Drawing Clustered Graphs as Topographic Maps

Martin Gronemann and Michael Jünger

Institut für Informatik, Universität zu Köln, Germany
{gronemann,mjuenger}@informatik.uni-koeln.de

Abstract. The visualization of clustered graphs is an essential tool for the analysis of networks, in particular, social networks, in which clustering techniques like community detection can reveal various structural properties.

In this paper, we show how clustered graphs can be drawn as topographic maps, a type of map easily understandable by users not familiar with information visualization. Elevation levels of connected entities correspond to the nested structure of the cluster hierarchy.

We present methods for initial node placement and describe a tree mapping based algorithm that produces an area efficient layout. Given this layout, a triangular irregular mesh is generated that is used to extract the elevation data for rendering the map. In addition, the mesh enables the routing of edges based on the topographic features of the map.

1 Introduction

Clustered graphs are able to express relationships between entities and, at the same time, hierarchies on those entities in the form of a nested system of sets of entities called clusters. This special combination has turned out to be very useful in many areas. Our main motivation, however, has been the analysis and visualization of social networks. Here, the entities are persons whose bilateral relationships are captured by edges while nested communities can be expressed by clusters. In particular, citation and collaboration networks have been subjects of increased attention in recent years. Such networks can offer insights in the publication behavior in different research areas or can be used to measure the performance of people, institutions or other entities.

Clustered graph drawing is difficult because the relational data and the hierarchy must be visualized in one picture. For the relational data alone, a variety of graph drawing methods is available, and for the cluster structure alone, various tree drawing methods can be applied, because by their nested nature, the cluster system corresponds to a tree with the set of entities as the root. Graphs are usually drawn using geometric shapes for the nodes corresponding to the entities, and the relations correspond to the edges that connect nodes by lines or curves. For trees, similar methods can be applied, but there exist alternatives that can be used to create appealing drawings. Combining both in one drawing without producing visual clutter or occlusion is a difficult task.

Users of clustered graph drawing in the analysis and visualization of large software systems are typically software engineers, i.e., people who are used to (and happy with) technical drawings. In contrast, users of social network analysis and visualization are much less likely familiar with technical drawings. We have been wondering what would

be the most intuitive clustered graph visualizations for these “non-technical” people and came up with the idea of trying to visualize social network data as topographic maps. In fact, hardly anyone can avoid learning to read “real” maps already as a child. Modern services like Google Earth or Google Maps have certainly boosted this ability. Therefore, such an approach is likely to find interest also outside social network analysis and visualization.

Topographic maps are detailed graphic representations of the physical features of an area. This makes them a crucial tool for land, air and nautical navigation. The graphical representation often includes contours or isoclines. These lines follow the contours of the terrain at a specific elevation level to visualize terrain features like hills and valleys. In addition, elevation levels have natural color encodings.

In this paper, we describe how we succeeded in drawing clustered graphs as topographic maps. The basic idea is to generate a landscape where the nodes in different subtrees of the cluster hierarchy are separated by water, a valley, or a rift. With increasing distance from the root cluster, nodes will be located in the lowlands, the highlands, and ultimately, on mountain peaks. Given this idea, we present a method for visualizing the complete hierarchy and the underlying graph as a topographic map.

Any method implementing this idea must meet a few requirements: First, the usage of elevation levels requires that the nodes are placed in a nested manner, i.e. nodes of the same cluster have to form a compact area such that they can be placed on one plateau. Since treemap-like techniques have this property by definition, they are an obvious choice as a basis for the layout. Second, for technical reasons, the drawing area should be as small as possible and this supports our choice. Most of the treemap approaches are based on orthogonal primitives that are most unlikely going to produce a natural looking map. This led us to use convex polygons as subdivision primitives. While *Voronoi tree-maps* satisfy all our needs, it turned out in practice that many clusters contain only two sub clusters. The method proposed by Balzer et al. [2] uses an iterative algorithm to find an approximate solution for convex polygon partitioning that can easily be solved exactly for two subclusters. The *fat polygon partitioning* approach of de Berg et al. [3] exploits this property while finding a partition with a good aspect ratio.

Before describing their algorithm, our modifications and how these can be used to generate a topographic map, we briefly review in the next section related work on tree maps, drawing clustered graphs and edge routing. In Section 3, we then present the major components of our method: As our main contribution, we will then

- describe the method for node placement based on the cluster hierarchy,
- show how we generate a 2.5D triangle mesh based on this placement and the cluster hierarchy,
- and explain how we use this mesh as a basis for an edge routing graph with whose help the edges can be drawn as quadratic curves following the terrain features.

