Discrete Mathematics and Theoretical Computer Science DMTCS vol. **16:3**, 2014, 215–228

On the Number of Regular Edge Labelings[*]

Kevin Buchin[1] Bettina Speckmann[1] Sander Verdonschot[2]

[1]*Department of Mathematics and Computing Science, TU Eindhoven, Eindhoven, The Netherlands*
[2]*School of Computer Science, Carleton University, Ottawa, Canada*

received 27th Dec. 2011, revised 28th July 2014, accepted 8th Oct. 2014.

We prove that any irreducible triangulation on n vertices has $O(4.6807^n)$ regular edge labelings and that there are irreducible triangulations on n vertices with $\Omega(3.0426^n)$ regular edge labelings. Our upper bound relies on a novel application of Shearer's entropy lemma. As an example of the wider applicability of this technique, we also improve the upper bound on the number of 2-orientations of a quadrangulation to $O(1.87^n)$.

Keywords: regular edge labeling, counting, Shearer's entropy lemma

1 Introduction

An *irreducible triangulation* is a plane graph G such that (i) G is triangulated and the exterior face is a quadrangle, and (ii) G has no separating triangles (a 3-cycle with vertices both inside and outside the cycle). A *regular edge labeling* of an irreducible triangulation G is a partition of the interior edges of G into two subsets of red and blue directed edges such that: (i) around each inner vertex in clockwise order we have one contiguous non-empty set each of incoming blue edges, outgoing red edges, outgoing blue edges, and incoming red edges; (ii) the left exterior vertex has only outgoing blue edges, the top exterior vertex has only incoming red edges, the right exterior vertex has only incoming blue edges, and the bottom exterior vertex has only outgoing red edges (see Fig. 1, red edges are dashed). Regular edge labelings are also known as *transversal pairs of bipolar orientations* (Fusy, 2009).

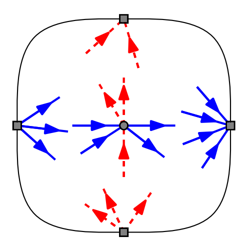

Fig. 1: The local conditions on a regular edge labeling.

Much of the importance of regular edge labelings stems from their connection to rectangular partitions. A *rectangular partition* or *floorplan* is a partition of a rectangle into rectangular faces such that no four rectangles meet at a common point (see Fig. 2, left). Rectangular partitions find applications in various areas: from floor plans of electronic chips (Koźmiński and Kinnen, 1985; Yeap and Sarrafzadeh, 1995) or architectural designs (Earl and March, 1979; Rinsma,

[*]A preliminary version of these results appeared in (Buchin et al., 2011) and (Verdonschot, 2010); those works have a strong focus on algorithmic applications of regular edge labelings, while we focus solely on the combinatorial bounds. K.B. and B.S. are supported by the Netherlands Organisation for Scientific Research (NWO) under project nos. 612.001.207 and 639.023.208, respectively. S.V. is partially supported by NSERC.

1988), to rectangular cartograms (Raisz, 1934; van Kreveld and Speckmann, 2007) and treemaps (Bruls et al., 2000).

Two rectangular partitions with n (interior) rectangles are *(strongly) equivalent* if the partitions contain (*i*) the same adjacencies between rectangles and (*ii*) these adjacencies have the same direction (vertical or horizontal). There has been considerable effort in counting the number $R(n)$ of equivalence classes of rectangular partitions with n rectangles. The currently known lower bound is $R(n) = \Omega(11.56^n)$ by Amano et al. (2007). The known upper bound is $R(n) \leqslant 13.5^{n-1}$ by Fujimaki et al. (2009). Inoue et al. (2009) computed $R(n)$ for $n \leqslant 3000$. These bounds consider all possible adjacencies between rectangles, but in practice, these adjacencies are often determined by the application (e.g., for rectangular cartograms they correspond to shared borders between regions), while the relative positions of regions can be still chosen. In such a setting we are interested in the number of equivalence classes of rectangular partitions with given adjacencies. In this article we give upper and lower bounds on this number.

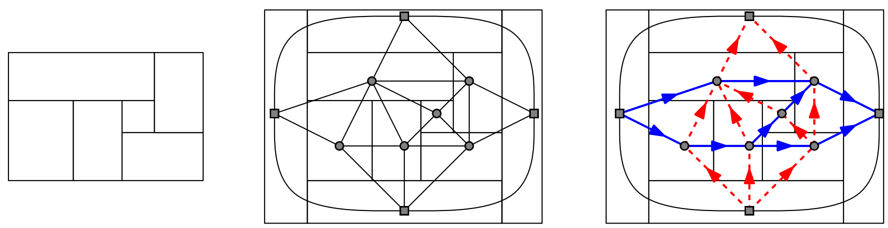

Fig. 2: A rectangular partition, its dual irreducible triangulation and the corresponding regular edge labeling.

The equivalence of rectangular partitions can be formulated in graph notion as follows. Given a rectangular partition, we augment it by adding four surrounding rectangles (see Fig. 2, middle). Two partitions now contain the same adjacencies between rectangles if and only if the dual graphs of the augmented partitions are the same. A plane graph can occur as dual if and only if it is an irreducible triangulation (Bhasker and Sahni, 1987; Koźmiński and Kinnen, 1985). A rectangular partition dual to an irreducible triangulation G is also referred to as a *rectangular dual* of G. A given irreducible triangulation might have many corresponding equivalence classes of rectangular partitions.

