Topological Coordinates

Alan Saalfeld

ABSTRACT: We present a new methodology for concisely encoding a map's topological structure based on an order-theoretic characterization of planar graphs. We describe how to implicitly store and recover a map's topological connectivity information by folding information about the edge structure of its embedded line-segment graph into new numerical topology-based coordinates chosen for the graph's vertices. We examine properties of our alternative geometric realization and show how to build ***cartograms*** by systematically altering our topological coordinates to modify the regions' relative areas.

KEYWORDS: barycentric coordinates, cartograms, geometric graphs, graph drawing, map topology.

Introduction

The goal of this paper is to bring some recent results in combinatorial geometric graph theory to the attention of computer mappers and GIS algorithm developers. Some key mathematical concepts that we will use are ***abstract graphs***, ***graph drawing***, ***planar graphs***, and ***plane graphs***. An ***abstract graph*** **G(V,E)** consists of two finite sets **V** and **E**. The set $\mathbf{V = \{v_1, v_2, \ldots, v_n\}}$ consists of **n** elements called ***vertices***, and the second set $\mathbf{E = \{e_1, e_2, \ldots, e_k\}}$ consists of **k** elements, each of which is a pair of the aforementioned vertices: $\mathbf{e_s = \{v_{1s}, v_{2s}\}}$, for $\mathbf{s = 1, 2, \ldots, k}$. The vertices $\mathbf{v_1}$, and $\mathbf{v_{2s}}$ are called the endpoints of e_s. The vertex-pair elements of the second set **E** are called ***edges***. An abstract graph is not a drawing of labeled points representing vertices and curved or straight, crossing or non-crossing lines linking those pairs of points that represent pairs of vertices that are edges, although such a drawing is frequently used to communicate the makeup of the vertex and edge sets. Such a drawn depiction of an abstract graph is called a ***graph drawing***. If all edges are depicted as straight line segments that do not cross each other, then the drawing will be called a ***proper graph drawing***. An abstract graph may have many different proper graph drawings; or it may have none. A ***planar graph*** is an abstract graph for which there exists some proper graph drawing. A planar graph ***realization*** is the actual drawing itself. The drawing is also called a ***plane graph*** or a ***planar graph embedding in*** $\mathbb{R}^2$. The so-called primitives or basic building blocks of our embedded plane graph maps will be points or graph vertices (**0-D** elements), straight line segments or graph edges (**1-D** elements), and connected (polygonal shaped) components (**2-D** elements) of the set complement of the graph in the plane.

When we refer to a map's ***topology***, we are referring to the adjacency and connectivity relations among the map's labeled primitive **0-D**, **1-D**, and **2-D** elements (points, line segments, and polygonal regions). The abstract graph alone accounts for the **0-D** through **1-D** topology. The way a drawn graph separates the plane gives us the remaining topological relations.

When we refer to a map's ***geometry***, we are referring to the actual positioning or location within the map's plane coordinate frame of the labeled points or vertices and the labeled segments or edges. ***Geometry*** refers to a specific graph drawing, and not to the underlying abstract graph.

Schnyder's Key Constructions and Results

In the late 1980s, Walter Schnyder discovered many beautiful and unexpected properties of realizations of planar graphs that will provide the basis for our exposition here. One of Schnyder's key discoveries was how to construct a proper graph drawing for a planar abstract graph from just the graph's edge/vertex structure. We summarize some of Schnyder's key findings and new approaches, starting with some basic definitions.

A ***triangular graph*** is a maximal plane graph on a set of **n** vertices in the plane having exactly three of the **n** vertices in its convex hull. A triangular graph is a triangulation of the **n** vertices; and a triangular graph always has **[2n – 5]** interior triangles and **[3n – 6]** edges.

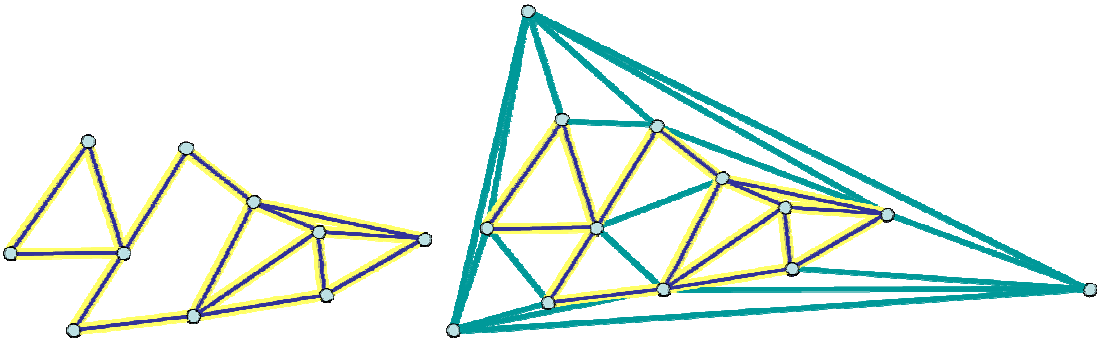

Figure 1. A plane graph and a triangular graph containing it

Every plane graph may be systematically extended to a triangular graph. One such systematic extension for a plane graph can be built, for example, from the graph alone with only three new hull vertices by first constructing a constrained Delaunay triangulation of the graph's convex hull and then placing that triangulation in a large containing triangle, as illustrated in Fig. 1.

A ***normal angle labeling*** of a triangular graph is an assignment of a number from the set **{1,2,3}** to each of the **[6n – 15]** angles of the interior triangles such that two conditions hold:

[Normal labeling condition for triangles] Every triangle contains all three different labels with "**2**" following "**1**", and "**3**" following "**2**" in clockwise order in each triangle.
[Normal labeling condition for interior vertices] At every interior vertex, there is a single non-empty sequence of "**1**" labels, followed in clockwise order by one non-empty sequence of "**2**" labels, followed in clockwise order by one non-empty sequence of "**3**" labels.

