THE FRACTIONAL CHROMATIC NUMBER OF THE PLANE

DANIEL W. CRANSTON AND LANDON RABERN

ABSTRACT. The chromatic number of the plane is the chromatic number of the uncountably infinite graph that has as its vertices the points of the plane and has an edge between two points if their distance is 1. This chromatic number is denoted $\chi(\mathbb{R}^2)$. The problem was introduced in 1950, and shortly thereafter it was proved that $4 \leq \chi(\mathbb{R}^2) \leq 7$. These bounds are both easy to prove, but after more than 60 years they are still the best known. In this paper, we investigate $\chi_f(\mathbb{R}^2)$, the fractional chromatic number of the plane. The previous best bounds (rounded to five decimal places) were $3.5556 \leq \chi_f(\mathbb{R}^2) \leq 4.3599$. Here we improve the lower bound to $76/21 \approx 3.6190$.

1. INTRODUCTION

A proper coloring of the plane assigns to each of its points a color, such that points at distance 1 get distinct colors. The smallest number of colors that allows such a coloring is the *chromatic number of the plane*, denoted $\chi(\mathbb{R}^2)$. This problem was introduced in 1950, by Edward Nelson, a student at the University of Chicago. In the same year John Isbell, a fellow student, observed that $\chi(\mathbb{R}^2) \leq 7$. This upper bound comes from a result of Hadwiger [11], who showed that it was possible to partition the plane into hexagons, each of diameter slightly less than 1, and color each hexagon with one of seven colors, such that hexagons with the same color are distance more than 1 apart (see Figure 2).

The lower bound $\chi(\mathbb{R}^2) \geq 4$ comes from the observation of William and Leo Moser [16] that the graph in Figure 1(A) is a unit distance graph, i.e., it can be drawn in the plane with all edges of length 1. This graph is now known as the Moser spindle; since it has chromatic number 4, the lower bound follows. Around the same time [20, p. 19], Solomon Golomb
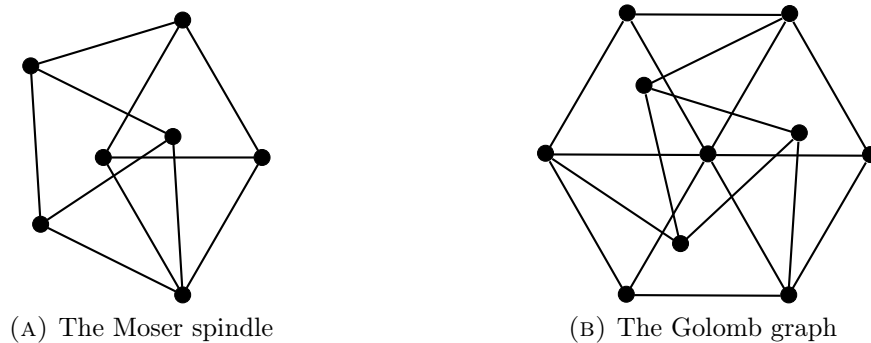

(A) The Moser spindle　　　(B) The Golomb graph

FIGURE 1. Two 4-chromatic unit distance graphs.
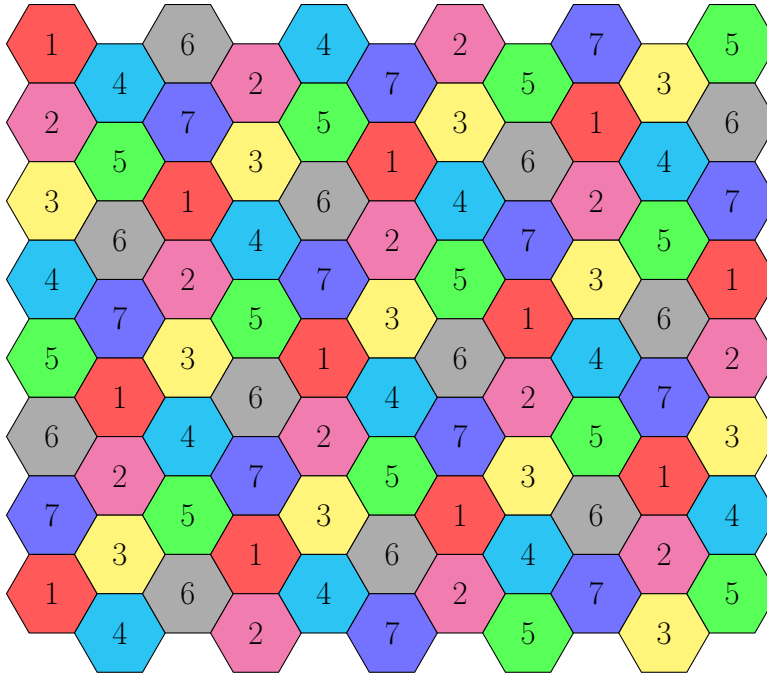
DANIEL W. CRANSTON AND LANDON RABERN

FIGURE 2. A 7-coloring of the plane.

discovered the unit distance graph in Figure 1(B), which also has chromatic number 4. The problem first appeared in print in 1960 in Martin Gardner's *Mathematical Games* column [9].

A seemingly helpful result of de Bruijn and Erdős [2] implies that the chromatic number of the plane is achieved by some finite subgraph (this proof does assume the Axiom of Choice). But unfortunately we have no reason to expect that such a subgraph will have fewer than (say) a billion vertices. By constructing a 6-coloring of nearly all of the plane, Pritikin [18] showed that if a 7-chromatic unit distance graph does exist, then it has at least 6198 vertices. As for the lower bound, Erdős wrote in 1985 [4, p. 4] "I am almost sure that $\chi(\mathbb{R}^2) > 4$."

The history of the problem has many more interesting twists than we can recount here, but Soifer records nearly all of them in his comprehensive and entertaining *The Mathematical Coloring Book* [20]. Although this problem has been widely popularized in the last half century, the best known bounds remain $4 \leq \chi(\mathbb{R}^2) \leq 7$.

The notion of fractional chromatic number was introduced in the early 1970s, with the goal of amassing more evidence in support of the Four Color Conjecture, or possibly disproving it. In a fractional coloring of a graph G, we assign to each independent set in G a nonnegative weight, such that each vertex appears in independent sets with weights summing to at least 1. The *fractional chromatic number*, $\chi_f(G)$, is the minimum sum of weights on the independent sets that allows such a coloring (for infinite graphs, take the infimum of this sum). This definition comes from solving the linear relaxation of the integer programming formulation of chromatic number. When G is infinite, $\chi_f(G)$ may also be infinite, although certainly $\chi_f(G) \leq \chi(G)$. Interestingly, no analogue of the above-mentioned de Bruijn–Erdős result holds for fractional coloring [19, p. 74].

In 1992, Fisher and Ullman [8] first investigated the fractional chromatic number of the plane, denoted $\chi_f(\mathbb{R}^2)$. (Formally, we define a graph that has as its vertices the points of $\mathbb{R}^2$, where two vertices are adjacent if they are distance 1 in the plane. Now we consider the fractional chromatic number of this graph.) They observed that the fractional chromatic number of the Moser spindle is 3.5, which gives a lower bound on $\chi_f(\mathbb{R}^2)$. They also gave a coloring that proved the upper bound $\chi_f(\mathbb{R}^2) \leq 8\sqrt{3}/\pi \approx 4.4106$. Using a similar approach, Hochberg and O'Donnell [12] improved the upper bound to 4.3599. The construction they used was actually discovered much earlier by Croft [1]. The lower bound was first improved by Shawna Mahan [14], who found a unit distance graph with fractional chromatic number $144/41 \approx 3.5122$. This bound was significantly improved by Fisher and Ullman [19, p. 63–66], who found a unit distance graph with fractional chromatic number $\frac{32}{9} \approx 3.5556$. The bounds $3.5556 \leq \chi_f(\mathbb{R}^2) \leq 4.3599$ were the best known, until now. In this paper, we improve the lower bound to $76/21 \approx 3.6190$.

2. A First Lower Bound

In this section we provide a unit distance graph with fractional chromatic number greater than 3.6. Our construction builds heavily on an example of Fisher and Ullman, so we present that as well. To begin, we consider the fractional chromatic number of the two unit distance graphs we have already seen, the Moser spindle and the Golomb graph. The Moser spindle has 7 vertices and independence number 2, which show that $\chi_f \geq \frac{7}{2} = 3.5$. To prove that this lower bound holds with equality, it suffices to find 7 independent sets such that each vertex appears in two of them. This task is straightforward, once we put the bottom vertex into independent sets with each of its nonneighbors. The Golomb graph has 10 vertices and independence number 3, which show that $\chi_f \geq \frac{10}{3} = 3.\overline{3}$. In fact, this bound also holds with equality. The matching upper bound comes from finding 10 independent sets, such that each vertex appears in 3 of them. We leave this as an easy exercise.

The Moser spindle shows that $\chi_f(\mathbb{R}^2) \geq 3.5$. The intuition behind the Fisher–Ullman construction is the following. Consider a unit distance graph that contains many copies of the Moser spindle, along with as many edges as possible between these copies of the spindle. These edges between the copies of the spindle, as well as vertices that appear in more than one copy, ensure that some copy of the spindle must be colored suboptimally. This will prove some lower bound greater than 3.5. The details forthwith.