Now, given a rectangular partition, the directions of the adjacencies between rectangles can be represented as a labeling of the edges of the dual irreducible triangulation. In this labeling, all edges corresponding to horizontal (resp. vertical) adjacencies are colored blue (resp. red) and directed from left to right (resp. from bottom to top). The result is a regular edge labeling. The equivalence classes of the rectangular partitions with the same irreducible triangulation G as dual correspond one-to-one to the regular edge labelings of G.

Our motivation to bound the number of regular edge labelings of an irreducible triangulation stems from their application to drawing rectangular cartograms. To draw such a cartogram one needs to select one rectangular dual of a given irreducible triangulation, that is, one needs to select a regular edge labeling. The design of an algorithm that does so by exploring the space of regular edge labelings (Buchin et al., 2011, 2012) can be aided by a better understanding of this space and in particular its size.

In this article we prove the following bounds. Any irreducible triangulation G with n vertices has $O(4.6807^n)$ regular edge labelings (Section 2). Furthermore, there are irreducible triangulations with $\Omega(3.0426^n)$ regular edge labelings (Section 3).

Related work. Besides the notion of strong equivalence of rectangular partitions, there is also a notion of *weak equivalence*, for which two partitions are considered to be equivalent when the incidence structure between rectangles and maximal line segments is the same. The number of weak equivalence classes is in $\Theta(8^n/n^4)$ (Ackerman et al., 2006).

Instead of counting the total number of rectangular partitions, we are interested in the maximum number of rectangular partitions that all share the same dual graph. This is the same approach taken by Felsner and Zickfeld (2008), who count different ways of orienting the edges of planar graphs. The most general of these is called an α-orientation: given a planar graph $G = (V, E)$ and a function $\alpha : V \to \mathbb{N}$, an orientation of the edges of G is an α-orientation if every vertex v has out-degree $\alpha(v)$. They showed that a planar graph has at most $O(3.73^n)$ α-orientations, for any fixed function α. While regular edge labelings and α-orientations may not seem directly related, Fusy (2009) showed that there is a function α_0 such that the regular edge labelings of an irreducible triangulation G are in bijection with the α_0-orientations of the angular graph of G (the angular graph of G adds a vertex in every interior face of G and connects it to the three vertices of that face, then removes all original edges). However, because the angular graph has $3n - 6$ vertices, applying the general bound on α-orientations only gives a bound of $O(51.90^n)$ on the number of regular edge labelings – far from the bound we achieve.

A bipolar orientation of a connected graph G is an acyclic orientation of the edges of G such that there is exactly one vertex with no incoming edges (the source) and one with no outgoing edges (the sink). Felsner and Zickfeld (2008) showed that any connected planar graph has at most $O(3.97^n)$ bipolar orientations, while also showing that there are planar graphs with $\Omega(2.91^n)$ bipolar orientations. As a regular edge labeling consists of two disjoint bipolar orientations (one on the red edges and one on the blue edges), one might expect that the number of regular edge labelings is related to the number of bipolar orientations. Indeed, any regular edge labeling can be turned into a bipolar orientation by adding a source and a sink vertex and connecting the new source to the red and the blue sources and connecting the red and the blue sinks to the new sink. However, in this way many regular edge labelings can be mapped to the same bipolar orientation, as some regular edge labelings differ only in edge colors. Conversely, although Kant and He (1997) developed an algorithm that produces a regular edge labeling from the directions of the edges, not every bipolar orientation can be turned into a regular edge labeling this way. This is caused by the fact that bipolar orientations only require each (non-source and non-sink) vertex to have an in- and outdegree of at least one. Regular edge labelings on the other hand, require an in- and outdegree of at least two, one blue and one red.

Bipolar orientations are also related to α-orientations. Specifically, Rosenstiehl (1989) showed that bipolar orientations of a graph G are in bijection with 2-orientations (α-orientations where every vertex has out-degree 2) of the angular graph of G. As this angular graph is always a quadrangulation, Felsner and Zickfeld (2008) proved an upper bound of $O(1.91^n)$ and a lower bound of $\Omega(1.53^n)$ on the maximum number of 2-orientations a quadrangulation can have. Note that even though 2-orientations are in bijection with bipolar orientations, the bounds differ because the number of vertices differs.

Our upper bound relies on Shearer's entropy lemma (Chung et al., 1986), which was recently used by Björklund et al. (2008) to obtain $O\big((2 - \varepsilon)^n\big)$ time algorithms for the traveling salesperson problem. In contrast to our application of the lemma, they count vertex sets with certain properties and crucially rely on bounded maximum degree. As an example of the wider applicability of this technique, we also use it to slightly improve the upper bound on the number of 2-orientations of a quadrangulation to $O(1.87^n)$.

Many other interesting substructures have been counted in planar graphs. Aichholzer et al. (2007) list the known upper bounds for various subgraphs contained in a triangulation: Hamiltonian cycles,

cycles, perfect matchings, spanning trees, connected graphs and so on. Several of these bounds have been improved recently (Buchin et al., 2007; Buchin and Schulz, 2010). Some of the techniques used to count these substructures can also be used to count regular edge labelings. The bounds obtained this way (see end of Section 2 and end of Section 3), however, are far from the best bounds that we obtain. Also related to the problem of counting substructures in planar graphs is the problem of counting structures like triangulations on planar point sets (Dumitrescu et al., 2011; Hoffmann et al., 2013; Sharir and Sheffer, 2011).