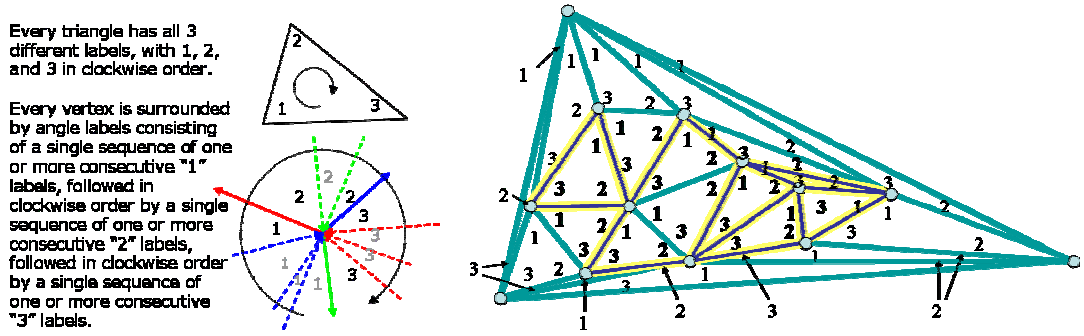

Figure 2. Criteria for an angle labeling to be normal and a normal labeling of the triangular graph

Schnyder showed by an inductive argument that normal labelings are always possible for any triangular graph. He also showed that the three extreme or hull vertices of the triangular graph each have only one label assigned to all angles at the vertex; and that one extreme vertex has all angle labels equal to "1", followed in clockwise order by the extreme vertex with only angle labels equal to "2", followed in clockwise order by the third extreme vertex having only angle labels equal to "3". We say that the extreme vertex whose angles all have the label "i" is the extreme vertex associated with the label "i".

Property 1 (Schnyder, 1989). Every triangular graph has a normal labeling of its angles.

Although the normal labeling is not necessarily unique, each normal labeling of angles gives rise to a unique edge labeling of all edges having at least one interior vertex as an end point. Such an induced edge labeling has the same three labels as the normal angle labeling.

Every interior edge of a triangular graph participates in two triangles and four angles. Since there are only 3 different labels, at least one label must appear twice among the four labels associated with an interior edge. If our angle labeling is *normal*, it turns out that there are only three possibilities for the combinations of the four angle labels that appear adjacent to an edge in a normal labeling of a triangular graph. On the right-hand side of Fig. 3, we see the possible angle labels for a "1-edge" or edge of type 1 (the blue edge shown with two associated "1" labels), a "2-edge" or edge of type 2 (the yellow edge shown with two associated "2" labels), and a "3-edge" or edge of type 3 (the red edge shown with two associated "3" labels), respectively.

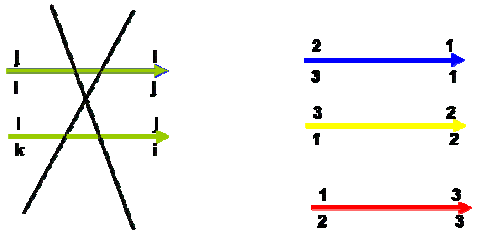

Figure 3. Impossible (left) and possible (right) angle labelings around an edge

Property 2 (Schnyder, 1989). The normal labeling of angles induces a *directed-edge labeling* on the same 3 labels, where the direction of the edge is toward the vertex with duplicate labels.

Conversely, an appropriate directed-edge labeling gives rise to a normal angle labeling: if we start with a directed edge labeling of a triangular graph that satisfies the two conditions: (1) at any interior vertex there are exactly three out-directed edges, one of each different type **1**, **2**, and **3,** appearing in clockwise order; and (2) for **{i, j, k} = {1, 2, 3}**, all in-directed edges of type **i** come into the vertex in question between the out-directed edges of the other two types, **j** and **k**, as shown in Fig. 4, then the directed edge labeling induces a unique normal angle labeling of the triangular graph.

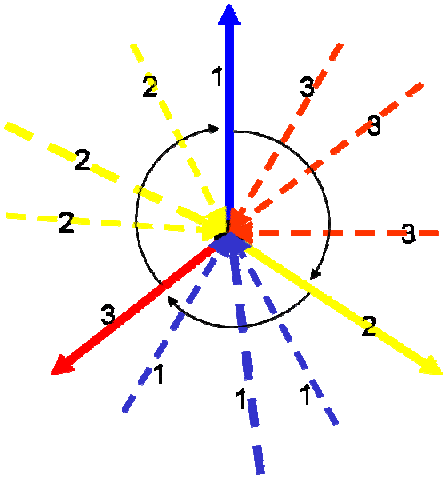

Figure 4. At each interior vertex there are exactly three out-edges

Property 3. Every internal vertex has out-degree three, and all three different edge labels each occur once on the three out-directed edges.

The set of all directed edges that have the same label, say label "i", together with all of the **(n – 3)** interior vertices plus the one extreme vertex

associated with the label "i", together form a graph that is a tree or directed acyclic graph (DAG) connecting all interior vertices to the root that is the extreme vertex associated with "**i**". If we ignore the three hull or extreme edges of the triangular graph, then we see that the graph edges are partitioned into three disjoint tree edge sets. Every interior edge belongs to exactly one of the three trees $\mathbf{T_i}$ associated with edge label "**i**". We summarize this result in the following property noted by Schnyder and illustrated in Fig. 4.

Property 4. For any edge label **i**, the set of edges with that label form a tree $\mathbf{T_i}$ linking all interior vertices to the hull vertex corresponding to that label.

