

Optimizing Regular Edge Labelings^{*}

Kevin Buchin¹, Bettina Speckmann¹, and Sander Verdonschot²

¹ Dep. of Mathematics and Computer Science, TU Eindhoven, The Netherlands

k.a.buchin@tue.nl, speckman@win.tue.nl

² School of Computer Science, Carleton University, Canada

sverdons@connect.carleton.ca

Abstract. A regular edge labeling (REL) of an irreducible triangulation G uniquely defines a rectangular dual of G . Rectangular duals find applications in various areas: as floor plans of electronic chips, in architectural designs, as rectangular cartograms, or as treemaps. An irreducible triangulation can have many RELs and hence many rectangular duals. Depending on the specific application different duals might be desirable. In this paper we consider optimization problems on RELs and show how to find optimal or near-optimal RELs for various quality criteria. Furthermore, we give upper and lower bounds on the number of RELs.

1 Introduction

A *rectangular partition* of a rectangle R is a partition of R into a set \mathcal{R} of non-overlapping rectangles such that no four rectangles in \mathcal{R} meet at one common point. A *rectangular dual* of a plane graph G is a rectangular partition \mathcal{R} , such that (i) there is a one-to-one correspondence between the rectangles in \mathcal{R} and the nodes in G ; (ii) two rectangles in \mathcal{R} share a common boundary if and only if the corresponding nodes in G are connected. Rectangular duals find applications in various areas: as floor plans of electronic chips or in architectural designs, as rectangular cartograms, or as treemaps.

Not every plane graph has a rectangular dual. A plane graph G has a rectangular dual \mathcal{R} with four rectangles on the boundary of \mathcal{R} if G is an *irreducible triangulation*: (i) G is triangulated and the exterior face is a quadrangle; (ii) G has no separating triangles (a 3-cycle with vertices both inside and outside the cycle) [6,19]. A plane triangulated graph G has a rectangular dual if and only if we can augment G with four external vertices such that the augmented graph is an irreducible triangulation.

The equivalence classes of the rectangular duals of an irreducible triangulation G correspond one-to-one to the *regular edge labelings* (RELs) of G . An REL 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 four contiguous sets of incoming blue edges, outgoing red

^{*} K. Buchin and B. Speckmann are supported by the Netherlands Organisation for Scientific Research (NWO) under project no. 639.022.707.




Fig. 1. A subdivision and its augmented dual graph G , a regular edge labeling of G , and a corresponding rectangular dual

edges, outgoing blue edges, and incoming red edges; (ii) the left exterior vertex has only blue outgoing edges, the top exterior vertex has only red incoming edges, the right exterior vertex has only blue incoming edges, and the bottom exterior vertex has only red outgoing edges (see Fig. 1, red edges are dashed). Kant and He [17] show how to find a regular edge labeling and construct the corresponding rectangular dual in linear time. Regular edge labelings are also studied by Fusy [15] who calls them *transversal pairs of bipolar orientations*.

An irreducible triangulation can have many RELs and hence many rectangular duals. Depending on the specific application different duals might be desirable. For example, *sliceable* duals—which can be obtained by recursively slicing a rectangle by horizontal and vertical lines—are popular in VLSI design. Not every irreducible triangulation has a sliceable dual. A full characterization of those graphs that do is lacking, but Yeap and Sarrafzadeh [24] prove that irreducible triangulations without separating 4-cycles have a sliceable dual. *Area-universal* duals have the nice property that any assignment of areas to rectangles can be realized by a combinatorially equivalent rectangular dual. Again, not every irreducible triangulation has an area-universal dual, but Eppstein *et al.* [12] show how to find such a dual if it exists.




Fig. 2. Two different rectangular duals of the dual graph of a map of Europe. Luxembourg and Moldavia have been removed and “sea regions” have been added to ensure that the dual graph is an irreducible triangulation.

We are particularly interested in the application of rectangular duals to rectangular cartograms. A rectangular cartogram is a thematic map where every region is depicted as a rectangle. The area of the rectangles corresponds to a geographic variable, such as population or GDP. In the context of rectangular cartograms it is desirable that the direction of adjacency between the rectangles of the dual follows the spatial relation of the regions of the underlying map as closely as possible. Consider the two rectangular duals of the dual graph of a map of Europe shown in Fig. 2. The left dual will lead to a recognizable cartogram, whereas the right dual (with France east of Germany and Hungary north of Austria) is useless as basis for a cartogram. Both rectangular duals stem from the same graph G and correspond to two different valid RELs of G .