Regular edge labelings are not only important because of their connection to rectangular partitions but also because of their connection to 4-connected plane graphs. Many plane graphs, in particular 4-connected plane graphs with at least four vertices on the exterior face, can be extended to an irreducible triangulation Biedl et al. (1997). Regular edge labelings can then be used to obtain straight-line drawings of these graphs on a small grid (Fusy, 2009).

2 Upper bound

We will use the following properties of regular edge labelings. A *regular edge coloring* is a regular edge labeling, with the directions of the edges omitted.

Lemma 1 (Fusy (2009), Proposition 2 and Lemma 1)

(a) A regular edge coloring uniquely determines a regular edge labeling.

(b) A regular edge labeling (of an irreducible triangulation) induces no monochromatic triangles.

Let $G = (V, E)$ be an irreducible triangulation on n vertices. By Lemma 1(a), a regular edge labeling is uniquely determined by the colors (red and blue) of the edges. Thus, the number of regular edge labelings of G is bounded by the number of edge colorings with two colors. Since G has less than $3n$ edges by Euler's formula, a simple upper bound is $2^{3n} = 8^n$. However, most of the colorings that we obtain by coloring the edges independently red or blue do not correspond to a valid regular edge labeling. In the following we refine our bound using Shearer's entropy lemma.

Lemma 2 (Shearer's entropy lemma (Chung et al., 1986)) *Let S be a finite set and let $A_1, \ldots, A_m$ be subsets of S such that every element of S is contained in at least k of the $A_1, \ldots, A_m$. Let $\mathcal{F}$ be a collection of subsets of S and let $\mathcal{F}_i = \{F \cap A_i : F \in \mathcal{F}\}$ for $1 \leqslant i \leqslant m$. Then we have*

$$|\mathcal{F}|^k \leqslant \prod_{i=1}^{m} |\mathcal{F}_i|.$$

Shearer's entropy lemma allows us to use the local conditions on regular edge colorings to bound the number of regular edge labelings.

Theorem 3 *The number of regular edge labelings of an irreducible triangulations is in $O(4.6807^n)$.*

Proof: Let $G = (V, E)$ be an irreducible triangulation on n vertices. Let S be E with the four edges on the exterior face excluded. For a regular edge labeling $\mathcal{L}$ of G let $E(\mathcal{L})$ be the set of blue edges in $\mathcal{L}$. Let

$$\mathcal{F} := \{E(\mathcal{L}) \mid \mathcal{L} \text{ is a regular edge labeling of } G\}.$$

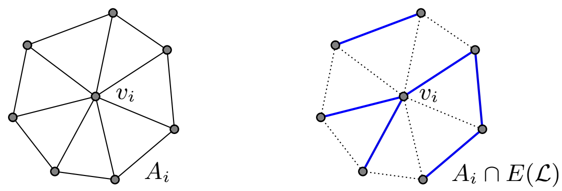

Fig. 3: A vertex v_i with the corresponding set of edges A_i, a locally consistent choice of blue edges, and an edge e with the four vertices that include e in their A_i.

Since $E(\mathcal{L})$ determines $\mathcal{L}$, the number of regular edge labelings is $|\mathcal{F}|$.

For the vertices v_i of G $(i = 1, 2, \ldots, n)$ let A_i be the set of edges belonging to S and to a triangle adjacent to v_i (see Fig. 3). Every edge $e \in S$ is in four of the sets A_i, namely in the four sets corresponding to the vertices of the two triangles with e as edge. Let $\mathcal{F}_i$ be the set of intersections of the set A_i with the sets $E(\mathcal{L})$, i.e., $\mathcal{F}_i$ contains all possible ways to choose blue edges around v_i consistent with a regular edge labeling. By Lemma 2 the number of regular edge labelings is bounded by $\prod_{i=1}^{n} |\mathcal{F}_i|^{1/4}$.

For a vertex v_i on the outer face there is only one way to choose the colors for the edges in A_i, since the adjacent edges must all have the same color. Since by Lemma 1(b) a regular edge labeling has no monochromatic triangle, the remaining edges in A_i need to be of the other color. Now let v_i be a vertex that is not on the outer face. We first bound the number of ways in which the edges adjacent to v_i can be colored. We color these edges starting with an arbitrary one and going clockwise from there. For the first edge, we have at most 2 choices and moving clockwise, we need to switch colors exactly four times by the local conditions on regular edge colorings. Therefore the number of choices for the edges adjacent to v_i is bounded by $2\binom{d_i}{4}$, where d_i is the degree of v_i. After coloring the adjacent edges, every triangle in A_i already has two colored edges. We have no choice for the third edge if these two edges have the same color, so we only have a choice for the four places where we switched colors.

Therefore, the number of regular edge labelings of G is bounded by

$$\prod_{i=1}^{n} \left(2^5 \binom{d_i}{4} \right)^{1/4} = \left(32^n \prod_{i=1}^{n} \binom{d_i}{4} \right)^{1/4} .$$

Jensen's inequality states that given a concave function f and a set of k values x_i in its domain, $\sum f(x_i) \leqslant kf(\sum x_i/k)$. Since the function $\log \binom{d}{4}$ is concave, this gives us $\sum_{i=1}^{n} \log \binom{d_i}{4} \leqslant n \log \binom{\sum_{i=1}^{n} d_i/n}{4}$. Since by Euler's formula the average degree $\sum_{i=1}^{n} d_i/n$ is bounded from above by 6, we get $\prod_{i=1}^{n} \binom{d_i}{4} \leqslant \binom{6}{4}^n = 15^n$. This yields the bound on the number of regular edge labelings of G of $480^{n/4} < 4.6807^n$. $\square$

2-orientations of quadrangulations

Next we improve the upper bound on the number of 2-orientations of quadrangulations. This demonstrates that the techniques of this section are more generally applicable. A quadrangulation is a plane graph where all faces are quadrangles (cycles of length four). A 2-orientation of a quadrangulation G is an orientation of the interior edges of G such that every interior vertex has out-degree 2, and the four vertices of the outer face have out-degree 0. By the bijection between bipolar orientations of planar graphs and 2-orientations of quadrangulations, counting the latter is the same as counting the former in the number of vertices and faces combined. This is the approach we take here.