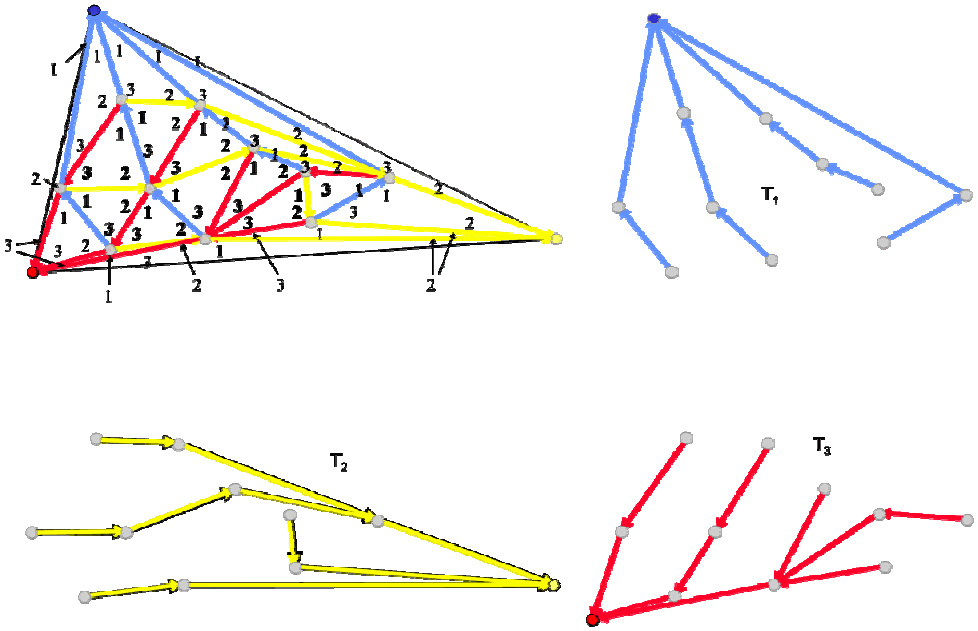

Figure 4. A normal labeling induces a partition of the directed edges into 3 trees

One important property of a tree is that, for any two vertices of the tree, there always exists a unique path between them. In particular, for any interior vertex **v**, there is always exactly one **i**-path made up of **i**-edges of $\mathbf{T_i}$ from **v** to the extreme vertex associated with the label "**i**".

For any interior vertex **v**, the three paths (one from each tree $\mathbf{T_i}$ as illustrated in Fig. 5 by the thick dashed polylines) linking **v** to the three hull vertices separate or partition the space of the triangular graph into three regions, each region being made up of some subset of all of the interior triangles.

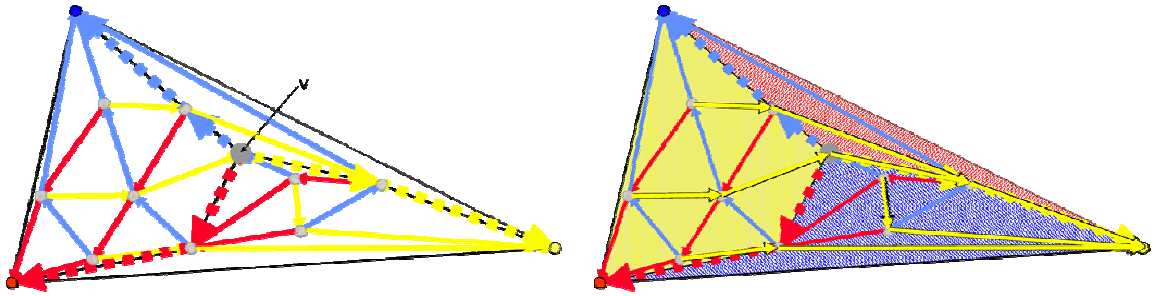

Figure 5. For every vertex v, the paths to the 3 corners 3-partition the region

If, for example, for every interior vertex **v**, we count the number of triangles in each of its three regions, and then we divide each of the three counts by the total number of triangles **[2n – 5]**, we wind up with, for every vertex **v**, three non-negative rational numbers whose sum is always equal to **1**. In our example in Fig. 5, we have **n = 13**, **[2n – 5] = 21**, and vertex **v** is associated with **7/21, 11/21,** and **3/21**, representing the fraction of triangles in the three regions opposite the extreme vertices labeled (and colored) "1", "2", and "3", respectively.

The three resulting triangle-count-fractions look like barycentric coordinates in that they sum to **1** and are all between **0** and **1**. If we use triangle-count-fractions as actual barycentric coordinate locations for the vertices in our directed-edge-labeled triangular graph, we can construct new graph drawings of our abstract triangular graph with an amazing property:

Theorem (Schnyder, 1989) If we take any three non-collinear points $\mathbf{p_1}$, $\mathbf{p_2}$, and $\mathbf{p_3}$ in the plane, and if we relocate every vertex **v** from our directed-edge-labeled triangular graph to a new position so that its triple of triangle-count-fractions are the actual barycentric coordinates with respect to $\mathbf{p_1}$, $\mathbf{p_2}$, and $\mathbf{p_3}$, then the drawing that we obtain by connecting with straight line segments every pair of vertices that constituted an edge in our original triangular graph will be a proper graph drawing of our triangular graph.

It turns out that triangle-count-fractions are not the only numerical values that can serve as barycentric coordinates for proper graph drawings. In fact, any probability measure that we can apply to the set of all regions that arise from all possible vertex 3-partitions will produce barycentric coordinates for a realization of our triangular graph with a **proper** graph drawing.