Recall that for any assignment of weights to the vertices, the fractional chromatic number is bounded below by the total weight on the vertices divided by the maximum total weight of any independent set. (We used this argument above to bound χ_f for the Moser spindle and Golomb graph; there we implicitly gave each vertex weight 1.) Thus, for any weight assignment, to bound χ_f from below, we need only bound from above the maximum weight of any independent set.

We construct the Fisher–Ullman graph in two stages. We begin with the subset of the triangular lattice shown in Figure 3(A), called the *core*; for now ignore the weights, which we will get to shortly. A *diamond* is the subgraph induced by two vertices at distance $\sqrt{3}$ and their two common neighbors. For each of the 5 vertical diamonds, we attach a copy of the Moser spindle; in each case, we identify the four vertices of the diamond with the four

(A) The core vertices and weights from the Fisher–Ullman construction; spindle weight 1 gives $\chi_f \geq \frac{32}{9} \approx 3.5555$.

(B) Vertices and weights for a bigger core; spindle weight 1 gives $\chi_f \geq \frac{168}{47} \approx 3.5744$.
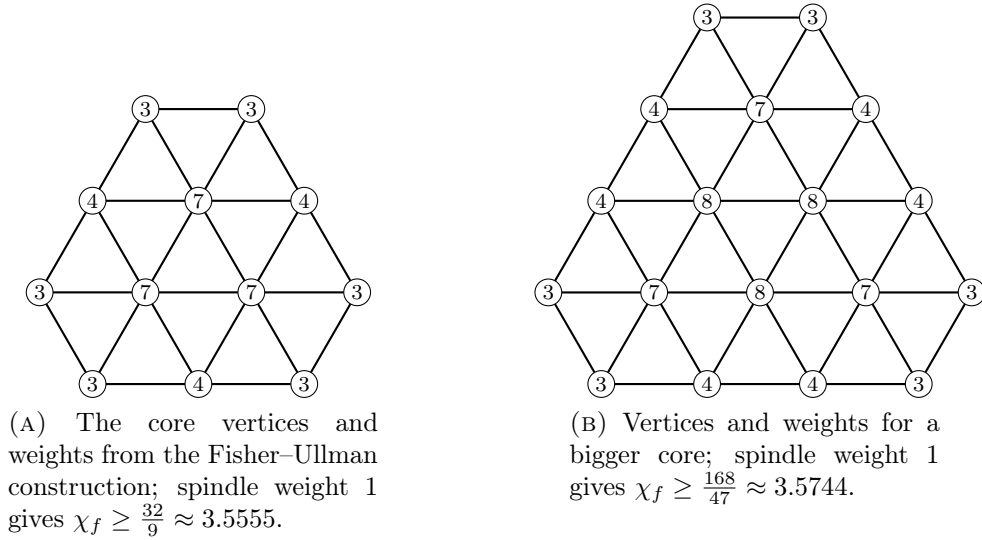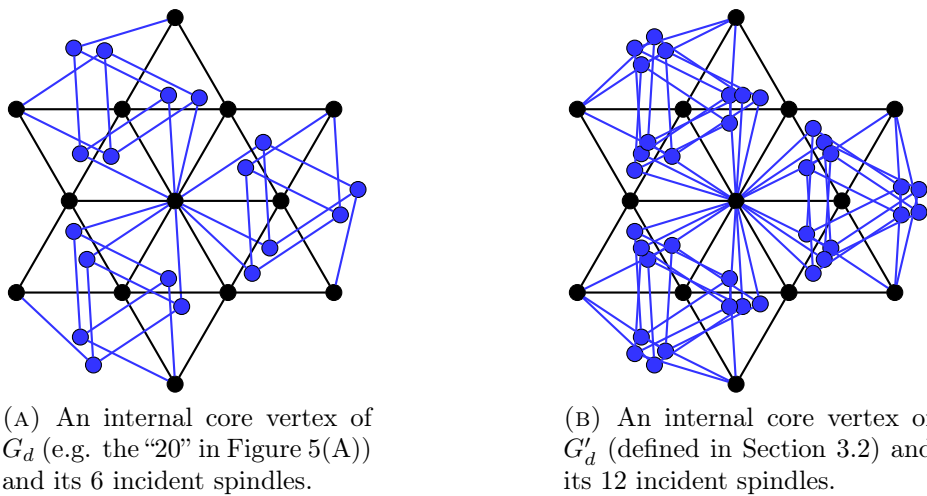
FIGURE 3. The two smallest cores that we consider.

vertices of the vertical diamond in the spindle, and we add three new vertices. These new vertices are *spindle vertices*. The two core vertices that are each adjacent to at least one of these spindle vertices are *incident to the spindle*. Nothing is special about the vertical direction in the core, so we also attach spindles to the five diamonds pointing down and to the left, as well as the five pointing down and to the right; before attaching these spindles, we rotate them 120 degrees clockwise and 120 degrees counterclockwise, respectively. Finally, we add all edges between pairs of vertices at distance 1.

The graph that results has 3-fold rotational symmetry. Each spindle adds 3 new vertices, so the 15 spindles add a total of 45 vertices. With the 12 core vertices, this makes a total of

(A) An internal core vertex of G_d (e.g. the "20" in Figure 5(A)) and its 6 incident spindles.

(B) An internal core vertex of G'_d (defined in Section 3.2) and its 12 incident spindles.

FIGURE 4. Spindles for G_d and G'_d.

(A) The core vertices and weights for a bigger core; spindle weight 2 gives $\chi_f \geq \frac{491}{137} \approx 3.5839$.

(B) The core vertices and weights for a bigger core; spindle weight 3 gives $\chi_f \geq \frac{1732}{481} \approx 3.6008$. This LP had over $25,000$ constraints.
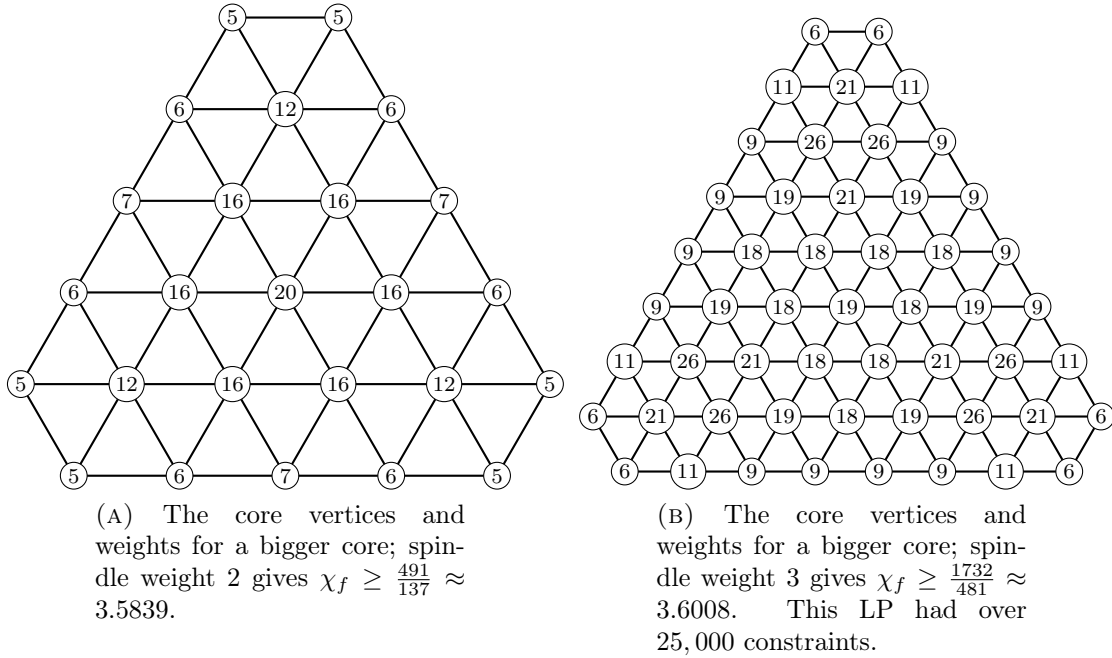
FIGURE 5. Two larger cores.

57 vertices. The graph has 24 edges among core vertices and 6(15) more edges within spindles; for each of three directions, it has 21 edges between vertices in different spindles that are oriented in the same direction. Finally, it has another 21 edges among pairs that seem to "accidentally" be at distance 1. (Understanding these last 21 edges is inessential, since they are not used in proving the lower bound.)

Now we assign weights to the vertices. The weights on the core vertices are shown in Figure 3(A). To each spindle vertex, we assign weight 1. We have 45 spindle vertices, and the weights on the core sum to 51, so the total weight is 96. Thus, to prove a lower bound of $\frac{32}{9}$, it suffices to show that every independent set has weight at most 27. This requires a short case analysis, and Scheinerman and Ullman [19, p. 64–65] gives most of the details.

Following this proof is fairly easy, but where do the weights come from? They come from solving a linear program (hereafter LP). Specifically, each weight is a variable, each independent set has weight at most 1, and the sum of the weights is to be maximized. (To simplify our presentation above, we multiplied all weights by 27, but that does not affect the lower bound.) So what about larger cores?

We can easily generalize the Fisher–Ullman construction to start from a larger core (as suggested in [19, p. 75]), and we will do exactly this. But first, it is helpful to comment on the obvious symmetry of the weights in Figure 3(A). Suppose we are given an optimal assignment of weights, i.e., an optimal solution to the LP in the previous paragraph. Each automorphism of the graph yields another optimal assignment of weights. Further, the average of all these weight assignments is again optimal. Thus, we may assume that the same weight is given to

all vertices in each orbit when the vertices are acted on by the automorphism group. This observation [13, p. 7] dramatically reduces the size of our LP, thus making it tractable.