Theorem 4 *The number of 2-orientations of a quadrangulation is in $O(1.87^n)$.*

Proof: Given a plane graph G with indicated source and sink, we intend to apply Shearer's entropy lemma with S being the edges of G. To define $\mathcal{F}$, we first define a total ordering on the vertices of G, for example by ordering them lexicographically by their coordinates in the plane. We can then uniquely identify an orientation of the edges of G by the subset of edges that are oriented from their lower ranked endpoint to their higher ranked endpoint. Let $\mathcal{F}$ be the collection of subsets of edges that correspond to a valid bipolar orientation on G.

As we are counting the number of bipolar orientations in terms of both faces and vertices, we consider two different neighbourhoods: all edges incident on a vertex and all edges around a face (see Figure 4). Since each edge is incident on two vertices and two faces, we have $k = 4$. In a bipolar orientation, the edges around each face form two directed paths, with one source and one sink. Thus, there are $f(d_f) = d_f(d_f - 1)$ options for each face, where d_f is the number of vertices on the face. Likewise, the edges around a vertex form two bundles, one with all incoming edges and the other with all outgoing edges. As the edges in each bundle are consecutive, the only choices are which vertex starts each bundle. This also gives $f(d_v)$ options per vertex, where d_v is the degree. As each edge is incident to two faces and two vertices, the total number of edge-face and edge-vertex incidences is four times the number of edges. Thus, the average degree d over all faces and vertices is 4. Using the concavity of $\log(f(d))$ and Jensen's inequality, we obtain

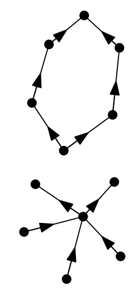

Fig. 4: The neighbourhoods around a face and vertex.

$$\prod_i f(d_i) \leqslant f(d)^n = 12^n.$$

Thus, by Shearer's entropy lemma, we get that the maximum number of 2-orientations of a quadrangulation is in $O(12^{n/4}) = O(1.87^n)$. $\square$

Probabilistic technique

In the remainder of this section we discuss an alternative derivation of an upper bound on the number of regular edge labelings. Although it is worse than the bound given in Theorem 3, we hope that the reasoning may contribute to improving the upper bound in future. If we call the average degree of the inner vertices d, the exterior vertices have total inner (so not counting the exterior edges) degree $(6 - d)n + 2$ by the Euler characteristic. The number of triangles adjacent to the exterior vertices is equal to this total degree. For each of these triangles, the coloring is already fixed. Again using the Euler characteristic, we can see that there are $(d - 4)n$ non-fixed triangles left. Each of these triangles must choose one of its corners to be monochromatic, since by Lemma 1(b) there can be no fully monochromatic triangles. Thus we obtain a bound of $3^{(d-4)n}$ labelings.

Not every assignment of monochromatic corners produces a valid regular edge labeling, as every inner vertex v must have exactly 4 dichromatic corners and $d_v - 4$ monochromatic corners, where d_v is the degree of v. If we now look at a random assignment of monochromatic corners, we can say something about the probability that all vertices will satisfy this condition. Since there are $4n$ dichromatic corners in total, asking that all inner vertices have at most four dichromatic corners is the same as asking that they have exactly four. Let X_i denote the event that vertex i has at most four dichromatic corners, and let I be a subset of the vertices. Then we have for any $j \in I$ that

$$P\Big(\bigcap_{i \in I} X_i\Big) \;=\; P\Big(\bigcap_{i \in I \setminus \{j\}} X_i \mid X_j\Big) \cdot P(X_j).$$

Next we argue that $P(\bigcap_{i \in I \setminus \{j\}} X_i \mid X_j) \leqslant P(\bigcap_{i \in I \setminus \{j\}} X_i)$. Intuitively, the probability on the left-hand side is smaller, because the condition X_j forces j to have "few" dichromatic corners, making it more likely for the neighbors of j to have "many" dichromatic corners. We prove the inequality by showing how to sample an assignment of colors to corners conditioned under X_j and by observing that in this process the probability of being dichromatic does not decrease for all corners not incident to j.

For all triangles not incident to j we can assign the monochromatic corner at random as before, since this is independent of X_j. Now assume that K corners are incident to j. We index these corners and their corresponding triangles (arbitrarily) by $1, \ldots, K$, and define Y_k as the random variable which is 1 if the kth corner of j is dichromatic, and 0 otherwise. Let $X_j(k, K')$ be the event that $Y_k + \ldots + Y_K \leqslant K'$. Note that $X_j = X_j(1, 4)$. We first assign a monochromatic corner to the first triangle at j conditioned under X_j. The condition X_j can only decrease the probability that $Y_1 = 1$, and therefore can only increase the probability for the other two corners of the triangle. We now continue with the assigning a monochromatic corner to the second triangle at j. If the first corner at j is dichromatic (resp. monochromatic), we now need to condition under X_j and $Y_1 = 1$ (resp. $Y_1 = 0$). This is equivalent to conditioning under $X_j(2, 3)$ (resp. $X_j(2, 4)$). Again this can only increase the probability of the corners of the second triangle that are not incident to j to be dichromatic. We proceed with the other triangles. Generally for triangle k we will condition under an event of the type $X_j(k, K')$ with $1 \leqslant K' \leqslant 4$, which then can only increase the probability of the corners not incident to j to be dichromatic. We conclude that the inequality above indeed holds. Applying it to the whole set of vertices as I and iteratively removing vertices from I, implies