Probability Measures on Directed-Edge-Labeled Triangular Graphs

Suppose we start with a directed-edge-labeled triangular graph **G(V,E)**. For each interior vertex **v**, the three paths $\mathbf{p_1(v)}$, $\mathbf{p_2(v)}$, and $\mathbf{p_3(v)}$ from **v** to the three hull vertices labeled "**1**", "**2**", and "**3**", respectively, divide the space of the triangular graph into three closed regions $\mathbf{R_1(v)}$, $\mathbf{R_2(v)}$, and $\mathbf{R_3(v)}$, whose pairwise overlap $\mathbf{R_i(v) \cap R_j(v) = p_k(v)}$, is one of the three paths. Consider the family $\mathcal{A}$ of all regions arising as partition regions: $\mathcal{A} = \{\mathbf{R_i(v) \mid v \in V, i = 1,2,3}\}$. We say that μ is a ***probability measure*** on these region sets if μ is a function, $\mu : \mathcal{A} \rightarrow \mathbf{[0,1]}$, satisfying:

$\mu(\mathbf{R_1(v)}) + \mu(\mathbf{R_2(v)}) + \mu(\mathbf{R_3(v)}) = \mathbf{1}$ for all interior vertices **v**. (1)

If, for some **i**, **v**, and **w**, $\mathbf{R_i(v) \subset R_i(w)}$, then $\mu(\mathbf{R_i(v)}) < \mu(\mathbf{R_i(w)})$. (2)

The first condition of the definition is called the "additivity-to-1 condition," and it makes our measure a *probability* measure. The second condition is called the "monotone condition." Every probability measure associated with the sets generated by the 3-partitions provides new generalized barycentric coordinates for the vertices, where each vertex **v** is associated with the triple of values $(\mu(\mathbf{R_1(v)}), \mu(\mathbf{R_2(v)}), \mu(\mathbf{R_3(v)}))$. If we re-locate the vertices as if the new generalized barycentric coordinates were actual barycentric coordinates (for any triple of non-collinear points), then the straight line segments connecting every pair of vertices that correspond to an edge will not cross each other and will produce a proper graph drawing of the original abstract triangular graph. The triangle-count-fractions are just one example of a valid barycentric-coordinate-generating probability measure. There are many others.

One simple probability measure that would extend and generalize the triangle-count-fraction measure would be to give the different **[2n – 5]** triangles different weights that nevertheless still sum to **1** (instead of simply giving each triangle the same weight of **1/[2n – 5]**). Since

every region $\mathbf{R_i(v)}$ is composed of whole triangles, any probability measure that we define on the individual triangles can be easily extended by additivity to the 3-partition regions $\mathbf{R_i(v)}$.

The key to the proof that the probability measures provide generalized barycentric coordinates for a proper graph drawing lies in the following observation: If the vertex pair **{v,w}** is an edge in the directed-edge-labeled triangular graph **G(V,E)**, then $\mu(\mathbf{R_i(v)}) \neq \mu(\mathbf{R_i(w)})$ for all **i** in **{1,2,3}**. In other words, if two vertices are adjacent in the triangular graph as they are in Fig. 6, then none of the three generalized barycentric coordinates assigned to **v** can equal the corresponding generalized barycentric coordinate assigned to **w**. The fact that the regions $(\mathbf{R_2(v) - R_2(w)})$ and $(\mathbf{R_3(v) - R_3(w)})$ can neither be empty nor have measure zero assures that the second and third barycentric coordinates of **v** are greater than the corresponding coordinates of **w**. By the additivity-to-**1** property, the first barycentric coordinate of **v** must be strictly less than the first barycentric coordinate of **w**.

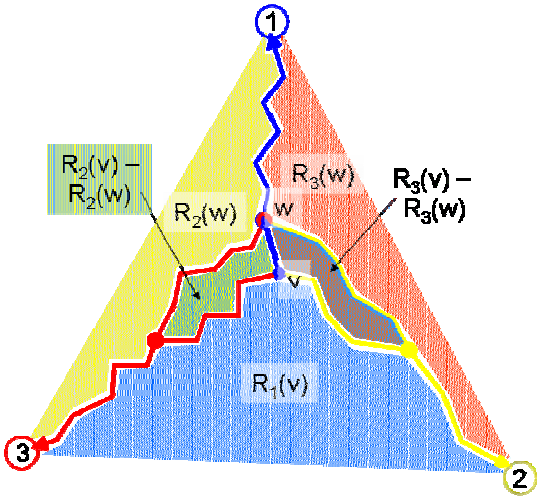

Figure 6. Endpoints of an edge have no common barycentric coordinates

We now formally define ***discrete barycentric coordinates*** (***DBC***s) for all vertices of a given directed-edge-labeled triangular graph (DELTG) **G(V,E)** in terms of some given probability measure (PM) μ applied to the three regions $\mathbf{R_1(v)}$, $\mathbf{R_2(v)}$, and $\mathbf{R_3(v)}$, of the DELTG 3-partition at **v** for each vertex **v**. The DBCs for each vertex **v** of DELTG **G(V,E)** with respect to the PM μ are the numbers of the triple $(\mu(\mathbf{R_1(v)}), \mu(\mathbf{R_2(v)}), \mu(\mathbf{R_3(v)}))$. These three numbers always sum to **1**. If we interpret the triple as actual barycentric coordinates for a new placement of each vertex **v** for all $\mathbf{v} \in \mathbf{V}$ with respect to any 3 non-collinear anchor points $\mathbf{p_1}$, $\mathbf{p_2}$, and $\mathbf{p_3}$, we may place the points and simultaneously be able to realize all of the edges of **E** as straight line segments in a proper graph drawing.

