriate color. An approximate figure, suitable for all arguments we make here, is found in Figure 7. When two decagon tiles are brought together to overlap, the uniqueness of the decomposition of overlapping regions ensures that different colors will be represented by disjoint sets of fixed points.[4]

Such tiles are quite porous: since the number of points that are used to represent the decagon tile are only countable, they have measure zero. This new decagon-shaped collection of points covers almost nothing, so these tiles—essentially structured dust—easily "pass through" each other. It is then necessary to introduce explicit edges in the same locations as the interior and exterior edges of the decagons of Gummelt. The matching rules, then are:

Rule 1. Each segment of an exterior edge must be covered by a segment of (interior or exterior) edge from

[4]Readers may find it useful to print a dozen copies of the decagon tile pattern—either black and white or color—on transparencies and experiment with their interaction. Kits can be found at `http://www.cs.williams.edu/~bailey/porous`.

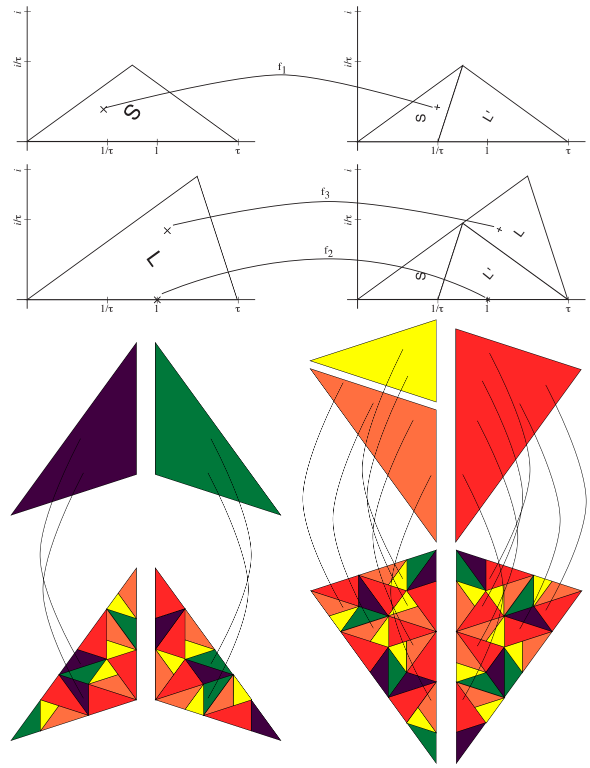

Figure 6: Above, the self-similar relationship between Robinson's S (half dart) and L (half kite) tiles. Mappings of left-handed tiles (L' and S') are mirrored. Below, the division of tiles into five colors (above) and the tile as it is decomposed after three steps. The associated points are fixed in this mapping.

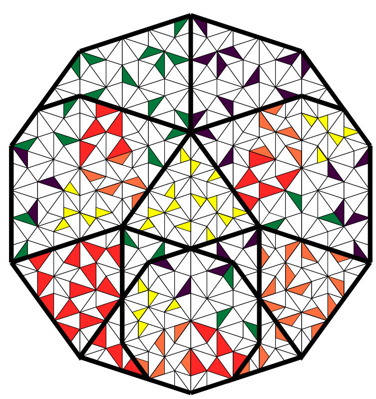

Figure 7: The first approximation to the coloring of the decagon tile. Colored tiles identify the location of the first fixed points in the region and suggest the interleaving of the tiles. Internal and external edges appear bold. Though not dense, this approximation has the same tiling properties as the ideal lacy decagon.

some other tile, and

Rule 2. No non-edge point of a tile may be coincident with a point of another tile.

The required overlap of edges forces the alignment of tiles using Type A or Type B regions. Elsewhere we fully demonstrate that the matching rules of the lacy decagon tile are equivalent to the decagon matching rules.

As a tiling mechanism this prototile could be improved. In the next section we introduce a means of coloring points that provides more "heft" to the prototile.