$$P\big(\bigcap_i X_i\big) \leqslant \prod_i P\big(X_i\big).$$

Now let N_i be the number of dichromatic corners of vertex i. Then we can expand $P\big(X_i\big) = P\big(N_i \leqslant 4\big) = P\big(N_i = 0\big) + \cdots + P\big(N_i = 4\big)$. And since each triangle has exactly 2 dichromatic corners, N_i is a binomial random variable, with $P\big(N_i = x\big) = \binom{d_i}{x}(2/3)^x(1/3)^{d_i - x}$.

Finally, let $p(d_i) = P\big(X_i\big) = \sum_{x=0}^{4} \binom{d_i}{x}(2/3)^x(1/3)^{d_i - x}$. Since $\ln p(d)$ is concave for $d \geqslant 4$, we can use Jensen's inequality to show that $\prod_i P\big(X_i\big) \leqslant p(d)^n$. Thus the probability that any of those $3^{(d-4)n}$ labelings is valid is at most $p(d)^n$. With the average degree being at most 6, this gives us a bound of $p(6)^n \cdot 3^{(6-4)n} \approx 5.84^n$. Unfortunately, this is not even close to the bound of $O(4.6807^n)$ given in Theorem 3.

3 Lower bound

Next, we give lower bounds for the number of regular edge labelings in the triangulated square grid. We refer to the number of rows of a triangulated grid as its *height* h and to the number of columns as its *width* w. Each square is triangulated in the same way. To obtain an irreducible triangulation, we add four vertices to the outside of the grid, connecting one to every vertex on the top row, one to the bottom row, one to the left column, and one to the right column. The total number of vertices of the augmented grid is $n = hw + 4$.

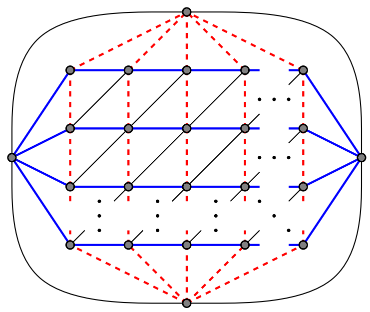

Fig. 5: The diagonals in the triangulated grid can be colored arbitrarily if all horizonal edges are blue, and all vertical edges red.

A simple lower bound on the number of regular edge labelings in a triangulated grid is $2^{n-h-w-3}$, which is in $2^{n-O(\sqrt{n})}$ for $h = w$. To see this, color all horizontal edges blue and all vertical edges red as shown in Fig. 5. Now all vertices already have the required four intervals of alternating red and blue edges and these intervals cannot be broken up by the diagonals, as these are all adjacent to intervals of both colors. Therefore all $n - h - w - 3$ diagonals can be colored independently blue or red, which yields the lower bound. This already shows that the number of regular edge labelings is exponential in n, but one would expect that there are many more regular edge labelings in a triangulated grid, because fixing the color of all horizontal and vertical edges, i.e., two-thirds of the edges, seems very restrictive.

We will show a better lower bound by only prescribing the edges of every h'th row to be blue, and not prescribing any colors for the edges of columns. (Nonetheless, the edges of the first and last column will be colored red in any regular edge labeling.) We color the parts between the blue rows independently. Therefore, we can assume for the moment that $h = h' + 1$. The reason for the introduction of h' will become clear later; suffice it to say that larger values of h (relative to h') do not change the analysis, but do improve the lower bound.

Our bounds require the analysis of large matrices, so part of the proof is by computer. We first describe all steps for $h' = 1$, i.e., the edges of all rows are blue. In this case, we can do all calculations by hand. Then we show how to generalize the method for larger values of h'.

We color the triangulated grid from left to right. The edges of the first and last column will need to be colored red, since by Lemma 1(b), a regular edge labeling has no monochromatic triangles. Assume we have colored the triangulated grid up to the ith column. We call the edges of the ith column and the diagonals connecting to this column from the left the ith *extended column*. Assuming we have no restriction from the right, our options for coloring the $(i + 1)$st extended column are determined by the colors of he ith extended column.

If $h' = 1$, the previous column can be either red or blue, while the color of the previous diagonal does not influence our choices for this column. If the previous column is red, we can make this column red, too, and choose either color for the diagonal. We can also make this column blue, but then the diagonal

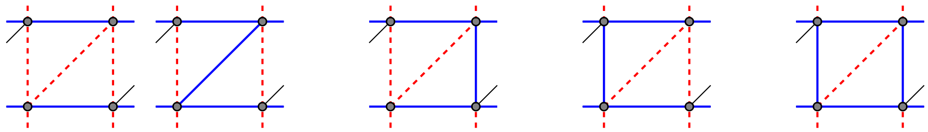

Fig. 6: The possible labelings of an extended column for the single row case: two options for red to red and one option each for red to blue, blue to red and blue to blue.

needs to be red to satisfy the constraints around the top vertex of this column. Likewise, if the previous column was blue, our diagonal needs to be red to satisfy the constraints around the bottom vertex of the previous column. These possibilities are depicted in Fig. 6.