It is worth pointing out that our particular applications for Schnyder's theory involving maps and mapping are seeking to go from one placement of features to a different placement of features while preserving topological structure. We will see in the section on cartograms one situation in which a purposeful repositioning of features may be carried out. We will also show in a later section that merely by repositioning vertices according to discrete barycentric coordinates, we have actually implicitly encoded all of the edge information of our triangular graph. This rather amazing result is yet another justification for our labeling our coordinates "topological." Who would have suspected that telling where to place the vertices could also simultaneously serve to uniquely specify one and only one compatible edge structure?

It is perhaps also worth noting that Schnyder's original work was not focused on repositioning graph features, but rather on finding some good legitimate placement (a construction) for the features, working only from the abstract graph **G(V,E)** itself and from non-constructive proofs of the existence of realizations. Although our exposition here uses proper graph drawings to highlight many of our ideas, Schnyder was able to utilize properties on the abstract graph level that are equivalent to our drawing-dependent notions of normal angle labelings and directed-edge labelings.

Before we proceed further with our exposition on generating new kinds of barycentric coordinates, let us review briefly how barycentric coordinates come into play in more traditional settings.

Relation of Cartesian Coordinates to Barycentric Coordinates

Cartesian coordinates are used more than barycentric coordinates, but the two have some similarities that are useful.

Cartesian coordinates **(x,y)** permit every point **p** in the plane to be uniquely represented by two real numbers **x** and **y**, where **x** is the horizontal distance of the point from the vertical **y**-axis (negative, if **p** is to the left), and **y** is the vertical distance of the point **p** from the horizontal **x**-axis (**y** is negative if **p** is below the **x**-axis). The left-hand side of Fig. 7 illustrates that points belonging to a vertical line all have the same **x** coordinate, and all points lying on the same horizontal line all share the same **y** coordinate.

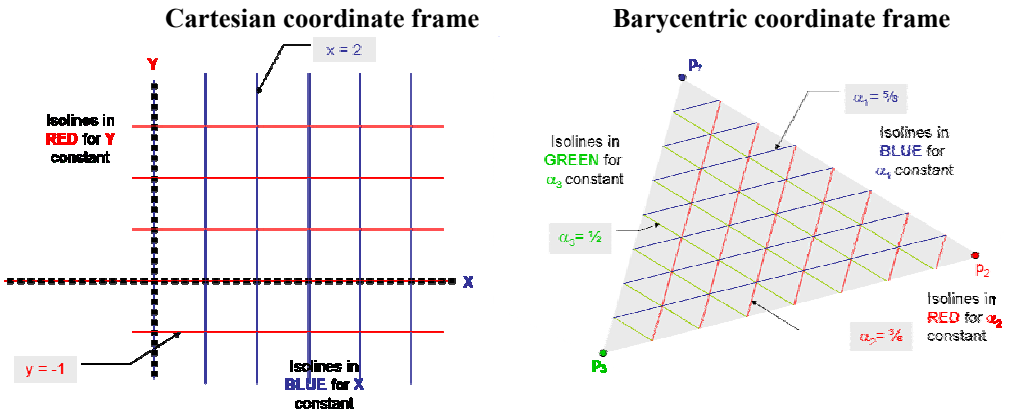

Figure 7. Lines of a constant coordinate value

Barycentric coordinates $(\alpha_1, \alpha_2, \alpha_3)$ permit every point **p** in a given triangle $\Delta \mathbf{p_1p_2p_3}$ to be uniquely represented by three real numbers α_1, α_2, and α_3, where α_1 is a weighted distance of the point **p** from the line $\mathbf{p_2p_3}$, α_2 is a weighted distance of the point **p** from the line $\mathbf{p_1p_3}$, and α_3 is a weighted distance of the point **p** from the line $\mathbf{p_1p_2}$. The weights in the three cases are proportional to the length of the side in question. The right-hand side of Fig. 7 shows that points belonging to a line parallel to $\mathbf{p_2p_3}$ all have the same α_1 coordinate, that all points belonging to a line parallel to $\mathbf{p_1p_3}$ have the same α_2 coordinate; and that points belonging to a line parallel to $\mathbf{p_1p_2}$ all have the same α_3 coordinate.

Barycentric coordinates provide a unique location to points within a triangle, and, by natural extension, to every point in the plane. Given any three points in the plane $\mathbf{p_1} = (\mathbf{x_1}, \mathbf{y_1})$, $\mathbf{p_2} = (\mathbf{x_2}, \mathbf{y_2})$, and $\mathbf{p_3} = (\mathbf{x_3}, \mathbf{y_3})$, we form ***linear combinations*** of the points as follows:

$$\alpha_1\mathbf{p}_1 + \alpha_2\mathbf{p}_2 + \alpha_3\mathbf{p}_3 = (\alpha_1x_1 + \alpha_2x_2 + \alpha_3x_3, \alpha_1y_1 + \alpha_2y_2 + \alpha_3y_3),$$

where α_1, α_2, and α_3 are any real numbers.

If the three points $\mathbf{p_1}$, $\mathbf{p_2}$, and $\mathbf{p_3}$ are not collinear, then every point $\mathbf{p}$ in the plane may be written as some (not necessarily unique) linear combination of the three points.

If, however, we further require that $\alpha_1 + \alpha_2 + \alpha_3 = \mathbf{1}$, then for any point $\mathbf{p}$, the choice of the α_1, α_2, and α_3 will still exist, but the triple of numbers will also be uniquely determined.
If all of the α_1, α_2, and α_3 are non-negative and add up to **1**, then they must also all be less than **1**, and the point $\alpha_1\mathbf{p_1} + \alpha_2\mathbf{p_2} + \alpha_3\mathbf{p_3}$ will lie inside the triangle formed by $\mathbf{p_1}$, $\mathbf{p_2}$, and $\mathbf{p_3}$.