Previous work on finding RELs that lead to cartograms with geographically suitable adjacency directions has focused on finding RELs that satisfy *user-specified* constraints on a subset of the edges of the input graph. Eppstein and Mumford [11] show how to find RELs that satisfy user-specified orientation constraints, if such labelings exist for the given set of constraints. Van Kreveld and Speckmann [23] search through a user-specified subset of the RELs. Every labeling in this subset is considered acceptable with respect to adjacency directions. In contrast, we consider quality measures that take all edges of G into account and do not concentrate on a fixed, user-specified subset.

Results and organization. In this paper we consider optimization problems on RELs and show how to find optimal or near-optimal RELs for various quality criteria. Furthermore, we give upper and lower bounds on the number of RELs.


Let G be an irreducible triangulation with n vertices. Fusy [15] proves that the RELs of G form a distributive lattice. Hence, one can use reverse search to enumerate all RELs of G and so find optimal RELs for any given quality measure (see Section 2). G can have exponentially many RELs; simple upper and lower bounds are 8^n and $2^{n-O(\sqrt{n})}$. Since the running time and hence the feasibility of any enumeration algorithm depends on the number of RELs, we next give much tighter bounds. In Section 3 we show that G has less than $O(4.6807^n)$ RELs and that there are irreducible triangulations with $\Omega(3.0426^n)$ RELs. Our upper bound relies on Shearer’s entropy lemma [10]. Björklund *et al.* [7] recently used this lemma to obtain $(2 - \varepsilon)^n$ algorithms for the TSP problem. In contrast to our application of the lemma, they count vertex sets with certain properties and crucially rely on bounded maximum degree.

In Section 4 we show how to find optimal or near-optimal RELs for rectangular cartogram construction. This step of the construction pipeline has been performed essentially manually in previous work and can now finally be fully automated. We consider two quality criteria: (i) the relative position of the rectangles and (ii) the *cartographic error* of the resulting cartogram. For smaller maps enumeration of all RELs is feasible and we can find optimal solutions. For larger maps enumeration is infeasible. But the diameter of the distributive lattice of RELs is comparatively small and hence simulated annealing performs well. We present experimental results that show that our method can find RELs which result in visually pleasing cartograms with small cartographic error.

2 Useful Facts and Reverse Search

In this section we first collect some useful facts and definitions from previous work and then show how to use reverse search to enumerate RELs.

Regular Edge Labelings. Fusy [15] gives the following facts. A *regular edge coloring* is an REL, with the directions of the edges omitted. A regular edge coloring uniquely determines an REL. An *alternating 4-cycle* is an undirected 4-cycle in which the colors of the edges alternate between red and blue. There are two kinds of alternating 4-cycles, depending on the color of the interior edges incident to the cycle. If these are the same color as the next clockwise cycle edge the cycle is *right alternating*, otherwise it is *left alternating*. The set of RELs of a fixed irreducible triangulation form a distributive lattice. The flip operation consists of switching the edge colors inside a right alternating 4-cycle, turning it into a left alternating 4-cycle. An REL with no right alternating 4-cycle is called *minimal*; it is at the bottom of the distributive lattice. An REL induces no monochromatic triangles.



Perron-Frobenius theory. We need the following matrix theory for our lower bound in Section 3.2. For details refer to textbooks on matrices [16,20]. We assume from now on that A is a nonnegative $n \times n$ matrix. A matrix is *nonnegative* if all its elements are nonnegative. 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.

Theorem 1 ([16,20]). *Let A be a primitive non-negative matrix with maximal eigenvalue λ .*

- (a) λ is positive and the unique eigenvalue of largest absolute value. λ has a positive eigenvector and is the only eigenvalue with nonnegative 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 $f_A(x) \leq \lambda$ for all nonnegative non-zero vectors x , and $g_A(x) \geq \lambda$ 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 nonnegative non-zero vector. Then $A^t x / \|A^t x\|$ converges to an eigenvector with eigenvalue λ . Thus, $\lim_{t \rightarrow \infty} f_A(A^t x) = \lim_{t \rightarrow \infty} g_A(A^t x) = \lambda$.