We can represent these coloring options as a transition matrix $M = \left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)$, using the column colors as state. Now we can compute the number of colorings up to the ith extended column, by starting in the red state (represented as $(1, 0)$) and repeatedly multiplying it with M. The resulting vector gives us the number of colorings ending in a red or a blue edge.

To get from the transition matrix M to a lower bound we use the Perron–Frobenius Theorem, which we recall in the following. Let A be a non-negative $n \times n$ matrix. A matrix is *non-negative* if all its elements are non-negative. The matrix A is *irreducible* if for each (i, j) there is a $k > 0$ such that $(A^k)_{ij} > 0$. Consider the directed graph with adjacency matrix A, where we interpret every non-zero element as an adjacency. The matrix A is irreducible if and only if the associated graph is strongly connected. The matrix A is *primitive* if there is a $k > 0$ such that all elements of A^k are positive. An irreducible matrix with a positive diagonal entry is primitive. For more background on the following theorem we refer to textbooks on matrices (Horn and Johnson, 1985; Minc, 1988).

Theorem 5 (Perron–Frobenius) *Let A be a primitive non-negative matrix with maximal eigenvalue λ.*

(a) The eigenvalue λ is positive and it is the unique eigenvalue of largest absolute value. It has a positive eigenvector, and it is the only eigenvalue with non-negative eigenvector.

(b) Let $f_A(x) = \min_{x_i \neq 0} \frac{(Ax)_i}{x_i}$ and $g_A(x) = \max_{x_i \neq 0} \frac{(Ax)_i}{x_i}$. Then $f_A(x) \leqslant \lambda$ for all non-negative non-zero vectors x, and $g_A(x) \geqslant \lambda$ for all positive vectors x. If $f_A(x_0) = \lambda$ then x_0 is an eigenvector of A corresponding to λ.

(c) Let x be a non-negative non-zero vector. Then $A^t x / \|A^t x\|$ converges to an eigenvector with eigenvalue λ. Consequently,

$$\lim_{t \to \infty} f_A(A^t x) = \lim_{t \to \infty} g_A(A^t x) = \lambda.$$

To apply this theorem to the transition matrix M, let us first note that since M has only positive elements, it is primitive. By Theorem 5(c) the ratio between the number of labelings ending in a red column up to the ith extended column and up to the $(i + 1)$th extended column converges towards the largest eigenvalue of M. This eigenvalue is $\phi + 1 > 2.61803$, where ϕ is the golden ratio. So for any $\varepsilon > 0$, we obtain more than $(\phi + 1 - \varepsilon)^w$ labelings for sufficiently large w. Since we need to add two vertices to add a single column, this yields a lower bound of $(\phi + 1 - \varepsilon)^{(n-4)/2} > 1.61803^n$ for sufficiently large width w. However, by increasing h (and thereby glueing multiple copies of h' rows together) we only need to add h vertices to add one column, now with $h - 1$ rows. This essentially amortizes the additional vertex over all copies of h' rows. Thus we get a lower bound of $(\phi + 1 - \varepsilon)^{(n-4)(h-1)/h}$, which for sufficiently large h and w is larger than 2.61803^n.

Next we consider the case where $h' > 1$, using $h' = 2$ as an example. Since we prescribe the color of fewer edges, this will yield a better bound. The biggest change from $h' = 1$ is that we need to extend the states with information about the vertices, as just using the colors of the edges is not sufficient to decide how the coloring can be extended. Therefore, we include for each vertex the number of color switches that should be in the next extended column. With this information, we can reconstruct the colors of all column edges from a single colored edge, so we will describe the state by the color of the bottom column edge, followed by the number of color switches at each vertex, moving upwards. As an example, the states for $h' = 2$ are given in Fig. 7.

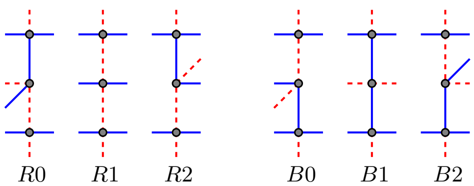

$$\begin{pmatrix} 1 & 4 & 1 & 1 & 2 & 5 \\ 0 & 5 & 2 & 2 & 1 & 4 \\ 0 & 2 & 2 & 1 & 1 & 1 \\ 0 & 2 & 2 & 2 & 2 & 1 \\ 0 & 1 & 2 & 1 & 2 & 2 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \qquad \begin{pmatrix} 5 & 2 & 2 & 1 \\ 2 & 2 & 1 & 1 \\ 2 & 2 & 2 & 2 \\ 1 & 2 & 1 & 2 \end{pmatrix}$$

Fig. 7: The states for $h' = 2$, the corresponding transition matrix and reduced transition matrix.

Some states that can be described in this way cannot in fact be part of a regular edge labeling. We call such states *infeasible*. A state is feasible if it can be reached from the initial all-red state ($R1$ in the case of $h' = 2$) and if the all-red state can be reached from it. Thus a state is feasible if and only if it is in the strongly connected component of the all-red state. For example, looking at the transition matrix for $h' = 2$ given in Fig. 7, we can see that the first state, $R0$, doesn't have any incoming transitions from other states and the last state, $B2$, doesn't have any outgoing transitions to other states. Therefore these states are both infeasible and the matrix can be reduced to include only the feasible states. In our implementation[(i)], we use two depth-first searches through the adjacency graph corresponding to the transition matrix and starting from the all-red state to determine the feasible states. The first search traverses the edges in the usual way to mark all states that are reachable from the all-red state, while the second search traverses each edge backwards, to mark all states from which the all-red state is reachable. The feasible states are exactly those states marked by both searches.