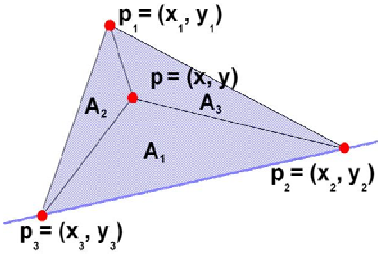

Figure 8. The geometry of barycentric coordinates

Moreover, for every point p on or inside the triangle, there is a unique triple of numbers $(\alpha_1, \alpha_2, \alpha_3)$ that satisfy:

$$\mathbf{p} = \alpha_1\mathbf{p_1} + \alpha_2\mathbf{p_2} + \alpha_3\mathbf{p_3} \qquad (1)$$
$$\mathbf{1} = \alpha_1 + \alpha_2 + \alpha_3 \qquad (2)$$
$$\mathbf{0} \le \alpha_i \le \mathbf{1}, \text{ for } \mathbf{i = 1, 2, 3} \qquad (3)$$

The three numbers α_1, α_2, and α_3 that make up this triple are called the ***barycentric coordinates*** of $\mathbf{p}$ with respect to $\mathbf{p_1}$, $\mathbf{p_2}$, and $\mathbf{p_3}$.

The values of α_1, α_2, and α_3 are easy to interpret and evaluate: If $\mathbf{\{i,j,k\}=\{1,2,3\}}$, and if $\mathbf{A_i}$ is the area of $\Delta\mathbf{pp_jp_k}$, the inside triangle opposite the vertex $\mathbf{p_i}$, then:

$$\alpha_i = \mathbf{A_i}/(\mathbf{A_i}+\mathbf{A_j}+\mathbf{A_k}) = \frac{\begin{vmatrix} x & y & 1 \\ x_j & y_j & 1 \\ x_k & y_k & 1 \end{vmatrix}}{\begin{vmatrix} x_i & y_i & 1 \\ x_j & y_j & 1 \\ x_k & y_k & 1 \end{vmatrix}}.$$

Thus the α_i's are the relative sizes of the three interior triangles. If we drop condition (3) and allow the values of the α_i's to take on negative values and values larger than **1**, but continue to insist that the α_i's sum to **1**, we can still uniquely represent any point p in the entire plane (inside, on, or outside the triangle) by $\mathbf{(\alpha_1, \alpha_2, \alpha_3)}$, and we will call the triple the ***plane barycentric coordinates*** of $\mathbf{p}$ with respect to $\mathbf{p_1}$, $\mathbf{p_2}$, and $\mathbf{p_3}$. We can now state quite simply a relation between barycentric and Cartesian coordinates

Proposition: There exist three points $\mathbf{p_1}$, $\mathbf{p_2}$, and $\mathbf{p_3}$ in the plane such that the plane barycentric coordinates $\mathbf{(\alpha_1,\alpha_2,\alpha_3)}$ of any point $\mathbf{p}$ with respect to $\mathbf{p_1}$, $\mathbf{p_2}$, and $\mathbf{p_3}$ are related to the Cartesian coordinates $\mathbf{(x,y)}$ of that same point $\mathbf{p}$ by: $\alpha_1 = \mathbf{x}$, and $\alpha_2 = \mathbf{y}$.

Proof: Fig. 9 shows that one may choose $\mathbf{p_1 = (1, 0)}$, and $\mathbf{p_2 = (0, 1)}$, and $\mathbf{p_3 = (0, 0)}$, so that: $\alpha_1\mathbf{p_1} + \alpha_2\mathbf{p_2} + \alpha_3\mathbf{p_3} =$
$\alpha_1(1, 0) + \alpha_2(0, 1) + \alpha_3(0, 0)$
$= (\alpha_1, 0) + (0, \alpha_2) + (0, 0)$
$= (\alpha_1, \alpha_2)$
$= (x, y)$.
Thus, $\alpha_1 = \mathbf{x}$, and $\alpha_2 = \mathbf{y}$.

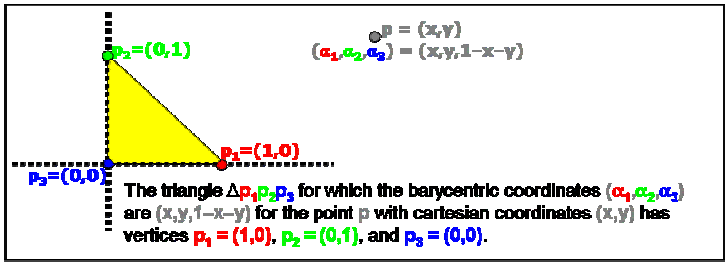

Figure 9. The first two (plane) barycenctric coordinates equal the Cartesian coordinates

The last proposition states that we may simply treat the first two barycentric coordinates as if they were Cartesian coordinates if we wish to make some barycentric coordinate graph drawing for some choice of 3 non-collinear anchor points. Suppose we were to take 3 different non-collinear points as our anchor points, but keep the same barycentric coordinate values for our features, and apply them to the new anchor points. This would amount to nothing more or less than applying an affine transformation of the whole plane that moves the old three points to the new three points (and all our features accordingly). In particular, any affine transformation sends lines to lines, parallel lines to parallel lines, intersecting line segments to intersecting line segments, and non-intersecting line segments to non-intersecting line segments. Consequently an affine transformation of the plane will always send any proper graph drawing to a proper graph drawing. Therefore, the choice of anchor points for drawing our triangular graph will not matter. Our graph drawing will either be always proper or never proper.