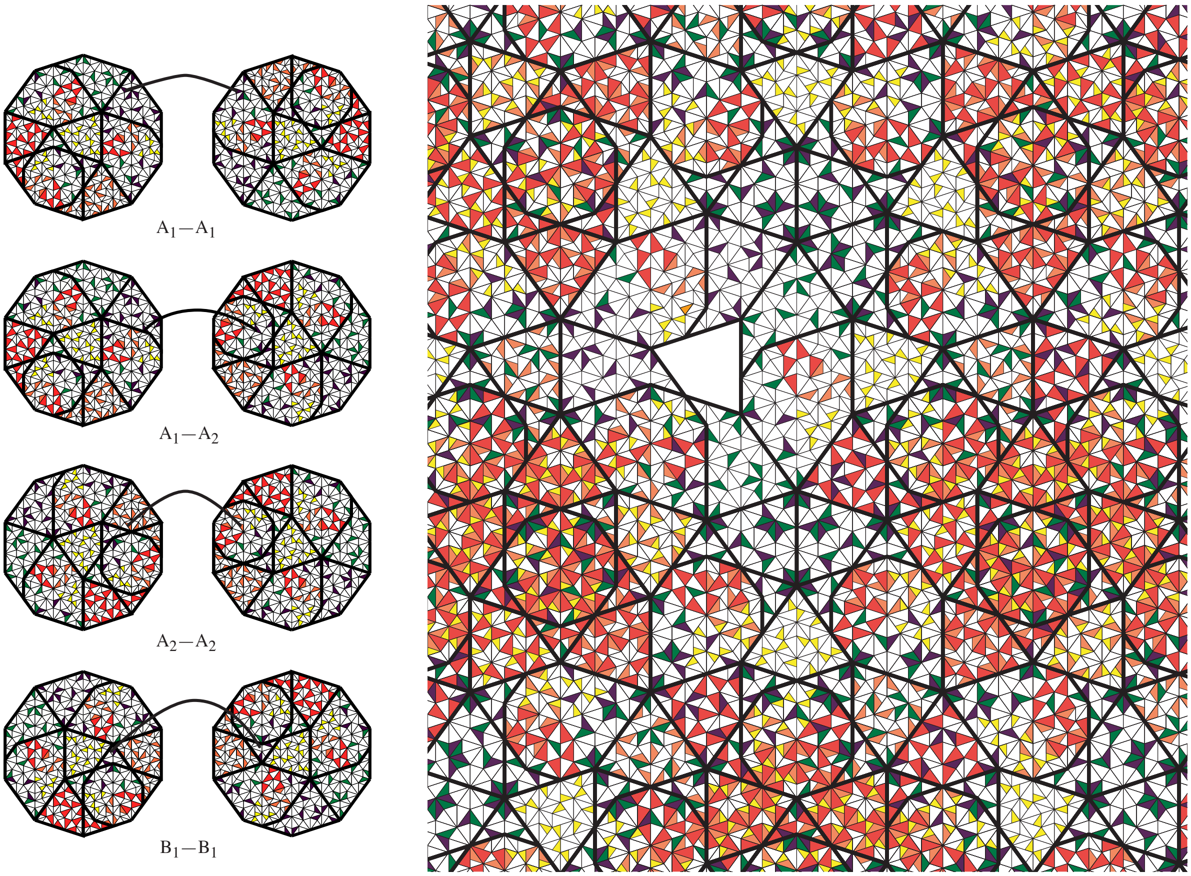

Figure 8: Left, locations of overlapping colored points in various representative illegal overlaps. Right, the beginnings of a consistent decagon tiling. The hole in the center can be covered with a tile in the obvious manner.

3.2 A Hefty Cantor-Like Decagon Tile

A general method for increasing the measure of our tile in any particular neighborhood depends on the notion of constructing multiple *well-mixed* sets of points: sets that are non-zero measure and dense in the same region.

We first consider the simpler problem of constructing these sets in one dimension[9], and then lift the results to two dimensions.