The resulting reduced matrix is primitive, since it is irreducible by construction, and there is always at least one transition from the all-red state to itself by coloring all diagonals and horizontal edges between the two columns blue. When constructing a regular edge coloring, we start with the all-red column and color the columns one by one. By Theorem 5(c) the number of regular edge colorings (and therefore the number of regular edge labelings) increases with each new column by a factor that converges to the largest eigenvalue of the reduced transition matrix. Therefore, a strict lower bound on this eigenvalue $\lambda_{h'}$ of this matrix gives us a strict lower bound for the growth rate per column (ignoring a constant number of initial columns).

We obtain this strict lower bound on $\lambda_{h'}$ by taking a non-negative non-zero state vector x, multiplying it with the transition matrix and determining the minimum growth rate for the non-zero elements. If

[(i)] Our code for generating the transition matrices and estimating the eigenvalues can be found at http://www.win.tue.nl/~speckman/demos/TransitionMatrixComputerRelease.zip the vector is not an eigenvector of A (i.e., the growth rate is not the same for all non-zero states) then the minimum growth rate is a strict lower bound on $\lambda_{h'}$ by Theorem 5(b). Tab. 1 gives minimum and maximum growth rates for $h' = 2$, where A is the reduced transition matrix and x_0 is the vector with a 1 for the all-red state and 0 otherwise. It shows that the growth rates converge quite rapidly to $\lambda_{h'}$. We use $x = x_0 A^{100}$ for our lower bounds. Since in all the cases that we consider the vector x is positive, we also obtain an upper bound on $\lambda_{h'}$ by the maximum growth factor.

t	1	2	3	4	100
$f_A(x_0 A^t)$	6.8	7.56	7.80	7.87	7.89167
$g_A(x_0 A^t)$	13	8.77	8.10	7.94	7.89167

Tab. 1: Minimum and maximum growth rates for $h' = 2$. The values for the smaller values of t are rounded to 2 decimal places, while the values for $t = 100$ are rounded to 5 decimal places, although they only start to deviate at the 58th decimal place.

As for the case $h' = 1$, we now use several copies of h' rows beneath each other to obtain a larger triangulated grid. The growth rate per vertex in this way approaches $\lambda_{h'}^{1/h'}$. Our results are given in Tab. 2.

h'	1	2	3	4	5	6	7
$\lambda_{h'}$	2.61803	7.89167	24.5036	76.8353	241.977	763.785	2414.05
$\lambda_{h'}^{1/h'}$	2.61803	2.80921	2.90453	2.96067	2.99746	3.0233	3.04263

Tab. 2: Lower bounds on the growth rate per column and per vertex for different values of h'. Note that these are rounded down, and that our upper bounds on $\lambda_{h'}$ (and $\lambda_{h'}^{1/h'}$) equal the (unrounded) lower bounds up to at least 10 significant digits.

Theorem 6 *The number of regular edge labelings of the triangulated grid is in* $\Omega(3.04263^n)$.

4 Discussion

The technique used by Felsner and Zickfeld (2008) to bound the number of bipolar orientations crucially relies on an encoding of bipolar orientations that uses one bit of information per face of the graph. If such an encoding were found for regular edge labelings, this would immediately improve the upper bound to $O(4^n)$.

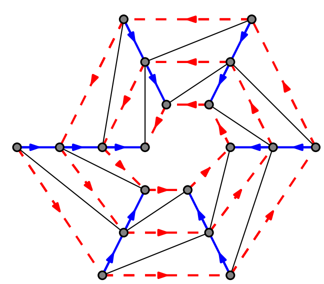

Fig. 8: A planar drawing of the twisted cylinder with width 6.

We would like to end with a discussion of why a *twisted cylinder* is not suitable to derive an interesting lower bound. A twisted cylinder is sometimes a good alternative for counting structures that are numerous on the triangulated grid. For example, it has been used to count simple and Hamiltonian cycles on planar graphs by Buchin et al. (2007). Imagine taking a piece of squared paper and bending it to make the ends meet and form a cylinder. Now instead of lining up the rows with each other, shift everything one square to the right. This produces a single line of squares, twisting itself around the cylinder. The resulting graph can be drawn in the plane without crossings, as if you were looking through the cylinder (depicted to the right). The advantage over the triangulated grid is that now cells can be added one at a time, leading to a far less complicated transition matrix than if you would add an entire column at a time.

Unfortunately, the twisted cylinder has only few regular edge labelings. Although it resembles the triangulated grid in many ways, the transformation used when drawing the graph in the plane causes the graph to have a very limited number of regular edge labelings. This is easiest to see when looking at the regular edge labeling for the first simple lower bound we gave using the triangulated grid, where every vertical edge was colored red and oriented from bottom to top and every horizontal edge was colored blue and oriented from left to right. Applying this labeling to the twisted cylinder makes both the red and blue edges move inwards. This can never lead to a valid regular edge labeling, as this requires that the blue edges form a bipolar orientation with the west vertex as source and the east vertex as only sink, both on the outside of the spiral.

Acknowledgements

We are grateful to the anonymous referees for their detailed comments.

References

E. Ackerman, G. Barequet, and R. Y. Pinter. A bijection between permutations and floorplans, and its applications. *Discrete Applied Mathematics*, 154(12):1674–1684, 2006.