Figures 3(B) and 5(A) show the results from solving the corresponding LPs for the next few sizes of cores. In each case, the lower bound on χ_f improves, but more slowly. We set for ourselves the goal of proving a lower bound of 3.6, and using the weights in Figure 5(B) we achieved it, just barely. In fact, these weights prove $\chi_f \geq 3.6008$. We would have liked to consider still larger cores, but the size of the derived LP was growing exponentially (along with the number of maximal independent sets), and the LP for the core in Figure 5(B) involved already more than 25,000 constraints. For this lower bound of 3.6008, we offer no proof. However, the interested reader can download our code for generating the LP from: `https://github.com/landon/WebGraphs/blob/master/Analysis/SpindleAnalyzer.cs`. To do much better, and certainly to give a human-checkable proof, we needed a new approach.

3. An Improved Lower Bound

3.1. **A Lower Bound by Discharging.** In the previous section, we proved that $\chi_f(\mathbb{R}^2) > 3.6$. To do so, we chose a unit distance graph and gave its vertices weights summing to more than 3.6, such that the weights on every independent set summed to at most 1. However, this proof has three drawbacks. First, it is not practically human-checkable, since that graph has more than 25,000 maximal independent sets. So to verify that each independent set has weight at most 1, we have a lot of cases to check. Second, we have no reason to believe that this bound of 3.6008 is actually near the fractional chromatic number of the plane. It is simply the best bound we could prove before the number of maximal independent sets became unmanageable. Finally, and perphaps most disturbingly, the proof offers no real insight. We have simply found some weights "that work".

In this section, we prove a stronger lower bound. In the process, we address all of these concerns, as well. We take a similar approach to our previous proof, but with a few key differences. As before, we start with a subset of the triangular lattice, called the core. Now everywhere that we possibly can, we add on spindles, much like in the previous proof. The first main difference is that we don't worry much about optimizing the weights. We give the same weight to every vertex in the core, and the same weight to every spindle vertex. (We will optimize the ratio of these two weights, but that problem is much easier.)

The second main difference—really the key that allows the proof to work—is that when bounding the weight of an independent set, we don't consider the set all at once. Rather, for an arbitrary maximal independent set I, we partition the graph into subgraphs, and bound the fraction of the weight on each subgraph that is in I. Now the fraction of weight on the whole graph that is in I is no more than the maximum fraction on any of these subgraphs. By partitioning into subgraphs of bounded size, we avoid the combinatorial explosion we faced in the previous section, when each maximal independent set generated its own constraint in the LP. An important decision is how to choose these subgraphs, so that we have relatively few cases, but we also get a good bound on the fraction of the total weight in I. We return to this question later.

To illustrate our approach, we start with a short proof that $\chi_f(\mathbb{R}^2) \geq 3.5$. Obviously this bound is weak (recall that the Moser spindle has fractional chromatic number 3.5), but its proof elucidates many of the key features of our method.

We define a graph G_d and its weights as follows. To begin, specify an arbitrary vertex v in the triangular lattice. For our core, take all of the triangular lattice induced by vertices that are distance at most d from v. For our spindles, add every possible spindle in each of the three directions we used in the previous section. This is our graph G_d. Let C_d denote the subgraph of G_d induced by its core vertices. For our weight function, we assign weight 12 to every core vertex and weight 1 to every spindle vertex. Now, given an arbitrary independent set I, we will discharge all of its weight to core vertices, so that each core vertex has weight at most 6. We use the following two discharging rules.

(R1) Each core vertex in I gives weight 1 to each of its neighbors in the core.
(R2) Each spindle vertex in I splits it weight equally between the core vertices incident to its spindle that are *not* in I.

Note that if a spindle vertex is in I, then at least one of the core vertices incident to the spindle is not in I; hence $(R2)$ does, indeed, move all charge from spindle vertices to core vertices. Now clearly all the weight in I ends on vertices in the core. We must verify that each core vertex finishes with charge at most 6. In addittion to core vertices in I, we have four possibilites for a core vertex not in I; it can have 3, 2, 1, or 0 core neighbors in I. We now check these five cases.

Note that each core vertex has at most 6 incident spindles (exactly 6 if the core vertex is not too close to the "outside" of the core). So a naïve upper bound on the final charge at each core vertex v is 9, since v receives 1 from each of at most 3 core neighbors, and it receives at most 1 from each of 6 spindles. However, this bound can be improved. Suppose that v is a core vertex not in I, and let u be some core neighbor of v that is in I. Now u has two neighbors, say w_1 and w_2, that are each distance $\sqrt{3}$ from v. Note that v shares one spindle each with w_1 and w_2. The key observation is that $u \in I$ implies $w_1 \notin I$ and $w_2 \notin I$. Thus, v receives weight at most $1/2$ from each of the spindles it shares with w_1 and w_2. Repeatedly applying this insight leads to the following upper bounds on the final charge at each core vertex.

Core vertex in I: $12 - 6(1) + 0 = 6$.
Core vertex with three neighbors in I: $0 + 3(1) + 6(1/2) = 6$.
Core vertex with two neighbors in I: $0 + 2(1) + 4(1/2) + 2(1) = 6$.
Core vertex with one nbr in I: $0 + 1(1) + 2(1/2) + 4(1) = 6$.
Core vertex with zero neighbors in I: $0 + 0(1) + 0(1/2) + 6(1) = 6$.

So we have shown that each core vertex finishes with weight at most 6. To compute a lower bound on χ_f, we divide the total weight on G_d by (an upper bound on) the total weight of any independent set. To simplify the computations, we neglect the effect of vertices near the outside of the core, which have fewer than six incident spindles (this choice can be justified, since the number of "interior" core vertices is asymptotically greater than the number of "boundary" core vertices). Let M denote the number of core vertices. Each core vertex is incident to six spindles and each spindle is incident to two core vertices; hence, the number of spindles is $M(3 - o(1))$. Since each spindle has 3 spindle vertices, the number of spindle

vertices is $M(9 - o(1))$. Thus, the total weight on the graph is $M(12 + 9(1) - o(1))$. The final total weight on the core is at most $6M$. This proves a lower bound of $\frac{21}{6} = 3.5$.

It is reasonable to ask why we give the same weight to every core vertex, since this differs from our choice in the previous section. A simple answer is that "it works", but this is unenlightening. A somewhat better answer is that "by weighting each vertex identically, we dramatically reduce the number of cases we must consider." This is true, but still misses the real point. To understand more fully, we return to the weights used in Section 2. There we weighted two vertices identically whenever the graph had an automorphism mapping one to the other. In that context, our automorphisms consisted of reflections and rotations, and compositions of the two.

In the present context, since we consider G_d as d grows without bound, we are essentially considering the fractional chromatic number of an infinite graph. (Our choice to consider the graph in ever larger pieces, rather than all at once, is mainly a concession with the goal of simplifying the averaging argument.) If we are indeed considering an infinite graph, then we have an additional type of automorphism: translations. Now it is clear that we can map any core vertex to any other core vertex; so the choice to weight them equally is quite natural.

Once we understand the proof of this lower bound, it's reasonable to ask how we can strengthen it. One obvious choice is to change the weight we give to each core vertex. To make this approach work, we need to consider the core vertices in larger groups, rather than in isolation, as in the previous proof. Our idea is to partition the subgraph induced by the core vertices into smaller subgraphs, which we call *tiles*. As our tiles grow bigger, we gain more information about the neighborhoods of their vertices, which facilitates a better bound on their average final weight. However, as the tiles grow bigger, they also grow more numerous, and require more case analysis. As a compromise, we choose the tiles to be (essentially) as small as possible, subject to each corner of each tile being a vertex of I, the independent set.

A priori, the distance could be large between a core vertex in I and its nearest core vertex in I. However, we are interested only in choices of I that potentially have maximum weight. By requiring that each core vertex gets weight greater than its number of incident spindles, we can make the following simplification. We only need to consider choices of I that are maximal independent subsets of the core. Suppose instead that for some core vertex v some independent set I contains neither v nor any of its core neighbors. We form I' from I by adding v and removing all spindle neighbors of v in I. Since the weight on v is greater than its number of incident spindles, I' weighs more than I.

In the following lemma, we formalize this idea of choosing the tiles to be "as small as possible, subject to each corner of each tile being a vertex of I, the independent set." We also show that this process results in only 8 distinct tiles, up to rotation and reflection.