3.2.1 Constructing Pairs of Well-Mixed Cantor Sets

A common example from measure theory is the Cantor set, $\mathcal{C}(I)$, on an interval I of $\mathbb{R}$. We begin with the first approximation by intervals, $\mathcal{C}_0(I)$ a set containing only the interval $I = [l, r]$. We may then recursively define successive approximations $\mathcal{C}_i(I)$, by removing the "middle thirds". The result, $\mathcal{C}(I)$, is "dust": a set of points with measure zero. For our purposes, we consider a less aggressive Cantor-set-like construction. This version leaves behind dust, but the measure of the portion removed is much smaller, allowing the dust to accumulate. Our first approximation to the set is, $\mathcal{C}'_0(I) = \{I\}$ where $I = [l, r]$. We then define

$$\mathcal{C}'_i(I) = \bigcup_{[a,b]\in\mathcal{C}'_{i-1}(I)} \left\{ \left[a, a + \frac{(b-a)(3^i - 1)}{2 \cdot 3^i}\right], \left[b - \frac{(b-a)(3^i - 1)}{2 \cdot 3^i}, b\right] \right\}$$

The first step removes the middle third. The second step removes the middle ninth of the two intervals remaining after step 1. The i^{th} step removes the middle 3^{-i} of the remaining 2^{i-1} intervals, and so forth. The portion that remains in the limit, $C'(I)$, is the set of all points of I that are mentioned by some interval of each step. For the interval $I = [l, r]$, the portion removed consists of

$$\mu(I \backslash C'(I)) = \left(\frac{1}{3} + \frac{1}{9} \cdot \frac{2}{3} + \frac{1}{27} \cdot \frac{8}{9} \cdot \frac{2}{3} + \cdots \right) \mu(I) \approx 0.44\mu(I)$$

It follows, then, that the non-aggressive Cantor dust, $C'(I)$, saves $\sigma = 0.56$ of the measure $\mu(I)$. This positive measure set still enjoys many of the properties of the original Cantor set. It is interesting to point out that while the middle-thirds Cantor construction $C(I)$ can be thought of as a composition of two sets that are similar to $C(I)$, that is not true for the non-aggressive construction, $C'(I)$.

An important goal in our construction of Cantor-like sets is to make them *well mixed*:

Definition 2 *Two sets $S \subset I$ and $S^c = I \backslash S$ are* well mixed *if for any open interval $B \subseteq I$, $\mu(S \cap B) \neq 0$ and $\mu(S^c \cap B) \neq 0$.*

Well mixed sets are dense in each other and have positive measure everywhere. Since the set of Cantor dust that remains is not dense in the set that is removed (consider, for example, small open intervals centered in any interval removed) the two sets are obviously not well mixed. Indeed, the intervals removed work against the mixing. We now focus on post-processing the removed intervals.

As a basic step in our construction, we depend heavily on the non-aggressive Cantor construction on interval I, $C'(I)$. Notice that this construction is equivalent to computing $C'([0, 1])$ and rescaling the result in the natural manner to fit within the interval $I = [l, r]$. In addition, the construction on an open or half-open interval differs from the construction on a closed interval by only one or two points.

We now develop a countable number of sets, S_i that, together, cover the points of interval $R_0 = I$. First, let $S_0 = C'(I)$, a set with measure $\sigma\mu(I)$. This construction removed a countable number of open "remaindered" intervals, $R_1 = I \backslash S_0$ that have total measure $(1 - \sigma)\mu(I)$. We hope to break up these remainders, to support the mixing. Thus, for $i > 0$ we define an *approximant*

$$S_i = \bigcup_{(a,b)\in R_i} C'((a, b))$$

This constructs dust from R_i with positive measure $\sigma(1-\sigma)^i\mu(I)$, and was derived by removing the countable open intervals $R_{i+1} = R_i \backslash S_i = I \backslash (S_0 \cup S_1 \cup \cdots \cup S_i)$ with total measure $(1 - \sigma)^{i+1}\mu(I)$.