O. Aichholzer, T. Hackl, C. Huemer, F. Hurtado, H. Krasser, and B. Vogtenhuber. On the number of plane geometric graphs. *Graphs and Combinatorics*, 23:67–84, 2007.

K. Amano, S. Nakano, and K. Yamanaka. On the number of rectangular drawings: Exact counting and lower and upper bounds. Technical Report 2007-AL-115, IPSJ SIG, 2007.

J. Bhasker and S. Sahni. A linear algorithm to check for the existence of a rectangular dual of a planar triangulated graph. *Networks*, 7:307–317, 1987.

T. Biedl, G. Kant, and M. Kaufmann. On triangulating planar graphs under the four-connectivity constraint. *Algorithmica*, 19(4):427–446, 1997.

A. Björklund, T. Husfeldt, P. Kaski, and M. Koivisto. The travelling salesman problem in bounded degree graphs. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP), Part I*, volume 4051 of *Lecture Notes in Computer Science*, pages 198–209. Springer, 2008.

M. Bruls, K. Huizing, and J. J. van Wijk. Squarified treemaps. In *Data Visualization 2000: Proceedings Joint Eurographics and IEEE TCVG Symp. on Visualization*, pages 33–42. Springer, 2000.

K. Buchin and A. Schulz. On the number of spanning trees a planar graph can have. In *Proceedings of the 18th European Symposium on Algorithms (ESA)*, volume 6346 of *Lecture Notes in Computer Science*, pages 110–121. Springer, 2010.

K. Buchin, C. Knauer, K. Kriegel, A. Schulz, and R. Seidel. On the number of cycles in planar graphs. In *Proceedings of the 13th Computing and Combinatorics Conference (COCOON)*, volume 4598 of *Lecture Notes in Computer Science*, pages 97–107. Springer, 2007.

K. Buchin, B. Speckmann, and S. Verdonschot. Optimizing regular edge labelings. In *Proceedings of the 18th International Symposium on Graph Drawing (GD)*, volume 6502 of *Lecture Notes in Computer Science*, pages 117–128. Springer, 2011.

K. Buchin, B. Speckmann, and S. Verdonschot. Evolution strategies for optimizing rectangular cartograms. In *Proceedings of the 7th International Conference on Geographic Information Science*, volume 7478 of *Lecture Notes in Computer Science*, pages 29–42. Springer, 2012.

F. R. K. Chung, R. L. Graham, P. Frankl, and J. B. Shearer. Some intersection theorems for ordered sets and graphs. *Journal of Combinatorial Theory, Series A*, 43(1):23–37, 1986.

A. Dumitrescu, A. Schulz, A. Sheffer, and C. D. Tóth. Bounds on the maximum multiplicity of some common geometric graphs. In *Proceedings of the 28th International Symposium on Theoretical Aspects of Computer Science*, volume 9 of *LIPIcs*, pages 637–648. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2011.

C. F. Earl and L. J. March. Architectural applications of graph theory. In R. J. Wilson and L. W. Beineke, editors, *Applications of Graph Theory*, pages 327–355. Academic Press, 1979.

S. Felsner and F. Zickfeld. On the number of planar orientations with prescribed degrees. *Electronic Journal of Combinatorics*, 15(1), 2008. Research paper R77, 41 p.

R. Fujimaki, Y. Inoue, and T. Takahashi. An asymptotic estimate of the numbers of rectangular drawings or floorplans. In *Proceedings of the International Symposium on Circuits and Systems (ISCAS)*, pages 856–859, 2009.

É. Fusy. Transversal structures on triangulations: A combinatorial study and straight-line drawings. *Discrete Mathematics*, 309(7):1870–1894, 2009.

M. Hoffmann, A. Schulz, M. Sharir, A. Sheffer, C. D. Tóth, and E. Welzl. Counting plane graphs: flippability and its applications. In J. Pach, editor, *Thirty Essays on Geometric Graph Theory*, pages 303–326. Springer, 2013.

R. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, UK, 1985.

Y. Inoue, T. Takahashi, and R. Fujimaki. Counting rectangular drawings or floorplans in polynomial time. *IEICE Transactions*, 92-A(4):1115–1120, 2009.

G. Kant and X. He. Regular edge labeling of 4-connected plane graphs and its applications in graph drawing problems. *Theoretical Computer Science*, 172(1–2):175–193, 1997.

K. Koźmiński and E. Kinnen. Rectangular dual of planar graphs. *Networks*, 5:145–157, 1985.

H. Minc. *Nonnegative Matrices*. Wiley-Interscience, Hoboken, NJ, USA, 1988.

E. Raisz. The rectangular statistical cartogram. *Geographical Review*, 24(2):292–296, 1934.

I. Rinsma. Rectangular and orthogonal floorplans with required rooms areas and tree adjacency. *Environment and Planning B: Planning and Design*, 15:111–118, 1988.

P. Rosenstiehl. Embedding in the plane with orientation constraints: The angle graph. *Annals of the New York Academy of Sciences*, 555(1):340–346, 1989. ISSN 1749-6632.

M. Sharir and A. Sheffer. Counting triangulations of planar point sets. *Electronic Journal of Combinatorics*, 18(1), 2011.

M. van Kreveld and B. Speckmann. On rectangular cartograms. *Computational Geometry: Theory and Applications*, 37(3):175–187, 2007.

S. Verdonschot. Optimizing regular edge labelings. Master's thesis, Eindhoven University of Technology, 2010.

G. K. H. Yeap and M. Sarrafzadeh. Sliceable floorplanning by graph dualization. *SIAM Journal on Discrete Mathematics*, 8(2):258–280, 1995.