Reverse search, proposed by Avis and Fukuda [5], is a general method for enumerating structures that match two criteria: (i) there must be a concept of “neighboring” structures such that the structures form a graph; (ii) there must be a local search operation that moves through this graph in a deterministic way and ends up at a local optimum. The local search defines a forest on the graph, of which each tree is rooted at a local optimum. If the local optima are known and we have a way of enumerating all neighbors of a structure, then we can traverse

these trees by starting at a local optimum and testing for each neighbor if the local search ends up at our current structure when applied to that neighbor. If it does, we traverse the edge and recurse.

RELS fit the criteria for reverse search: the distributive lattice is the underlying graph structure and the flip operation is the local search that ends up at the minimal labeling. We need to ensure only that the local search is deterministic. We do so by imposing an ordering on the 4-cycles of the input graph. Then, if an REL has multiple right alternating 4-cycles, we choose the first one according to this ordering. One possible ordering is to sort the vertices lexicographically by their x - and y -coordinates, use this ordering to sort the edges lexicographically by their lower and higher endpoint, and finally use this order on the edges to sort the cycles lexicographically by their lowest edge and the non-adjacent edge.

Avis and Fukuda give an implementation of their algorithm if the graph is given by an *adjacency oracle*. For this we need an upper bound δ on the number of neighbors a labeling can have: δ is the number of 4-cycles in the input graph. The oracle returns the k -th neighbor of a labeling, or \perp if that neighbor does not exist. Using our ordering on the 4-cycles of the input graph, we let the oracle return the resulting labeling after flipping the colors of all edges inside the k -th 4-cycle, or \perp if this 4-cycle is not alternating.

The enumeration takes $O(\delta t(\text{oracle})\Lambda + t(\text{local search})\delta\Lambda)$ time, where Λ is the number of RELs. We have $\delta = O(n^2)$, since a 4-cycle is defined by two edges. The oracle takes linear time, as it might have to switch the color of linearly many edges, and the local search takes quadratic time, as it might have to evaluate all 4-cycles to find the first right-alternating one, for a total of $O(n^4\Lambda)$.

Note that RELs can be represented as lower sets of a directed acyclic graph with polynomial size (see [11,12]). Using this representation one can enumerate RELs more efficiently per REL than with reverse search.

3 Counting Regular Edge Labelings

Here we prove that every irreducible triangulation with n vertices has less than $O(4.6807^n)$ RELs and that there are irreducible triangulations with $\Omega(3.0426^n)$ RELs. Before we present our bounds, we review some additional related work.

Counting all RELs of all n -vertex irreducible triangulations yields the number of combinatorially different rectangular partitions with n rectangles which is in $\Omega(11.56^n)$ [4] and less or equal to 13.5^{n-1} [14]. If we consider partitions to be identical when the incidence structure between rectangles and maximal line segments is the same, then the number of different partitions is in $\Theta(8^n/n^4)$ [2]. RELs are related to *bipolar orientations*—orientations of the edges from a source to a sink—but there is no direct relation between their numbers. Felsner and Zickfeld [13] show that the number of bipolar orientations of a planar graph is in $O(3.97^n)$ and that there are planar graphs with $\Omega(2.91^n)$ bipolar orientations. Many other interesting substructures have been counted in planar graphs (see [3,8,9]), but the upper bounds we obtained by adapting the techniques used for these structures to RELs were far from the bounds that we present.

3.1 Upper Bound

Let $G = (V, E)$ be an irreducible triangulation on n vertices. Since an REL is uniquely determined by a regular edge coloring and G has less than $3n$ edges, we obtain a simple upper bound of 8^n on the number of RELs of G . In the following we refine this bound using Shearer's entropy lemma.