3.2. **Tiling Result.**

Lemma 1. *Let I denote a maximal independent subset in C_d. There exists a set $\mathcal{T}$ of 8 finite tiles (shown in Figure 6), independent of d and I, such that C_d can be tiled with tiles from $\mathcal{T}$ where each corner of each tile is a vertex of I and no vertex of I lies in the interior of any tile. In this tiling, each face of C_d is covered by exactly one or two tiles. (We do allow*
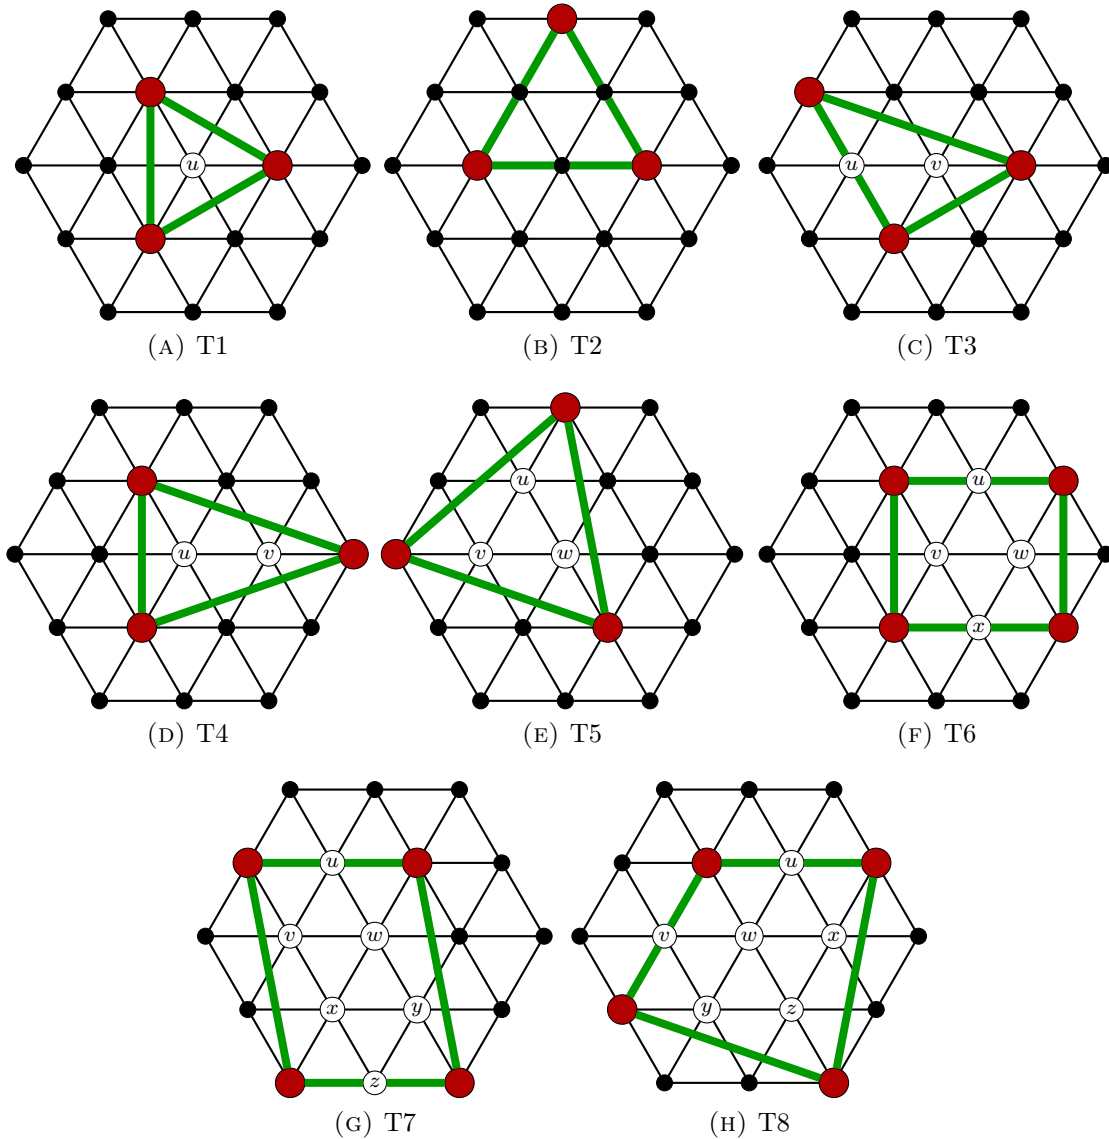
FIGURE 6. The 8 possible tiles (up to reflection and rotation).

tiles to extend past the boundary of G_d, though this allowance could be removed by adding more tiles to $\mathcal{T}$.)

Proof. For an example of such a tiling, see Figure 7.

Before we begin the formal proof, we note that every maximal subset I of C_d must contain at least $\frac{1}{7}$ of its vertices. Since I is maximal, for every vertex v, the set I contains either v or some neighbor of v. So to prove the lower bound $\frac{1}{7}$, we observe that there exists a set A of $\frac{1}{7}$ of all the vertices such that for every pair $u, w \in S$, we have $\text{dist}_{C_d}(u, w) \geq 3$. Thus, for each vertex u in A, the set I contains some element of $N[u]$; further, these sets are disjoint, i.e., for any pair $u, w \in A$, we have $N[u] \cap N[w] = \emptyset$. To see that such a set A exists, we view the
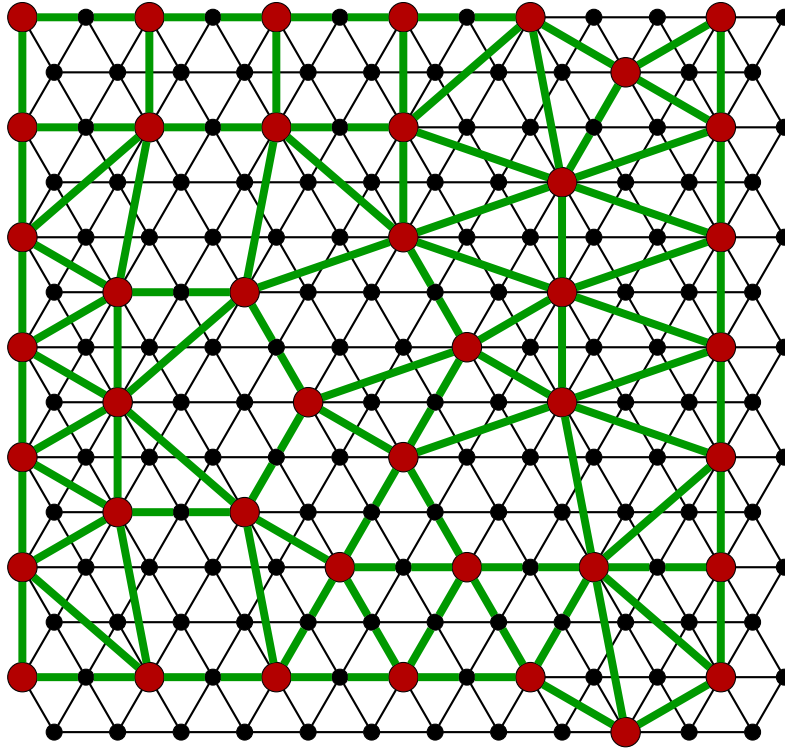
FIGURE 7. An example tiling of a portion of the triangular lattice.

infinite triangular lattice as the planar dual of the 7-face-coloring shown in Figure 2. This gives a 7-coloring of the triangular lattice, and hence a 7-coloring of C_d; we choose as A the vertices of a largest color class.

This reasoning actually proves that I must contain at least roughly $\frac{1}{7}$ of the vertices in any region of C_d. Intuitively, for each face f of C_d, some nearby vertices are in I, so it will be possible to cover f using a small tile, with all of its vertices nearby. Now we make this intuition rigorous.

First we describe the process for partitioning C_d into tiles; shortly, we will analyze it to show that we form only 8 tiles (up to rotation and reflection). To partition C_d into pieces, for each vertex pair $u, w \in I$, we add edge uw if and only if u and w have euclidean distance less than 3. Whenever two edges cross, we delete them both. This process clearly constructs a plane graph. In what follows, we show that the faces of this graph, which will become our tiles, have at most 8 distinct shapes and sizes. Let T be some arbitrary face of the plane graph constructed above. We first consider the case where T contains some edge of C_d in its interior (not boundary) that is incident to one of its corners; we will later show that if this is not the case, then T must be T2. By symmetry, we may assume that this corner is w, as shown in Figure 8. By rotational symmetry around w, we may assume that T contains edge ww_0 in its interior.

As noted above, I contains w_3 or one of its neighbors; similarly, I contains w_7 or one of its neighbors. Since ww_0 is in the interior of T, we may assume that $w_5 \notin I$. Since $w \in I$, none of its neighbors are in I. Thus, I contains some vertex in each of vertex sets A and
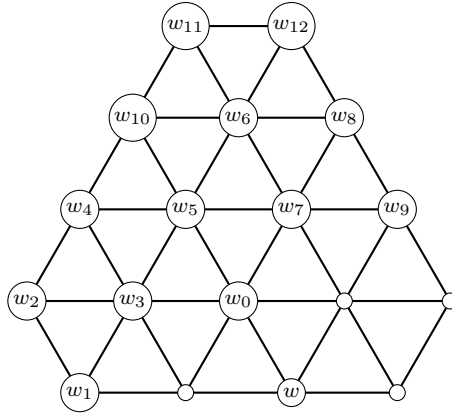
FIGURE 8. Names of key vertices for verifying that the tiles have only 8 shapes.

B, where $A = \{w_1, w_2, w_3, w_4\}$ and $B = \{w_6, w_7, w_8, w_9\}$. Naïvely, this gives 16 cases to consider. The number is actually smaller due to the symmetry across the line segment $\overline{ww_5}$, but we rarely make use of this symmetry, to keep the case organization clear. We consider these cases now.

Suppose that $w_3 \in I$. If $w_7 \in I$, then we have T1 (throughout the proof, we refer to each tile by its caption in Figure 6). If $w_8 \in I$, then we have T4. If exactly one of w_6 and w_9 is in I, then we have T3. If both of w_6 and w_9 are in I, then we have T6; note that this is the unique instance where initially two edges of our constructed graph crossed, so we deleted them.