The sets S_i are pairwise disjoint since they are defined on points of I remaindered from all previous steps. We may now define

$$S = \bigcup_{i=0}^{\infty} S_{2i}$$

and

$$S^* = \bigcup_{i=0}^{\infty} S_{2i+1}$$

Since the union operation here augments a Cantor set with dust from some of the removed area, this process is often called *refilling*. It is this refilling that eventually ensures the two sets are mixed. Sets S and S^* are disjoint, and avoid only a countable number of points in the original interval that correspond to endpoints a and b for each open interval removed during the Cantor set construction process. For our purposes, this countable collection of points is of marginal concern, and we note that these points augment S^* to provide: $S^c = I \backslash S$ The two sets S and S^c are well-mixed. While we have constructed S (S^*) by accumulating the even (odd) numbered S_i, our theorem holds for any partitioning of the sets that leaves a countably infinite number in each of the two piles. This approach can be used to achieve a measure for each of S and S^c that is arbitrarily close to equal division of the interval. The notions of well-mixed sets may be extended, as well, to any finite number of sets using similar techniques. A collection of n sets might be constructed, for example, by accumulating the stepwise approximates, S_j, into sets in round-robin order, $0 \leq i < n$: $S(i) = \bigcup_{j=0}^{\infty} S_{nj+i}$ We now make the analogy between open intervals and borderless Robinson tiles in $\mathbb{R}^2$.

3.2.2 The Well-Mixed Penrose Sponge

Our approach now is to construct different Cantor-like sponges from decomposed Robinson tilings. Each tile in a patch of Robinson tiles can be decomposed into similar tiles in a unique fashion. From this collection of subtiles we may elect to remove one or more leaving a portion of the original tile's decomposition. The remaining tiles are then decomposed in a similar manner with a portion of the subtiles removed. The process is then repeated until a sponge-like structure is developed.

As with the Cantor set in $\mathbb{R}$, removal of a fixed percentage of the subtiles found in the decomposition can cause the remaining dust to have measure zero. We seek a less aggressive removal process that leaves more "hefty" dust of positive measure. The essential concept is to remove a smaller percentage of the area at each step, based on the removal of tiles from deeper decompositions.

For the ease of controlling which subtiles are to be taken out after each decomposition, we define an *address* for each Robinson's tiles in the decomposition process. We denote Robinson's prototile types as we have elsewhere, with L and S tiles having left-handed versions denoted L' and S'. After decomposing a tile n times, we will obtain many subtiles, each identified by its containment in the stages of decomposition, read from left (before first decomposition) to right (after last decomposition). In Figure 9 the indicated tile is generated by decomposing Robinson's L tile three times: the tile resides in the S' tile that resulted from decomposing the L' tile, that resulted from decomposing the L' tile, that was a direct result of decomposing the original L. The address of the tile is, therefore, $LL'L'S'$. It is easy to see that addresses are in one-to-one correspondence with the tiles that appear through the decomposition process.

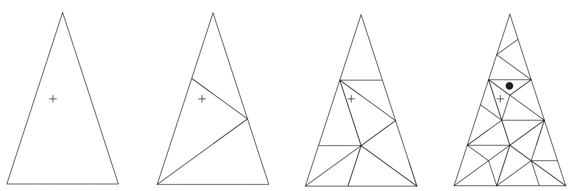

Figure 9: The address of the S tile marked with a + in patch at right is $LL'L'S'$. Each step is one decomposition of the original L tile that determines the address of the tile. Note, also, that there is only one tile whose address ends in 3 S terms (marked with a $\bullet$).

We are now prepared to algorithmically define the non-aggressive Cantor construction operation $K(P)$ on a patch of tiles P. Controlling the operation is process-specific parameter d that determines the minimum number of decompositions that occur at each step. Increasing d will allow us to arbitrarily reduce the total error in our computation of the measure of $K(P)$.