Lemma 1 (Shearer's entropy lemma [10]). *Let S be a finite set and let A_1, \dots, A_m be subsets of S such that every element of S is contained in at least k of the A_1, \dots, 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 $1 \leq i \leq m$. Then we have $|\mathcal{F}|^k \leq \prod_{i=1}^m |\mathcal{F}_i|$.*

Theorem 2. *The number of regular edge labelings of an irreducible triangulation is in $O(4.6807^n)$.*

Proof (sketch). 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 REL \mathcal{L} of G let $E(\mathcal{L})$ be the set of blue edges in \mathcal{L} . Let $\mathcal{F} := \{E(\mathcal{L}) \mid \mathcal{L} \text{ is a REL of } G\}$. Since $E(\mathcal{L})$ determines \mathcal{L} , the number of RELs is $|\mathcal{F}|$.

For the vertices v_i of G , $1 \leq i \leq n$, let A_i be the set of edges in S of the triangles adjacent to v_i . 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 F_i be the set of intersections of the set A_i with the sets $E(\mathcal{L})$, i.e., F_i contains all possible ways to choose blue edges around v_i consistent with a REL. By Lemma 1 the number of RELs is bounded by $\prod_{i=1}^n |F_i|^{1/4}$.

It is easy to see that $|F_i| \leq 2^5 \binom{d_i}{4}$, where d_i is the degree of v_i . Therefore, the number of RELs of G is bounded by $(32^n \prod_{i=1}^n \binom{d_i}{4})^{1/4}$, which by convexity (and using a bound of 6 on the average degree) is upper-bounded by $(32^n \binom{6}{4}^n)^{1/4} = 480^{n/4} < 4.6807^n$. \square

3.2 Lower Bound


Our lower bound construction for the number of RELs uses triangulated grids. 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 . We add four vertices to the outside of the grid to turn it into an irreducible triangulation. The total number of vertices of the augmented grid is $n = hw + 4$.

A simple lower bound stems from the following coloring: color all horizontal edges blue and all vertical edges red. Then all diagonals can be colored independently blue or red. This gives a lower bound of $2^{n-O(\sqrt{n})}$ for $h = w$. We prove a stronger bound by coloring only the edges of every h' th row blue (for some choice of h') and by not coloring the edges of columns. We color the parts between the blue rows independently. We assume for now that $h = h' + 1$. Larger values of h do not change the analysis, but do improve the lower bound.

We first describe all steps for $h' = 1$, i.e., the edges of all rows are blue. Then we 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 must be colored red, since a REL has no monochromatic triangles. Assume we have colored the triangulated grid up to the i th column. We call the edges of the i th column and the diagonals connecting to this column from the left the i th *extended column*. How we can color the $(i + 1)$ st extended column depends on the colors of the i th extended column (assuming we have no restriction from the right).

If $h' = 1$, the previous column can be either red or blue. 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 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 (see figure).



We represent these coloring options as a transition matrix M , using the column colors as state. With M we can compute the number of colorings up to the i th extended column, by starting in the red state and repeatedly multiplying it with M . The resulting vector gives us the number of colorings ending

	To	R	B
From		2	1
		1	1

in a red or a blue edge. Since M has only positive elements, it is primitive. By Theorem 1(c) the ratio between the number of labelings ending in a red column up to the i th extended column and up to the $(i + 1)$ th extended column converges towards the largest eigenvalue of M . This eigenvalue is $\phi + 1 > 2.61803$. So for any $\varepsilon > 0$, we obtain more than $(\phi + 1 - \varepsilon)^w$ labelings for sufficiently large w . Since we add two vertices to add a single column, this yields a lower bound of $(\phi + 1 - \varepsilon)^{(n-4)/2}$ for sufficiently large w . If we now increase h , we need to add h vertices to add $h - 1$ columns, i.e., 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 $h' > 1$. We need to extend the states with information about the vertices, specifically how many color-switches there should be in the next extended column. This information, together with the color of the bottom column edge incident to this vertex, fixes the color of the top edge. So all we need for the state is the color of the bottom edge of the column and the color switches for each vertex, moving upwards. Some states that can be described in this way cannot be part of an REL. We call such states *infeasible*. A state is feasible if it can be reached from the initial all-red state (i.e., all vertical edges of the column are red) 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. We remove all infeasible states and consider only the reduced matrix.