Suppose that $w_4 \in I$. If $w_6 \in I$, then we have T4. If $w_7 \in I$, then we have T3. If $w_8 \in I$, then we have T5. So, suppose that $w_9 \in I$ and $w_6 \notin I$. Recall that I must contain w_6 or one of its neighbors. So either $w_{11} \in I$ or $w_{12} \in I$. In the former case, we have T7; in the latter, we have T8.

Suppose that $w_2 \in I$. If $w_6 \in I$, then we have T5. If $w_7 \in I$, then we have T4. If $w_8 \in I$, then we must have $w_{10} \in I$, since I contains neither w_5 nor any of its other neighbors. Now we have T8. So suppose that $w_9 \in I$ and $w_6 \notin I$. Again, we must have $w_{10} \in I$, since I contains neither w_5 nor any of its other neighbors. Now we have T7.

Suppose that $w_1 \in I$ (and $w_4 \notin I$). By symmetry across $\overline{ww_5}$, we may assume that $w_9 \in I$ and $w_6 \notin I$. Again, this implies $w_{10} \in I$. Once again, we have T8.

Now we consider the case that T contains no edge of C_d in its interior. Let w be a corner of T. We observe, as follows, that each boundary segment of T must have length at most 2. Suppose, to the contrary, that T has a boundary segment that is longer. By symmetry, we may assume that it is $\overline{ww_2}$. By symmetry, we may also assume that T lies above this segment (the case where T lies below it is essentially the same). We will show that T contains the edge w_2w_3 in its interior. Since I is maximal, it contains w_5 or one of its neighbors; specifically, I contains w_5, w_6, w_7, or w_{10}. In the first three cases, T is a triangle and clearly contains w_2w_3 in its interior. So assume that $w_{10} \in I$ and $w_7 \notin I$. Now either $w_8 \in I$ or $w_9 \in I$. So T is either T8 or T7, respectively; in each case T contains edge w_2w_3 in its interior. Hence, we conclude that each boundary segment of T has length at most 2.

Suppose now that T has a boundary segment of length $\sqrt{3}$. By symmetry, say this is segment $\overline{ww_3}$ and that T lies above this segment. Since I is maximal, it contains either w_7 or (at least) one of its neighbors. In each case, T contains edge ww_0 in its interior. Thus, we conclude that T has no boundary segments of length $\sqrt{3}$. Hence, every boundary segment of T must have length 2. Further, the interior angle formed by boundary segments at a corner must be $\pi/3$. So we conclude that T must be T2. $\square$

The previous result shows that we can tile C_d with only 8 tiles. However, each of these tiles may be rotated or reflected in numerous ways. In order to discharge the weight from spindle vertices to core vertices, we will view each of these tiles together with the spindles incident to its vertices. If we keep with our earlier strategy of attaching spindles at each vertex v in only 3 directions with v as their bottom vertex (for a total of 6 incident spindles, including those with v as their top vertex), we will break some of the symmetries we required in our proof that 8 tiles suffices. In particular, the roles of vertices incident to a spindle at its top and bottom are inherently asymmetrical, since the bottom is adjacent to two spindle vertices, while the top is adjacent to only one.

One obvious approach for handling this difficulty is to consider multiple cases for each tile, depending on its orientation; but this solution is inelegant. So we prefer instead the following approach; its key feature is that symmetry is preserved under both rotations and reflections, as well as compositions of the two. We now attach spindles at each vertex in 6 directions (for a total of 12 incident spindles) as follows; see Figure 4(B). Each pair of vertices, say u and v, that shared a spindle before now share two spindles; for one spindle u is the top vertex and v is the bottom, and for the other spindle the roles are reversed. For a pair of vertices u and v that share a spindle, the positions of the spindle vertices in the second spindle come from reflecting the vertices in the first spindle across the perpendicular bisector of line segment $\overline{uv}$. Thus, if the spindle with u as its base is rotated θ radians clockwise (around u) past $\overline{uv}$ (for an appropriate θ), then the spindle with v as its base is rotated θ radians *counterclockwise* (around v) past $\overline{uv}$. Hence, each vertex v will be the bottom vertex for 6 spindles. Three of these six spindles will each be oriented $2\pi/3$ radians clockwise of the previous one; the other three will be oriented similarly (relative to each other). We call this graph G_d'.

In Section 3.1, we used discharging to prove a lower bound of 3.5. To rephrase that proof in this setting with twice as many spindles, we simply give each core vertex weight 12 and each spindle vertex weight $\frac{1}{2}$ (rather than 1, as before). A moment's reflection will show that the rest of that proof goes through essentially unchanged.

An attentive reader will perhaps wonder whether all of these supposedly distinct spindle vertices do indeed fall in different locations. The answer is yes, although it turns out not to matter. First, we should mention that the angle of rotation for each spindle, called θ above, is $\cos^{-1}(\frac{5}{6})$. In less than a page of computations, we can show that the spindle vertices really do have distinct locations. But we don't need to.

Recall that our proof that $\chi_f(\mathbb{R}^2) \geq \frac{76}{21}$ consists of two parts. In one part, we prove lower bounds on the fractional chromatic number for a sequences of graphs, and show that these lower bounds converge to $\frac{76}{21}$. In the other part, we show that each graph in the sequence is a unit distance graph. We have constructed our graph sequence G_d', and in what follows, we

will prove the desired lower bounds. So we need only show that each G'_d is a unit distance graph, by describing some embedding.

We take as our embedding of the core vertices the obvious one, from the triangular lattice. If it happens that two (or more) spindle vertices coincide, then we assign to this "combined" vertex at their common location the *sum* of the weights we had intended to assign to them individually. Each of the supposedly distinct spindle vertices now has more neighbors than we originally claimed, but none has fewer. If this combined spindle vertex v does appear in the independent set I, then, for each spindle S containing v, we discharge to the core vertices incident to S the portion of v's weight that was due to it appearing in S. Thus, the lower bound remains valid.

3.3. **Discharging Result.** In this section, we continue with the graphs G'_d of the previous section. We show that as d grows, the fractional chromatic number of G'_d is bounded below by a sequence converging to $\frac{76}{21} \approx 3.619047$. In what follows, when we write $o(1)$, we mean as d goes to infinity.

Theorem 2. *The fractional chromatic number of the plane is at least $\frac{76}{21}$, i.e., $\chi_f(\mathbb{R}^2) \geq \frac{76}{21}$.*

Proof. At each core vertex we attach spindles in 6 directions, as described above. Each core vertex gets weight $\frac{31}{5}$, and each spindle vertex gets weight $\frac{1}{2}$. Let I be an arbitrary maximal independent set. We will redistribute the weight in I so that the core vertices end with average weight at most $\frac{21}{5}$ and each spindle vertex ends with weight at most 0. To help bound the average final weight of the core vertices, we compute the average weight of the core vertices in each tile, and show that the average weight for each tile is at most $\frac{21}{5}$. As before, we let M denote the number of core vertices. The total weight on the core is $\frac{31}{5}M$, and each spindle vertex has weight $\frac{1}{2}$. Since the number of spindle vertices is $M(18 - o(1))$, the total weight on G'_d is $M(\frac{31}{5} + \frac{1}{2}(18) - o(1))$. Since the core vertices end with average weight at most $\frac{21}{5}$, the total weight in I is at most $\frac{21}{5}M$. This proves that $\chi_f(\mathbb{R}^2) \geq (\frac{31}{5} + \frac{1}{2}(18) - o(1))/(\frac{21}{5}) = \frac{76}{21} \approx 3.619047$.

Now we give the details of how to redistribute the weight. Each vertex of C_d has a target weight of (at most) $\frac{21}{5}$. The *target* weight for a tile T is $\frac{21}{5}(i_T + \frac{1}{2}b_T + 0c_T)$, where i_T, b_T, and c_T denote the number of core vertices (respectively) in the interior of T, on the boundary (but not corners) of T, and on the corners of T. Each corner of T is a vertex of I, and its final weight will be computed by itself. Each interior vertex of T has all of its target weight assigned to T; similarly, each boundary (but not corner) vertex of T has exactly one half of its target weight assigned to T. If T has more weight than its target, then this difference is its *excess*; otherwise, the difference is its *deficit*. Our goal is to show that each tile finishes with excess at most 0.

Each diamond has a spindle attached in two directions; we call these the *up* spindle and the *down* spindle. If a diamond has two vertices in I, then its spindles are *trivial*. A nontrivial spindle that has no spindle vertex in I is *missing*. We redistribute the charge using three discharging phases.

Phase 1.