We begin by decomposing P d times and removing the set, R_1, of all the small tiles—that is all tiles whose d-long addresses end in a single S. Let $P_1 = P \backslash R_1$. Patch P_1 is expanded $d+1$ times, and we remove the small tiles (a set called R_2) appearing as the result of directly decomposing small tiles—that is all tiles whose $d+1$-long addresses end in 2 S's. The process of constructing S_{i+1} continues by removing, more and more selectively, the set R_{i+1} of small tiles whose addresses end in $i+1$ S's from what remains in S_i.

As d becomes large the measure of $K(P)$ converges quickly to $\rho\mu(P)$, where $\rho = 0.605741$. Thus, the $K(P)$ operation leaves approximately 61 percent of the tile intact. The error in ρ is bounded by the worst cases of consistently overestimating and consistently underestimating the amount removed at each stage, and has total error no worse than

$$\frac{5}{\tau^{4d+2}}\mu(P)$$

Notice that $P - K(P)$ is a countable collection of open, S tiles. This set is dense in $K(P)$, but the opposite is not true. To achieve mixing, as in the one dimensional case, we refill each of the small tiles with the appropriate K construction on those tiles. We set $R_0 = P$, and compute $S_0 = K(R_0)$. This set has measure $\rho\mu(P)$ and the remainder, $R_1 = R_0 \backslash S_0$, has $(1-\rho)\mu(P)$. At each stage, S_i is the result of applying K to the remainder set R_i generated by the previous step. $S_i = K(R_i)$ has measure $\rho(1-\rho)^i\mu(P)$, and each remainder has measure $(1-\rho)^i\mu(P)$.

In a round-robin manner described before for patch P and $0 \le i < 5$ the well-mixed set representing color i is

$$S(i) = \bigcup_{k=0}^{\infty} S_{5k+i}(P)$$

The measures of these sponges are easily computed, as

$$\mu(S(i)) = \frac{\rho(1-\rho)^i}{1-(1-\rho)^5}\mu(P)$$

This round-robin distribution of sets sponges among the colors is not very equitable: sponge 0 takes up 61 percent, while sponge 4 accounts for 1.5 percent of the patches they appear in. With a little effort, it is can be seen these sponges are well mixed.

We have, then, a means of constructing five different sponges that intermingle without overlap, with the additional advantage that each sponge appears with significant measure at all locations. Our second, heftier decagon, is constructed using points from these equivalence classes and inherits the same edges and matching rules from the lacy construction.

While the five sponges can be used to cover a patch, fewer than five tiles overlap in a single region. In these areas, a significant measure of points is not covered (corresponding to colors missing in the region). This seems to be an essential difficulty with the approach, although any point that is not covered is, of course, surrounded arbitrarily closely by points that are covered by each of the participating colors.

The general approach to developing porous tiles from overlapping shapes should, of course, generalize to higher dimensions. One difficulty, of course, is the development of overlapping polyhedra that cover the space in an appropriate manner.

3.2.3 Balancing Sponge Measures

In our hefty decagon we constructed sponges that are positive measure everywhere. One unfortunate feature is that their measures are dramatically different. Aside from being aesthetically unappealing, the ability to control the density of the different sponges may lead to more realistic models of interactions in physical systems. A first-fit bin-packing algorithm is sufficient to generate sponges with densities that are equal.

We start with five bins, b_i $0 \le i < 5$, each with capacity 0.2. Ultimately each bin will hold the approximants that will determine the component sponges of one color set. During our construction, $\rho > 0.2$ so P_1 will not fit within any bin. We modify, then, our K operation so that instead of removing S tiles, we remove tiles according to the number of L's in the tile address. This results in a proportion of $\rho_L = 0.0146634$ and, $\mu(P_0) = 0.0157$, $\mu(P_1) = 0.0154$, $\mu(P_2) = 0.0152$, ...

The bin-packing algorithm assigns the measure of each point set into an appropriate bin in a first-fit fashion:

```
PackBins:
    i ← 0
    repeat forever
        j ← 0
        while μ(b_j) + μ(P_i) ≥ 0.2
            j ← j + 1
        add set P_i to bin b_j
        i ← i + 1
```

In order for the algorithm to generate sponges with perfectly balanced measures, we need to ensure two things. First, the inner loop in the algorithm can always exit before $j = 5$. In the i^{th} step, we are trying to find a slot in one of the bins to accommodate 1.57% of the total unfilled area. Since we have 5 bins, there is at least one bin whose unfilled area is no less than 20% of the total unfilled area. Therefore, we are always able to find such a slot. As a result, the inner loop will always exit before $j = 5$.