In Section 4 we provide details of our implementation and the tools we used, including a short discussion of the performance. Finally, we present the results in Section 5 by applying our technique to selected instances followed by a conclusion and outlook in Section 6.

2 Related Work

The method we propose builds on a number of previous developments in graph drawing and information visualization in general, of which we cover the most influential aspects.

Recently, Fabrikant et al. [9] examined the effect of using the landscape metaphor for information visualization in more detail. The topographic map metaphor has a long tradition in information visualization. Especially in visualizing point distributions arising from, e.g., *multidimensional scaling (MDS)* or *self-organizing maps (SOM)*. One of the most prominent examples in this area is *ThemeScape* [22] that is used to visualize patents as topographic maps. But even in graph drawing, using the topographic map metaphor is not new. Cortese et al. [7] use it for visualizing hierarchical networks in a radial style.

However, to the best of our knowledge, it has not been applied to drawing clustered graphs. The visualization of clustered graphs that combine ordinary graphs with cluster trees has been studied quite intensively and many different approaches and styles exist, see [8] for an introduction. Ganser et al. [10] have proposed *GMap*, a system for drawing graphs as maps. Given a node partition, they show how to draw a graph as a political map using a Voronoi diagram based shape. The initial graph layout neglects the additional information given by the partition, like basic force-directed methods. Furthermore, an algorithm is proposed to reduce fragmentation, i.e. reduce the number of clusters that do not form a compact shape. In addition, sophisticated map coloring techniques are applied.

Another way of visualizing large hierarchies are so called *level-of-detail (LOD)* techniques. The idea is to show only a few levels that are of current concern to the viewer. Balzer et al. [1] use implicit surfaces for visualizing clusters this way. In addition, they propose an LOD technique for drawing inter-cluster edges. A more treemap-like approach is proposed by Garland et al. [11], where a combination of Voronoi treemaps, earlier proposed by Balzer et al. [2], and a force directed layout is used to draw clustered graphs.

Treemaps, first proposed by Johnson and Shneiderman [17], have been very popular in information visualization, since they offer an alternative to the node-link representation. Most algorithms for generating treemaps follow the same principle. A predefined area/shape is recursively subdivided into smaller parts according to the input tree. The result is a nested set of shapes contained in the root shape. Most approaches use rectangular shapes, while some use convex polygons like the earlier mentioned Voronoi treemaps [2]. The work of de Berg et al. [3] provides a detailed analysis of how to partition convex polygons such that the resulting shapes have a good aspect ratio. We use their results in this paper and, therefore, describe these in more detail in Section 3.

Drawing edges on top of a hierarchical layout raises various problems like, e.g., visual clutter. *Edge bundling* is one way to overcome this problem. Holten et al. [16] propose a very generic way to integrate this technique into hierarchical visualizations. In [18], Lambert et al. use a hybrid quad tree/Voronoi approach for routing and bundling edges for a given layout. Besides the construction of the routing network, they describe various ways to speed up the shortest path based routing of edges. Our approach for routing edges is similar to their idea, but instead of using a dynamic cost function for bundling edges, the costs are based on the elevation levels.

3 The Algorithm

Before describing the different steps of our technique for drawing topographic maps, we briefly recall the definition of a clustered graph. A *clustered graph* consists of a graph $G = (V, E)$ with node set V and edge set E as well as a tree $T = (V_T, E_T)$ whose leaves are exactly the nodes in V . Every inner node $C \in V_T$ of the tree is referred to as a *cluster node* that defines a subset of V consisting of all leaves of the subtree of T rooted at C . Thus T defines a hierarchy consisting of a nested system of subsets of V called *clusters*, the *root cluster* V corresponds to the root of T . Every cluster is thus endowed with a *cluster hierarchy level* that is the graph-theoretic distance of its corresponding cluster node to the root of T , the root cluster's hierarchy level is 0 and clusters whose corresponding cluster nodes have maximum distance from the root are at the top of the hierarchy. We shall later assign an elevation level to each node of G equal to the maximum cluster hierarchy level of the clusters it belongs to.