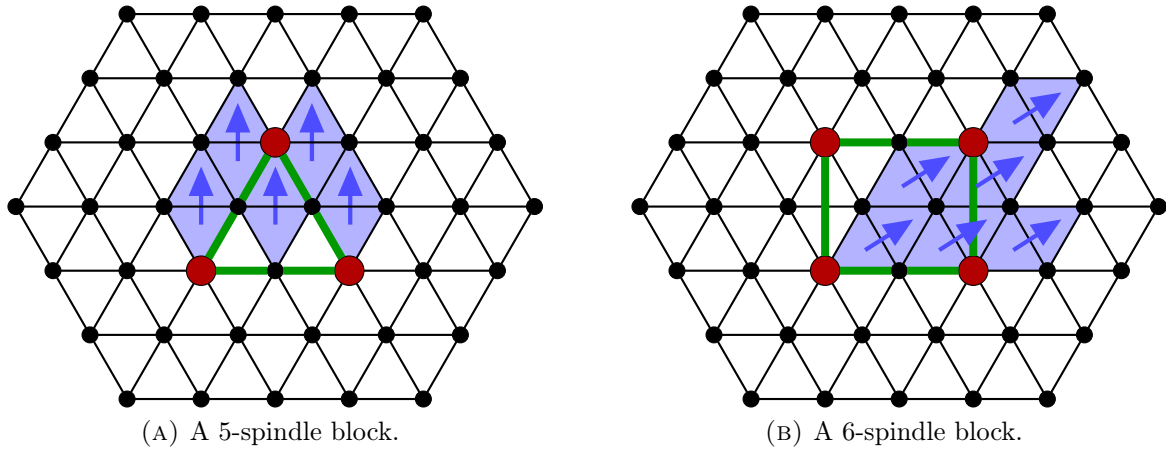(R1) Each core vertex in I gives weight $\frac{1}{3}$ to each core neighbor.

(A) A 5-spindle block.   (B) A 6-spindle block.

FIGURE 9. Spindle blocks used in Phase 3.

(R2) Each non-trivial spindle splits weight $\frac{1}{2}$ equally among the core vertices not in I that are incident to the spindle. If one of the two incident core vertices is in I, then the spindle sends weight $\frac{1}{2}$ to the other; otherwise it sends weight $\frac{1}{4}$ to each.

(R3) Each core vertex v not in I gives its weight to the tile containing it. If v is on the border of two tiles, then v splits its weight between them (but not quite equally). This is a little subtle, but crucial. For such a v, it has two neighbors in I; v does split the weight from its neighbors in I equally among its two tiles. Now v also has 12 incident spindles; so v sends any weight it got from 6 of the spindles to one tile and the weight from the other spindles to the other tile. Split the 12 spindles into two sets of 6 each, spaced 120 degrees apart (the up spindle and down spindle for each pair of vertices go into the same set). Now v gives each tile the weight it got from the set of spindles that includes the spindle pointing away from the border into that tile.

Phase 2.

(R4) Each copy of T3 takes $\frac{3}{10}$ from the tile on its long side.

(R5) Each copy of T1 splits its deficit equally among all adjacent copies of T6. In other words, it takes an equal amount of weight from each adjacent copy of T6, so that it ends with excess 0.

Phase 3.

(R6) Each copy of T2 gives $\frac{1}{4}$ to each missing spindle in its 5-spindle block (Figure 9(A)) in each of three directions. Each copy of T6 gives $\frac{1}{8}$ to each missing spindle in its 6-spindle block (Figure 9(B)) in each of four directions.

We now check that each vertex in I and each tile ends with average weight at most $\frac{21}{5}$. For vertices in I, this is immediate since each ends with weight $\frac{31}{5} - 6(\frac{1}{3}) = \frac{21}{5}$. We also check that each spindle finishes with weight at most 0.

Claim 1. *After Phase 1 the tiles have excess at most:*
T1: -1/5 T2: 7/10 T3: -3/10 T4: 3/5 T5: 2/5 T6: 2/5 T7: 0 T8: 0.

Proof. The main observation needed for this proof is essentially the same one we needed for the discharging proof in Section 3.1. If v is a core vertex not in I, then for each core neighbor u of v in I there exist four spindles incident to v that send v weight at most $\frac{1}{4}$ each. Specifically, u has two neighbors, say w_1 and w_2, that are each distance $\sqrt{3}$ from v. Further, v shares two spindles with each w_i. Since $u \in I$, we know that $w_1, w_2 \notin I$. Thus, each spindle shared between v and each w_i splits its weight equally between v and w_i. Hence, v receives weight at most $\frac{1}{4}$ from each such spindle.

T1: $3(\frac{1}{3}) + 12(\frac{1}{4}) = 4$. The target weight is $\frac{21}{5}$, so the excess is $4 - \frac{21}{5} = -\frac{1}{5}$.

T2: Since T2 has 3-fold rotational symmetry, we compute the weight received from one boundary vertex, then multiply by 3. This is $3(\frac{1}{3} + 4(\frac{1}{4}) + 2(\frac{1}{2})) = 7$. The target weight is $\frac{3}{2}(\frac{21}{5}) = \frac{63}{10}$, so the excess is $\frac{7}{10}$.

T3: We may assume that the tile $\widehat{T}$ bordering this copy T of T3 along its longest side is not another copy of T3 with each of its edges parallel to an edge of T. If it were, then these two copies of T3 would be merged into a single copy of T6. This assumption enables us to improve the bound on the total weight on T, which is $(\frac{1}{3} + 6(\frac{1}{4})) + (2(\frac{1}{3}) + 10(\frac{1}{4}) + 2(\frac{1}{2})) = 6$. The target weight is $\frac{3}{2}(\frac{21}{5}) = \frac{63}{10}$, so the excess is $-\frac{3}{10}$.

T4: $(2(\frac{1}{3}) + 12(\frac{1}{4})) + (\frac{1}{3} + 4(\frac{1}{4}) + 8(\frac{1}{2})) = 9$. The target weight is $2(\frac{21}{5}) = \frac{42}{5}$, so the excess is $\frac{3}{5}$.

T5: Since T5 has 3-fold rotational symmetry, we compute the weight received from one interior vertex, then multiply by 3. This is $3(\frac{1}{3} + 8(\frac{1}{4}) + 4(\frac{1}{2})) = 13$. The target weight is $3(\frac{21}{5}) = \frac{63}{5}$, so the excess is $\frac{2}{5}$.

T6: Since T6 has 2-fold rotational symmetry, we compute the weight received from vertices u and v, then multiply by 2. This is $2((\frac{1}{3} + 6(\frac{1}{4})) + (2(\frac{1}{3}) + 8(\frac{1}{4}) + 4(\frac{1}{2}))) = 13$. The target weight is $3(\frac{21}{5}) = \frac{63}{5}$, so the excess is $\frac{2}{5}$.

T7: Since T7 has 2-fold rotational symmetry, we compute the weight received from u, v, and w, then multiply by 2. This is $2((\frac{1}{3} + 6(\frac{1}{4})) + (\frac{1}{3} + 6(\frac{1}{4}) + 6(\frac{1}{2})) + (\frac{1}{3} + 10(\frac{1}{4}) + 2(\frac{1}{2}))) = 21$. The target weight is $5(\frac{21}{5}) = 21$, so the excess is 0.

T8: Since T8 has reflectional symmetry, we compute the weight received from w and z plus twice the weight received from u and x. This is $((\frac{1}{3} + 8(\frac{1}{4}) + 4(\frac{1}{2})) + (\frac{1}{3} + 12(\frac{1}{4}))) + 2((\frac{1}{3} + 6(\frac{1}{4})) + (\frac{1}{3} + 6(\frac{1}{4}) + 6(\frac{1}{2}))) = 21$. The target weight is $5(\frac{21}{5}) = 21$, so the excess is 0. $\square$

**Claim 2.** *After Phase 2 the tiles have excess at most:*
T1: 0 T2: 7/10 T3: 0 T4: 0 T5: 0 T6: 2/5 T7: 0 T8: 0.

Proof. Since Phase 2 only increases the weight for copies of T1 and T3, the bounds from Claim 1 remain valid for T2, T6, T7, and T8. By (R5), each copy of T1 ends with excess 0. Note that T3 has only a single long side. Since each copy ended Phase 1 with excess at most $-\frac{3}{10}$, it ends Phase 2 with excess at most 0.
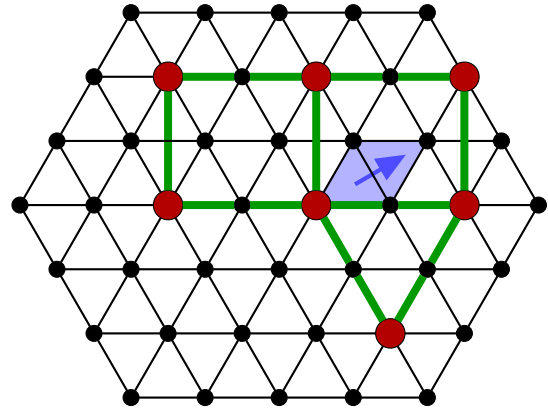
Consider a copy T of T4, and the tile $\widehat{T}$ that borders it along one of its two long sides. If $\widehat{T}$ is a copy of T3, then T4 gives away $\frac{3}{10}$. So assume instead that $\widehat{T}$ is a copy of T4, T5, T7, or T8. In this case, T saves at least $\frac{1}{2}$ over the bound we computed in Claim 1. Specifically,
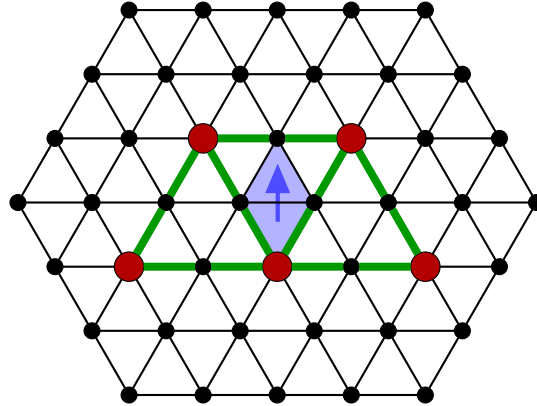
(A) Each spindle can appear in at
most two 5-spindle blocks.
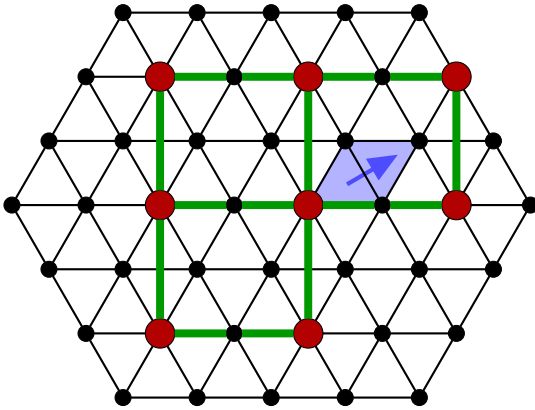
(B) Each spindle can appear in at
most three 6-spindle blocks.