Discrete Barycentric Coordinates

The values that we compute for our discrete barycentric coordinates arise from some combinatorial probability measure μ defined on our 3-partition regions. Since each interior vertex gives rise to **3** regions, and since there are [$\mathbf{n - 3}$] interior vertices, our probability measure μ needs only to be defined on [$\mathbf{3n - 9}$] regions and satisfy the "additivity-to-1" and monotonicity conditions there. Here are six candidate probability measures for producing new discrete barycentric coordinates while maintaining the same topology:

1. $\mu(R_i(v))$ = [# of Δs in $R_i(v)$]/[2n – 5]
2. $\mu(\mathbf{R_i(v)})$ = **[**# of vertices in $\mathbf{R_i(v)}$ **–** # of vertices on $\mathbf{p_{i+1}(v)}$**]/[n – 1]**
3. $\mu(\mathbf{R_i(v)})$ = **[**# of edges in $\mathbf{R_i(v)}$ **–** # of edges on $\mathbf{p_{i+1}(v)}$**]/[3n – 6]**
4. $\mu(\mathbf{R_i(v)})$ = **[**# of interior vertices in $\mathbf{R_i(v)}$ **+ ½(**# of boundary vertices in $\mathbf{R_i(v)}$**)]/[n + ½]**
5. $\mu(\mathbf{R_i(v)})$ = **[**Area of $\mathbf{R_i(v)}$**]/[**Area of the whole triangular graph]
6. $\mu(\mathbf{R_i(v)})$ = **[**Population of $\mathbf{R_i(v)}$**]/[**Total population**]**

In example 6 immediately above, for monotonicity, we must always have strictly smaller populations associated with proper subregions. This is equivalent to requiring that all difference regions, such as the two depicted in Fig. 6, must have positive populations.

Note especially that any conversion via any valid μ to barycentric coordinates followed by redrawing the graph with vertices in new positions does not change the topological relations of the original graph. If the computation of the new barycentric coordinates only depends on topological properties, as it does in examples 1 through 4 above, then using the newly drawn graph to once again try to generate new barycentric coordinates will result in the same coordinates generated by the first iteration. Another way of describing that outcome is to say that generating new barycentric coordinates is an idempotent operation for the probability measures in examples 1 through 4 above. Applying the operation twice gives the same output as applying the operation once.

Cartograms

Our ability to now change vertex placements (in a number of different ways using different probability measures) while still maintaining map topology is a natural impetus to ask ourselves if we can find a measure μ that generates a specific desired outcome such as being able to achieve pre-specified area ratios among map regions. Such a measure would be ideal

for building ***cartograms***, maps with areas altered to reflect the size of some variable associated with the regions. Although our search for good measures to accomplish area redistribution could start with a specific map and an embedding of that map into a triangular graph, we will focus, as did Tobler (2004), instead on how to modify some regular underlying grid to accomplish our desired area re-distribution. Once we know how to modify the graticule, so to speak, we can easily redraw map features with respect to their appropriate placement within the deformed graticule.

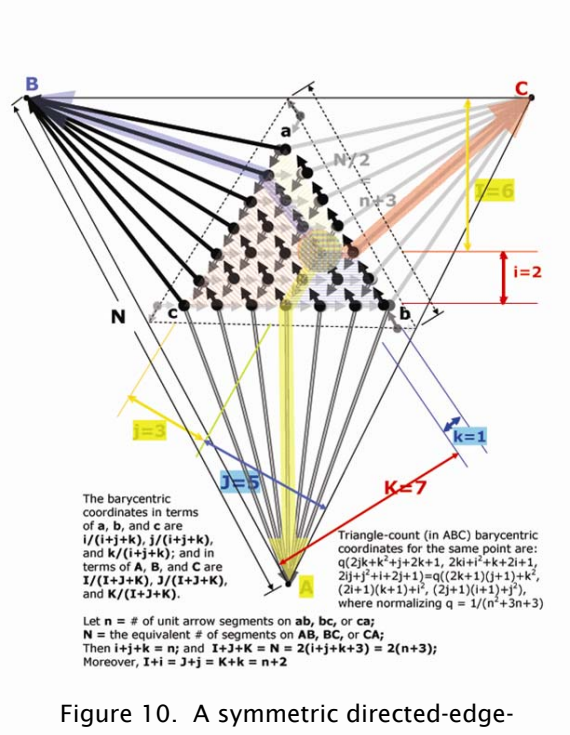

Figure 10. A symmetric directed-edge-labeled triangular graph for which μ will be

We will build a special triangular grid that has the property that its triangle-count-fraction barycentric coordinates are equal to its actual barycentric coordinates. To begin, we first build a nice symmetric directed edge-labeled triangular graph, illustrated in Fig. 10, for which we can easily count triangles in each of the regions $\mathbf{R_i(v)}$. The example drawn in Fig. 10, although small in complexity, is sufficiently large enough to allow us to describe the behavior for arbitrarily large symmetric triangular graphs with the same shape and structure (see inset notes in Fig. 10). Using the formulas shown in Fig. 10, and larger values for **n** and **N**, we redraw the triangular graph using triangle-count-fraction barycentric coordinates to produce Fig. 11. Finally, if we superimpose our special grid on a map, as shown also in Fig. 11, we may use geographic data to populate map regions, and hence grid triangles, with probability weights corresponding to geographic data sizes. These new weights will provide a probability measure as in example 6 of the previous section, from which we may compute new barycentric coordinates for the triangular grid points and consequently for map points by their positions within the initial and deformed grid.