The reduced matrix is primitive: it is irreducible by construction, and there is always at least one transition from the all-red state to the all-red state (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 1(c) the number of regular edge colorings

increases with each new column by a factor that converges to the largest eigenvalue of the 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 a strict lower bound on $\lambda_{h'}$ in the following way. We take a non-negative non-zero state vector x , multiply it with the transition matrix, and determine the minimum growth rate (for the non-zero elements). If 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 1(b). As vector x we choose $x_0 A^{100}$, where A is the reduced transition matrix and x_0 is the vector with a 1 for the all-red state and 0 otherwise. 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 rate.

We now again use several copies of h' rows beneath each other to obtain a triangulated grid with $w = h$. The growth rate per vertex in this way approaches $\lambda_{h'}^{1/h'}$. As strict lower bound on $\lambda_{h'}^{1/h'}$ we obtain 2.61803, 2.80921, 2.90453, 2.96067, 2.99746, 3.0233, and 3.04263 for $h' = 1, 2, 3, 4, 5, 6, 7$. Note that the lower bounds are rounded down, and that our upper bounds on $\lambda_{h'}^{1/h'}$ equal the (unrounded) lower bounds up to at least 10 significant digits.

Theorem 3. *The number of regular edge labelings of the triangulated grid is in $\Omega(3.04263^n)$.*

4 Optimizing RELs for Rectangular Cartograms

In this section, we describe how to find good RELs for rectangular cartogram construction and present experimental results. We follow the iterative linear programming method presented in [22] to build a cartogram from an REL.

Quality criteria. We consider two quality criteria: (i) the relative position of the rectangles and (ii) the cartographic error of the resulting cartogram. Furthermore, we bound the aspect ratio of all rectangles by 20.

To make a rectangular cartogram as recognizable as possible, it is important that the directions of adjacency between the rectangles of the cartogram follow the spatial relation of the regions of the underlying map. An REL specifies the relative directions between adjacent rectangles. We use two quality measures to quantify how well an REL matches the spatial relations between regions in the input map. The first method is based on region centroids [21]. It considers the direction between the centroids of two regions as the “true” direction of adjacency and expresses the quality of a labeling in terms of the deviation from this direction, measured as the smallest angle between the two directions (see Fig. 3 left). The centroid measure tends to perform quite well, although it can lead to counter-intuitive results in some cases. Our second measure is based on the bounding boxes of the regions. The *bounding box separation distance* (bb sep dist) measures the distance these bounding boxes would need to be moved to separate them in the direction indicated by the edge label (see Fig. 3 right). We




Fig. 3. The angle deviation and bounding box separation distance measures

consider both the average and the maximum error over all edges of a labeling, as well as a *binary* version of the measures: determine the “correct” color and direction for each edge and count the number of edges that are labeled correctly.

Another important quality criterion for cartograms is the *cartographic error* which is defined as $|A_c - A_s|/A_s$, where A_c is the area of the region in the cartogram and A_s is the specified area of that region, given by the geographic variable to be shown. As before, we consider both the maximum and average cartographic error over all regions of the cartogram.

We strive to construct cartograms of low cartographic error and high recognizability, hence we consider various ways to combine the two quality criteria. One possibility is to take weighted averages, another to bound the maximum cartographic error at 5%, while minimizing the maximum angle deviation or bounding box separation distance, which we call a *bounded* measure.

Enumeration. The augmented dual graph of the provinces of the Netherlands has only 408 RELs which can be enumerated in less than a second. Nevertheless the map is large enough to show interesting trends. Fig. 4 shows cartograms produced by enumerating all labelings and taking the best one according to various quality measures. The first data set shows total population on January 1st 2009, the second total livestock in 2009. Both were obtained from the Centraal Bureau voor de Statistiek. The color of a region corresponds to its cartographic error, with red indicating that the region is too small and blue indicating that it is too big. The saturation corresponds to the magnitude of the error, a white region has a cartographic error of at most 5%, while a fully saturated region has a cartographic error of over 30%. The figure clearly shows that combining recognizability measures with cartographic error leads to the best results.