(C) If a spindle appears in a 5-spindle
block, then it appears in at most two
6-spindle blocks.

FIGURE 10. Claim 3: Each missing spindle receives weight at most $\frac{1}{2}$.

an additional vertex at distance $\sqrt{3}$ from v is not in I. This means that v receives only one half of the weight from each spindle with these two vertices as its top and bottom (rather than all of it). So T saves $\frac{3}{10}$ if $\widehat{T}$ is a copy of T3, and saves $\frac{1}{2}$ otherwise. This savings applies to both of the long edges of T4, so T4 saves at least $2(\frac{3}{10}) = \frac{3}{5}$ over the bound given in Claim 1. Thus, each copy of T4 finishes with excess at most 0. A similar analysis holds for T5. Now T5 has three long edges, and saves at least $\frac{3}{10}$ along each of them, so finishes with excess at most $\frac{2}{5} - 3(\frac{3}{10}) = -\frac{1}{2}$. □

Claim 3. *Each spindle ends with nonnegative weight.*

Proof. If a spindle is trivial, then it begins and ends with weight 0. If a spindle is non-trivial and not missing, then it begins with weight $\frac{1}{2}$. By (R2), it ends with weight 0. So we need only consider missing spindles. By (R2), each such spindle finishes Phases 1 and 2 with

weight $-\frac{1}{2}$. So it suffices to show that in Phase 3 each missing spindle receives weight at most $\frac{1}{2}$. Recall that a missing spindle receives weight $\frac{1}{4}$ (from a copy of T2) for each 5-spindle block containing it and weight $\frac{1}{8}$ (from a copy of T6) for each 6-spindle block containing it. Thus, we must bound the number of 5-spindle and 6-spindle blocks containing each missing spindle.

First, if a missing spindle lies in two 5-spindle blocks, then it lies in no other 5-spindle or 6-spindle blocks; see Figure 10(A). So it receives weight at most $\frac{1}{4} + \frac{1}{4} = \frac{1}{2}$. Further, each missing spindle lies in at most three 6-spindle blocks. Such a spindle lies in no 5-spindle blocks; see Figure 10(B). So it receives weight at most $\frac{1}{8} + \frac{1}{8} + \frac{1}{8} \le \frac{1}{2}$. Finally, if a spindle lies in both a 5-spindle block and a 6-spindle block, then it lies in at most one 5-spindle block and at most two 6-spindle blocks; see Figure 10(C). So it receives weight at most $\frac{1}{8} + \frac{1}{8} + \frac{1}{4} = \frac{1}{2}$. $\qquad\square$

Claim 4. *In each 5-spindle block of a copy of T2, at least one spindle is missing.*

Proof. We number the (spindle) vertices of a spindle as 1, 2, 3, where vertices 1 and 2 are distance one from the bottom vertex, and vertex 3 is distance $\sqrt{3}$. We number the spindles in a 5-spindle block as $S_1, \ldots, S_5$, from left to right. Suppose that none of $S_1, \ldots, S_5$ is missing. Since the bottom vertices of S_1 and S_5 are in I, each of them must also have vertex 3 in I. In each of S_2, S_3, S_4 vertex 3 is distance 1 from vertex 3 of either S_1 or S_5. Since both of these vertices are in I, none of S_2, S_3, S_4 can have vertex 3 in I. However, among S_2, S_3, S_4, the three copies of vertex 1 are pairwise adjacent; similarly for the three copies of vertex 2. Hence, one of S_2, S_3, S_4 must be missing. $\qquad\square$

Claim 5. *For each copy of T6, the following are true.*
  (a) *For each 6-spindle block, either (i) at least one spindle is missing or (ii) at least two spindles are trivial.*
  (b) *If a copy of T6 has 2 trivial spindles on one side in the same direction, then that side is bordered by a copy of T1 that is adjacent to at most two copies of T6.*
  (c) *If a copy of T6 has 4 trivial spindles on one side, then that side is bordered by a copy of T1 that is adjacent to only one copy of T6.*

Proof. We prove each part in turn.
  (a) The proof is similar to that of Claim 4. We number the spindles as $S_1, \ldots, S_6$, from left to right (with the topmost spindle as S_5 and the bottom rightmost as S_6). We number the spindle vertices of each spindle 1, 2, 3, as in Claim 5. The only possible trivial spindles are S_5 and S_6; so suppose that at least one of them is neither trivial nor missing. Such a spindle must have vertex 3 in I. Whether this is S_5 or S_6, it forbids S_4 from having vertex 3 in I. Similarly, if S_1 is not missing, it has vertex 3 in I, which forbids both of spindles S_2 and S_3 from having vertex 3 in I. Now spindles S_2, S_3, S_4 all have vertex 3 forbidden from I. However, their copies of vertex 1 (resp. vertex 2) are pairwise adjacent. Hence, one of S_2, S_3, S_4 is missing.
  (b) Recall from (a) that the only possible trivial spindles in the 6-spindle block are S_5 and S_6. If S_6 is trivial, then the copy of T6 is bordered on its short side by a copy of T1. Further, if S_5 is also trivial, then that copy of T1 shares one of its sides with another copy of T1; hence, it can be bordered by at most two copies of T6.

(c) Each short side of a copy of T6 is crossed by two of its 6-spindle blocks. Recall that each 6-spindle block has only two possible trivial spindles. Thus, if a copy of T6 has 4 trivial spindles on one side, then both of its 6-spindle blocks crossing that side must have both possible trivial spindles. Now we apply (b) twice. So the copy of T6 is bordered along this side by a copy of T1, and that copy of T1 is bordered along its two other sides by other copies of T1. Hence, the copy of T1 is adjacent to only one copy of T6. □

Claim 6. *Each tile ends with excess at most 0.*

Proof. By Claim 2, we only need to check copies of T2 and T6. Further, at the end of Phase 2, each copy of T2 has excess at most $\frac{7}{10}$ and each copy of T6 has excess at most $\frac{2}{5}$. By Claim 4, each copy of T2 gives weight $\frac{1}{4}$ to a missing spindle in each of three directions. Thus, its final excess is at most $\frac{7}{10} - 3(\frac{1}{4}) = -\frac{1}{20}$.

If a copy of T6 has a missing spindle in all 4 directions, then it gives away to its missing spindles at least $4(\frac{1}{8}) = \frac{1}{2} \geq \frac{2}{5}$. If T6 has a missing spindle in 3 directions, then it gives away to its missing spindles at least $3(\frac{1}{8}) = \frac{3}{8}$. By Claim 5(a), in the direction where there is no missing spindle, T6 has 2 trivial spindles on one side in the same direction. By Claim 5(b), the copy of T6 is bordered by a copy of T1 that has at most 2 adjacent copies of T6. Hence by (R5), the copy of T6 gives at least $\frac{1}{2}(\frac{1}{5}) = \frac{1}{10}$ to the copy of T1. Thus the copy of T6 finishes with excess at most $\frac{2}{5} - \frac{3}{8} - \frac{1}{10} < 0$.

If a copy of T6 has a missing spindle in 2 directions, then it gives away to its missing spindles at least $\frac{2}{8}$. Using Claim 5(a) twice, we see that the copy of T6 has either 4 trivial spindles on one side or else 2 trivial spindles (in the same direction) on each side. In the first case, by Claim 5(c) and (R5), the copy of T6 gives $\frac{1}{5}$ to the copy of T1. In the second case, by Claim 5(b) and (R5), the copy of T6 gives away $\frac{1}{10} + \frac{1}{10}$. Hence the copy of T6 finishes with excess at most $\frac{2}{5} - \frac{1}{4} - \frac{1}{5} < 0$.