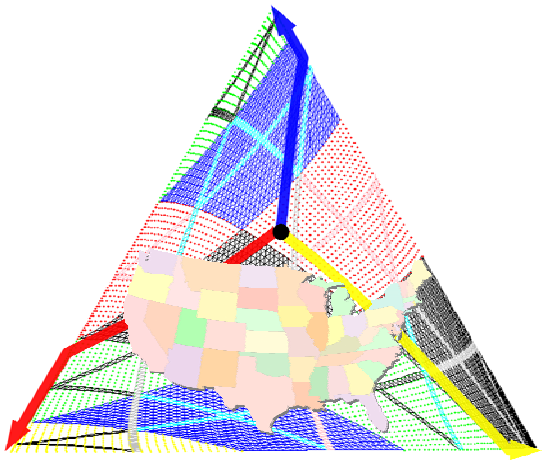

Figure 11. A transformed triangular graph for which triangle-count-fraction coordinates actually equal barycentric coordinates

There are obviously many different ways of using other discrete barycentric coordinates to redraw a topologically equivalent map. We encourage others to experiment with these techniques and further explore cartogram building with the tools we have outlined. Instead of examining other applications of discrete barycentric coordinates, we finish by highlighting the property that makes them truly topological coordinates—they encode the full edge structure of the triangular graph.

Edge recovery from discrete barycentric coordinates

Suppose that you are given some set of discrete barycentric coordinates for a collection of vertices. The fact that they are DBCs means that they came from a probability measure with the monotonicity property applied to a directed-edge-labeled triangular graph. The important consequence of that origin is that the triangular graph itself has a 3-tree edge-independent decomposition, as we illustrated in Fig. 4. We describe now how to recover the edges for each of those three trees. Here is what we need to check about barycentric coordinates to see if the vertex pair **{v, w}** corresponds to a type 1 (or blue) edge directed toward **w**. There are five other equivalent tests we can perform to cover all of the other edge possibilities: a blue edge directed toward **v**, a yellow edge directed toward **v** or **w**, and a red edge directed toward **v** or **w**. If the vertex pair fails all six tests, then there is no directed edge connecting the two vertices. If the vertex pair passes one of the tests, then we have found the edge (and its color and direction). So, here is the test (based on partial orderings of the barycentric coordinate triple):

Suppose **v** has DBCs equal to $(\alpha_1, \alpha_2, \alpha_3)$ and **w** has DBCs equal to $(\beta_1, \beta_2, \beta_3)$. Then a necessary condition for **{v, w}** to be an edge of type 1 (blue) direct toward **w** (as we can verify from Fig. 6) is that all three following strict inequalities must hold:

$$\alpha_1 < \beta_1 \qquad (1)$$
$$\alpha_2 > \beta_2 \qquad (2)$$
$$\alpha_3 > \beta_3 \qquad (3)$$

If even one of these inequalities fails to hold, then there can be no directed blue edge from **v** toward **w**. Suppose that all the inequalities hold. Is there some way that **{v, w}** can still fail to be an edge? Yes. The blue directed acyclic graph is transitive with respect to those inequalities. So the inequalities would hold if there exists a directed blue path from **v** to **w**. Therefore we need another test to rule out a non-direct connection from **v** to **w**. The other test would look for and rule out any intermediate vertices on the path:

For every other vertex **u** with barycentric coordinates equal to $(\gamma_1, \gamma_2, \gamma_3)$, it is <u>not</u> the case that:

$$\alpha_1 < \gamma_1 < \beta_1 \qquad (4)$$
$$\alpha_2 > \gamma_2 > \beta_2 \qquad (5)$$
$$\alpha_3 > \gamma_3 > \beta_3 \qquad (6)$$

In the language of partial orders on vertices with DBCs, we can define "$\mathbf{v} <_1^* \mathbf{w}$" to mean that conditions (1) – (3) in the preceding paragraph hold, and we say that **w** "$<_1^*$-covers" **v** if there exists no **u** such that $\mathbf{v} <_1^* \mathbf{u} <_1^* \mathbf{w}$. In other words, conditions (4) – (6) do not hold for any **u**.

Some Final Remarks and Conclusions

GIS users traditionally talk about "building map topology from geometry." What they actually do is fix both geometry and topology by identifying primitives or fundamental topological building blocks and their (new) geometry. They do this by first finding line intersections and polygonal regions from so-called input "spaghetti data," then by identifying the topological primitives either as subsets of the original input geometric entities or as connected components of the complement of the original input geometric entities. Our work does sort of the opposite: we have used a map's topology to compute a geometry, a placement of the geometric features. We

also "rebuilt geometry." We started with a geometry, built a topology, then built alternative geometries, different placements of the geometric features. We described ways of finding new coordinates for all of the map's vertices based only on the map's topology. Those newly calculated coordinates were used to designate where to position or *re-position* the vertex points (and, for that reason, we referred to the new position information as "topological coordinates"). We saw that if we connected by straight line segments every newly positioned pair of vertices that had been a graph edge of the original topological graph, then we would create a straight-line-segment realization of the original topological graph. Some of the tools we used to discover a vertex placement were based only on topology and did not use any *a priori* geometry knowledge. Such tools were first used by Schnyder to find a realization of an abstract planar graph.

We chose our title, "Topological coordinates," to arouse the reader's curiosity because the title is a seeming oxymoron. *Coordinates* usually provide a fixed space-regular geometric framework such as $\mathbb{R}^2$ for anchoring map feature locations, whereas *topology* usually deals with a relational structure among map features that goes beyond any particular location assignment of those features in $\mathbb{R}^2$ or other **2-D** coordinate space. We hope that the reader has gained a little bit more insight into some of the recently discovered interplay between geometry and topology through our exposition.

REFERENCES

S. Felsner, ***Geometric Graphs and Arrangements***, Veiweg, 2004.

B. Mohar and C. Thomassen, ***Graphs on Surfaces***, Johns Hopkins University Press, 2001.

W. Schnyder, Planar graphs and poset dimension, *Order*, 5(2): 323–343, 1989.

W. Schnyder, Embedding planar graphs on the grid, *Proceedings of the First ACM Symposium on Discrete Algorithms*: 138–148, 1990.

W. Tobler, Thirty-five Years of Computer Cartograms, *Annals of the Association of American Geographers*, March, 2004.