Second, no bin is ever completely filled during the iterations; each bin will be filled by an infinite number of point sets. The condition $\mu(b_j) + \mu(P_i) \geq 0.2$ effectively prevents any bins from being completely filled during our iterations. For the proof of well mixing to work, we must be able to, at any stage in the bin packing process, be assured that we will place a set in each of the five bins within a finite number of steps.

There is, of course, nothing special about the relative sizes of the bins. If colors could be consistently aligned with regions of influence in a physical quasicrystal, their relative densities could be adjusted to better reflect, say the degree of influence on neighboring unit cells.

4 Conclusions

It is, of course, always intriguing to develop a shape with aperiodic behavior. We have found the overlapping nature of the decagon tile has been useful in two distinct respects. First, it seems that physical systems with decagonal long-range symmetry are more adequately modeled if we allow this overlap. The interpretation of the overlap is the sharing of molecular regions in a unit cell. Secondly, we believe the use of overlap has served as a catalyst to better understand how other nontraditional tile shapes interact aperiodically.

In the prototiles presented here we demonstrate how to use the inflation-based similarities of tiles to construct dense tiling by a single prototile that is either zero measure or positive measure everywhere. In the constructions presented here, formal matching rules include interior and exterior edges that must meet in the manner inherited from the decagon cover. These edges are clearly necessary in the lacy tile that covers almost nowhere. It is not known, however, whether these edge rules are strictly necessary in the heftier constructions, where the overlapping sponges may inherit the interlocking geometry of the self-similar decompositions.

We also demonstrate how to adjust the relative densities of tile regions, perhaps to more accurately model the influence of physical quasicrystal systems. Work by Jeong and Steinhardt[8] has studied the stoichiometry of the unit cells of quasicrystals with 10-fold long-range symmetry, and is greatly facilitated by the decagon tiling and its offshoots. Their techniques compute precise densities of atoms per unit area of the crystal. It seems likely that the simplicity of the equal-density construction of Section 3.2.3 would be suitable for many purposes, though other density distributions are possible.

References

[1] C. Bandt and Petra Gummelt. Fractal Penrose tilings I: Construction and matching rules. *Aequationes Mathematicae*, 53:295–307, 1997.

[2] Martin Gardner. Extraordinary nonperiodic tiling that enriches the theory of tiles. *Scientific American*, pages 110–121, January 1977.

[3] Götz Gelbrich. Fractal Penrose tiles II: Tiles with fractal boundary as duals of Penrose triangles. *Aequationes Mathematicæ*, 54:108–116, 1997.

[4] Branko Grünbaum and G. C. Shephard. Is there an all-purpose tile? *The American Mathematical Monthly*, 93:545–551, August/September 1986.

[5] Branko Grünbaum and G. C. Shephard. *Tilings and Patterns*, chapter 10. W. H. Freeman and Company, New York, 1987.

[6] Petra Gummelt. Penrose tilings as coverings of congruent decagons. *Geometriae Dedicata*, 62(1):1–17, August 1996.

[7] Hyeong-Chai Jeong and Paul Steinhardt. Constructing Penrose-like tilings from a single prototile and the implications for quasicrystals. *Physical Review B*, 55(6):3520–3532, February 1997.

[8] Hyeong-Chai Jeong and Paul J. Steinhardt. Rules for computing symmetry, density and stoichiometry in a quasi-unit-cell model of quasicrystals. Technical Report 9901014, arXiv ePrint, December 2002.

[9] Cesar Silva. *Invitation to Ergodic Theory.* Manuscript in progress, 2001.

[10] Feng Zhu. The search for a universal tile. Technical report, Williams College Department of Computer Science, 2002.