If a copy of T6 has a missing spindle in 1 direction, then it gives away to that missing spindle $\frac{1}{8}$. Also, the copy of T6 has 4 trivial spindles on one side and 2 on the other. Hence, from (R5) it gives away at least $\frac{1}{5} + \frac{1}{10}$ to adjacent copies of T1. So, it finishes with excess weight at most $\frac{2}{5} - \frac{1}{8} - \frac{1}{5} - \frac{1}{10} < 0$. Finally, suppose the copy of T6 has no missing spindles. Now by Claim 5(c), the copy of T6 is bordered on each side by a copy of T1 with only one T6 neighbor. Hence, by (R5), the copy of T6 gives away $\frac{1}{5} + \frac{1}{5}$, so finishes with excess at most 0. □

Claim 6 completes the proof that we outlined prior to the discharging phases. Hence, $\chi_f(\mathbb{R}^2) \geq \frac{76}{21}$. □

The lower bound $\frac{76}{21} \approx 3.6190$ can be improved slightly to $\frac{105}{29} \approx 3.6207$, as follows. We replace $\frac{1}{3}$ in (R1) with $\frac{13}{42}$ and replace $\frac{3}{10}$ in (R4) with $\frac{2}{7}$. A proof nearly identical to that above shows that every tile (and spindle) finishes with excess at most 0, except for possibly copies of T4. In order to accomodate copies of T4, we need to consider another type of 5-spindle block, shown in Figure 11. Although this proof can be completed, the details are surprisingly complicated (the proof requires an additional 4 phases), so we decided the reader would benefit more from the version presented here.
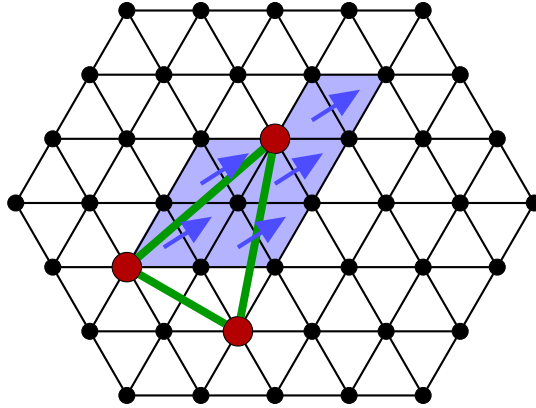
FIGURE 11. An additional 5 spindle block, used in the proof that $\chi_f(\mathbb{R}^2) \geq \frac{105}{29}$.

The bound $\frac{105}{29}$ seems to be the best possible using this method. The key obstruction to further improvements is that we need weight at least 6 on each core vertex. If a core vertex has less weight than the weight on an independent set among its spindle neighbors, then we cannot justify the assumption that I intersects the core vertices in a maximal independent set. This leads us to ask for the value of $\lim_{d\to\infty} \chi_f(G'_d)$. Specifically, is it larger than $\frac{105}{29}$?

4. Coloring Variants and Open Questions

In this section, we discuss a few variants on the problem of coloring (or fractionally coloring) the plane. Typically, we either put restrictions on each color class, or else we consider only a subset of the plane. We also include a few more open questions.

Falconer [6] proved that if we require color classes to be Lebesgue measurable, then any proper coloring of the plane requires at least 5 colors. The existence of non-measurable sets depends on the axioms of set theory we choose. In particular, when we drop the Axiom of Choice, there are (appealing) extensions of ZF in which all sets of reals are Lebesgue measurable [21]. So, basically, Falconer's result says that we can never hope to construct a 4-coloring of the plane. In the case of fractional coloring, Székely [22] proved in 1984 that a fractional coloring in which only measurable sets get non-zero weight must use total weight at least $\frac{43}{12} \approx 3.5833$. Recently, this was improved by Oliveira Filho and Vallentin [5] to ≈ 3.725.

Question. Do the fractional chromatic number of the plane and the measurable fractional chromatic number of the plane differ assuming ZFC?

The *j-fold chromatic number* $\chi_j(G)$ of a graph G is the least integer k so that it is possible to assign each vertex a set of j elements from $\{1, \ldots, k\}$ such that adjacent vertices get disjoint sets. Clearly, $\chi_f(G) \leq \frac{\chi_j(G)}{j}$, since we can assign each of the k color classes weight $\frac{1}{j}$. Recently, Grytczuk et.al. [10] studied j-fold coloring of the plane. Among other results, they generalized the construction of Hochberg and O'Donnell to give good j-fold colorings for small values of j.

Interesting results have been proved about chromatic numbers of extensions of $\mathbb{Q}$, the first being Woodall's result [23] that $\mathbb{Q} \times \mathbb{Q}$ is 2-colorable. Among other results, Fischer [7]

showed that $\chi\left(\mathbb{Q}(\sqrt{3}) \times \mathbb{Q}(\sqrt{3})\right) = 3$ and $\chi\left(\mathbb{Q}(\sqrt{11}) \times \mathbb{Q}(\sqrt{11})\right) \leq 4$. The graphs G_d' that we constructed have all vertices in $\mathbb{Q}(\sqrt{3}, \sqrt{11}) \times \mathbb{Q}(\sqrt{3}, \sqrt{11})$, so we have actually shown that $\chi_f\left(\mathbb{Q}(\sqrt{3}, \sqrt{11}) \times \mathbb{Q}(\sqrt{3}, \sqrt{11})\right) \geq \frac{76}{21}$. The natural extension of our construction is to attach spindles at all rotations that are integer multiples of $\cos^{-1}(\frac{5}{6})$, and this graph is still contained in $\mathbb{Q}(\sqrt{3}, \sqrt{11}) \times \mathbb{Q}(\sqrt{3}, \sqrt{11})$. Does this graph have larger fractional chromatic number? We do not know.

Question. What is the fractional chromatic number of $\mathbb{Q}(\sqrt{3}, \sqrt{11}) \times \mathbb{Q}(\sqrt{3}, \sqrt{11})$? What about its chromatic number?

Another intriguing direction of work is unit distance graphs with higher girth. Erdős [3] asked for which $k \geq 3$ there exist unit distance graphs of girth k with chromatic number 4. In his Ph.D. thesis [17], Paul O'Donnell showed that the answer is all $k \geq 3$. Mohar [15] extended this question to chromatic number 5 and 6. Much like the chromatic number (and fractional chromatic number) of the plane, this problem remains open.

References

[1] H. Croft. Incidence Incidents. *Eureka*, 30:22–26, 1967.

[2] N. G. de Bruijn and P. Erdös. A colour problem for infinite graphs and a problem in the theory of relations. *Nederl. Akad. Wetensch. Proc. Ser. A.* **54** = *Indagationes Math.*, 13:369–373, 1951.

[3] P. Erdős. Unsolved problems. *Congressus Numerantium XV–Proceedings of the 5th British Combinatorial Conference. 1975*, page 681, 1976.

[4] P. Erdős. Problems and results in combinatorial geometry. In *Discrete geometry and convexity (New York, 1982)*, volume 440 of *Ann. New York Acad. Sci.*, pages 1–11. New York Acad. Sci., New York, 1985.

[5] F. M. de Oliveira Filho and F. Vallentin. Fourier analysis, linear programming, and densities of distance avoiding sets in $\mathbb{R}^n$. *Journal of the European Mathematical Society*, 12(6):1417–1428, 2010.

[6] K. Falconer. The realization of distances in measurable subsets covering $\mathbb{R}^n$. *Journal of Combinatorial Theory, Series A*, 31(2):184–189, 1981.

[7] K. Fischer. Additive k-colorable extensions of the rational plane. *Discrete Mathematics*, 82(2):181–195, 1990.

[8] D. Fisher and D. Ullman. The fractional chromatic number of the plane. *Geombinatorics*, 2(1):8–12, 1992.

[9] M. Gardner. Mathematical games. *Scientific American*, 206:172–180, 10 1960.

[10] J. Grytczuk, K. Junosza-Szaniawski, J. Sokół, and K. Węsek. Fractional and *j*-fold colouring of the plane. *Bordeaux Graph Workshop 2014 Extended Abstracts*, pages 75–76, 2014.

[11] H. Hadwiger. Ueberdeckung des Euklidischen Raumes durch kongruente Mengen. *Portugaliae Math.*, 4:238–242, 1945.

[12] R. Hochberg and P. O'Donnell. A large independent set in the unit distance graph. *Geombinatorics*, 2(4):83–84, 1993.

[13] G. M. Levin. *Selected Topics in Fractional Graph Theory*. 1997. Thesis (Ph.D.)–The Johns Hopkins University.

[14] S. L. Mahan. The fractional chromatic number of the plane. Master's thesis, University of Colorado Denver, 1995.

[15] B. Mohar. *http://www.fmf.uni-lj.si/~mohar/Problems/P8UnitDistanceGraph.html*.

[16] L. Moser and W. Moser. Solution to problem 10. *Can. Math. Bull.*, 4:187–189, 1961.

[17] P. M. O'Donnell. *High-girth unit-distance graphs*. ProQuest LLC, Ann Arbor, MI, 1999. Thesis (Ph.D.)–Rutgers The State University of New Jersey - New Brunswick.

[18] D. Pritikin. All unit-distance graphs of order 6197 are 6-colorable. *Journal of Combinatorial Theory, Series B*, 73(2):159–163, 1998.

[19] E. R. Scheinerman and D. H. Ullman. *Fractional Graph Theory*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., New York, 1997. A rational approach to the theory of graphs, With a foreword by Claude Berge, A Wiley-Interscience Publication.

[20] A. Soifer. *The mathematical coloring book*. Springer, New York, 2009. Mathematics of coloring and the colorful life of its creators, With forewords by Branko Grünbaum, Peter D. Johnson, Jr. and Cecil Rousseau.

[21] R. Solovay. A model of set-theory in which every set of reals is Lebesgue measurable. *Annals of Mathematics*, pages 1–56, 1970.

[22] L. Székely. Measurable chromatic number of geometric graphs and sets without some distances in euclidean space. *Combinatorica*, 4(2):213–218, 1984.

[23] D. Woodall. Distances realized by sets covering the plane. *Journal of Combinatorial Theory, Series A*, 14(2):187–200, 1973.