Simulated Annealing. For larger maps—the countries of Europe or the contiguous states of the US—enumeration is infeasible (both have over four billion labelings). Fortunately the diameter of the lattice of RELs is comparatively small (115 for Europe and 278 for the US) and hence simulated annealing [18] performs well. We use a typical static cooling schedule and the original acceptance probability [1] for our experiments. Specifically, given two labelings with qualities q_1 and q_2 , the probability that our algorithm moves to the worse labeling is

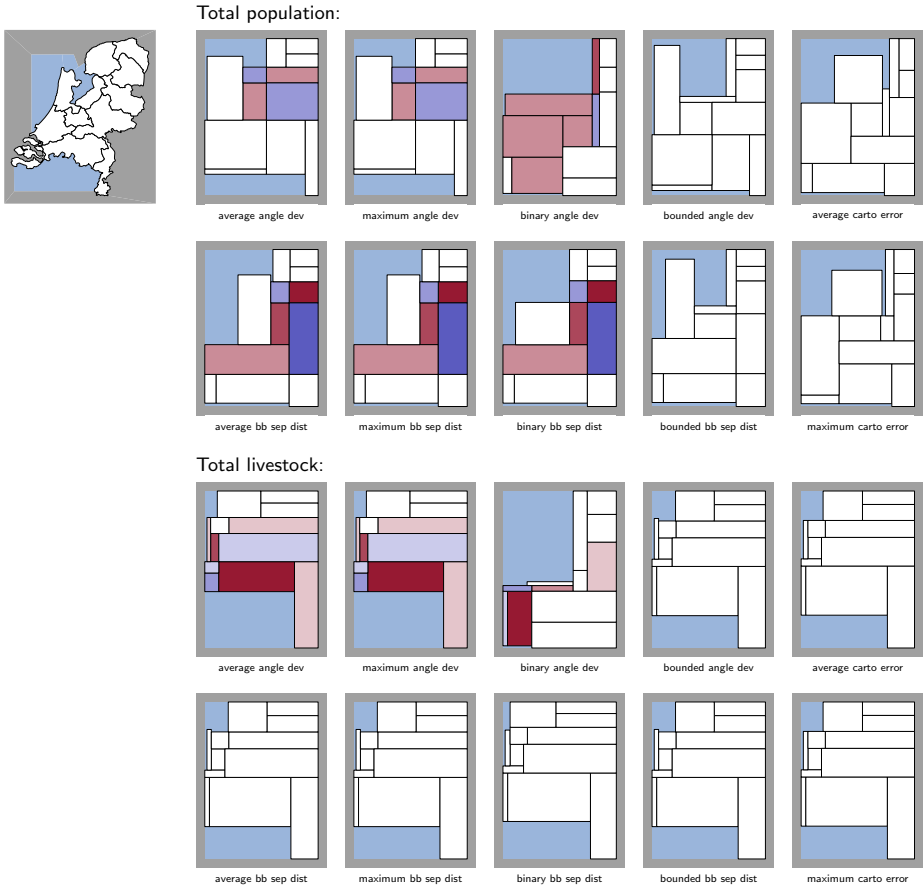


Fig. 4. Population and livestock cartograms of the provinces of the Netherlands

$e^{|q_1 - q_2|/T}$, where T is the current temperature. We let the temperature decrease exponentially as $T = 0.002^t$, where t is the current time, varying from 0 initially to 1 at the end of the process. The base factor of 0.002 can be increased to produce more random behaviour, or decreased to produce more greedy behavior.

Fig. 5 shows some results of our implementation¹. Note that we produce only cartograms with correct adjacencies. The top two figures show the total population of the countries of Europe on January 1st 2008, with the populations of Luxembourg and Moldova added to Belgium and Ukraine, respectively. The left cartogram was generated by bounding the maximum cartographic error on 5% and optimizing the average angle deviation, which results in an average

¹ Data from Eurostat <http://epp.eurostat.ec.europa.eu/portal/page/portal/eurostat/home>, the CIA World Factbook <https://www.cia.gov/library/publications/the-world-factbook/index.html>, and the the US Census Bureau <http://www.census.gov/>

References

1. Aarts, E., Korst, J., Michiels, W.: Simulated annealing. In: Search Methodologies, pp. 187–210. Springer, Heidelberg (2005)
2. Ackerman, E., Barequet, G., Pinter, R.Y.: A bijection between permutations and floorplans, and its applications. *Disc. Appl. Math.* 154(12), 1674–1684 (2006)
3. Aichholzer, O., Hackl, T., Vogtenhuber, B., Huemer, C., Hurtado, F., Krasser, H.: On the number of plane graphs. In: Proc. 17th SODA, pp. 504–513 (2006)
4. Amano, K., Nakano, S., Yamanaka, K.: On the number of rectangular drawings: Exact counting and lower and upper bounds. TR 2007-AL-115, IPSJ SIG (2007)
5. Avis, D., Fukuda, K.: Reverse search for enumeration. *Disc. Appl. Math.* 65(1), 21–46 (1996)
6. Bhasker, J., Sahni, S.: A linear algorithm to check for the existence of a rectangular dual of a planar triangulated graph. *Networks* 7, 307–317 (1987)
7. Björklund, A., Husfeldt, T., Kaski, P., Koivisto, M.: The travelling salesman problem in bounded degree graphs. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 198–209. Springer, Heidelberg (2006)
8. Buchin, K., Knauer, C., Kriegel, K., Schulz, A., Seidel, R.: On the number of cycles in planar graphs. In: Lin, G. (ed.) COCOON 2007. LNCS, vol. 4598, pp. 97–107. Springer, Heidelberg (2007)
9. Buchin, K., Schulz, A.: On the number of spanning trees a planar graph can have. In: Proc. 18th ESA (to appear, 2010), arXiv/0912.0712
10. Chung, F.R.K., Graham, R.L., Frankl, P., Shearer, J.B.: Some intersection theorems for ordered sets and graphs. *J. Comb. Theory, Ser. A* 43(1), 23–37 (1986)
11. Eppstein, D., Mumford, E.: Orientation-constrained rectangular layouts. In: Dehne, F., Gavrilova, M., Sack, J.-R., Tóth, C.D. (eds.) WADS 2009. LNCS, vol. 5664, pp. 49–60. Springer, Heidelberg (2009)
12. Eppstein, D., Mumford, E., Speckmann, B., Verbeek, K.: Area-universal rectangular layouts. In: Proc. 25th ACM Symp. Comp. Geom., pp. 267–276 (2009)
13. Felsner, S., Zickfeld, F.: On the number of planar orientations with prescribed degrees. *Electron. J. Comb.* 15(1), Research paper R77, 41 (2008)
14. Fujimaki, R., Inoue, Y., Takahashi, T.: An asymptotic estimate of the numbers of rectangular drawings or floorplans. In: Proc. ISCAS, pp. 856–859 (2009)
15. Fusy, E.: Transversal structures on triangulations: A combinatorial study and straight-line drawings. *Disc. Math.* 309(7), 1870–1894 (2009)
16. Horn, R., Johnson, C.R.: Matrix Analysis. Cambridge University Press, Cambridge (1985)
17. Kant, G., He, X.: Regular edge labeling of 4-connected plane graphs and its applications in graph drawing problems. *TCS* 172(1-2), 175-193 (1997)
18. Kirkpatrick, S.: Optimization by simulated annealing: Quantitative studies. *J. Statistical Physics* 34(5), 975–986 (1984)
19. Koźmiński, K., Kinnen, E.: Rectangular dual of planar graphs. *Networks* 5, 145–157 (1985)
20. Minc, H.: Nonnegative Matrices. Wiley-Interscience, Hoboken (1988)
21. Peuquet, D., Ci-Xiang, Z.: An algorithm to determine the directional relationship between arbitrarily-shaped polygons in the plane. *Pattern Rec.* 20(1), 65–74 (1987)
22. Speckmann, B., van Kreveld, M., Florisson, S.: A linear programming approach to rectangular cartograms. In: Progress in Spatial Data Handling: Proc. 12th International Symposium on Spatial Data Handling, pp. 529–546 (2006)
23. van Kreveld, M., Speckmann, B.: On rectangular cartograms. *Computational Geometry: Theory and Applications* 37(3), 175–187 (2007)
24. Yeap, G.K.H., Sarrafzadeh, M.: Sliceable floorplanning by graph dualization. *SIAM J. Discrete Mathematics* 8(2), 258–280 (1995)