3.1 Clustering Method

The clustering method establishes an indirect link between the edge set and hierarchy. We will later apply the fat polygon partition, which is not aware of the underlying graph, and requires this link. In the following we will give a brief description of the clustering algorithm we use.

For our instances we used the edge betweenness based algorithm proposed by Girvan and Newman [14]. Edge betweenness measures the number of shortest paths an edge is part of. See [5] for details on betweenness and variants. The algorithm of Girvan and Newman [14] for detecting communities is rather simple. It calculates betweenness for all edges and removes the one with the highest score. This procedure is repeated until the graph becomes disconnected. The connected components serve as clusters with cluster hierarchy level 1. The partitioning is recorded by creating the corresponding cluster subtrees and we can recurse on the connected components to obtain clusters with cluster hierarchy level 2, and so on, until no edges are left to remove.

The intuition is that if a part of the graph is well connected, the number of shortest paths using a given edge in this part is relatively low. Conversely, this number is relatively high in sparse parts of the graph. Thus edges connecting clusters will have a higher betweenness score compared to intra cluster edges. Removing edges with higher betweenness increases this effect until the graph becomes disconnected. We modified the above algorithm slightly such that in case there are multiple edges with the highest betweenness, we remove all of them simultaneously. This modification avoids artifacts in the hierarchy caused by some internal order of the edges.

We also use this algorithm for decomposing collaboration networks with weighted edges. The weights correspond to the amount of collaboration between two authors. For a weighted version of betweenness, see Brandes [5]. After some experiments we found that the most intuitive integration, that is to divide the final betweenness score of an edge by its weight, gives the best results. This modification matches the intuition of the clustering algorithm, that is light edges are more likely to be removed compared to heavy edges.

3.2 Fat Polygon Partitioning Placement

In the following, we describe the placement of the nodes based on the work of de Berg et al. [3]. The basic idea of *fat polygon partitioning* is to create a nested structure of convex polygons for a weighted binary tree. In our case, the weight is chosen as the area required for a subtree, i.e. the number of graph nodes contained in the corresponding cluster. At each internal node, a cutting line is chosen that subdivides the subtree's boundary polygon into two convex subpolygons with the appropriate area depending on the weights of the two children. We place the graph nodes, i.e., the leaves in the tree, in the centroid of the boundary polygon computed by the partitioning. The objective is to subdivide in a top down manner while obtaining polygons with a small aspect ratio. Following the notation of [3], the aspect ratio of a convex polygon P is defined as $\text{asp}(P) = \text{diam}(P)^2 / \text{area}(P)$, where $\text{diam}(P)$ is the diameter and $\text{area}(P)$ the area of P . The diameter of a polygon P is the maximum distance between two vertices of P .

In the following, we describe our implementation of the greedy algorithm presented in [3] to obtain a partition with a good aspect ratio. For more details we refer the reader to the original paper [3]. We are given a convex polygon P with k vertices and a parameter $0 < a \leq \frac{1}{2}$, where a is the fraction of the area we require for the smaller child. We want to find a direction for a cutting line that partitions P into two subpolygons P_1 and P_2 such that $\text{area}(P_1) = a \cdot \text{area}(P)$ and $\text{area}(P_2) = (1 - a) \cdot \text{area}(P)$.

When given a cut direction, we choose the orientation of the cut perpendicular to this direction. Finding such a cut is easy for a convex polygon. However, for a given cut direction and $a < 1/2$, we can always cut in two ways. In that case, we choose the cut where the maximum aspect ratio of the two polygons is minimized. In order to find the direction for a cut with a good aspect ratio, de Berg et al. [3] distinguish between two main cases.

In the first case, when we want to cut off a small piece, that is when $a \leq 1/k^2$, we take the bisector at the vertex with the smallest interior angle as cut direction. In the second case, when $a > 1/k^2$, we work a little harder for finding a balanced cut. In this case the proof in [3] contains two subcases that distinguish different shapes of P . When P has a good aspect ratio, i.e., $\text{asp}(P) \leq k^6$, we are allowed to choose any direction for a cut. In order to avoid too many diagonal cuts along the diameter, we cut horizontally or vertically, depending on which results in the best aspect ratio.

Otherwise, that is when $a > 1/k^2$ and $\text{asp}(P) > k^6$ holds, we choose the direction of the line representing the diameter of the polygon. This is the line that connects any two vertices of P with the maximum distance to each other.

The above algorithm constructs a proper polygon partition based on the structure and weights of the tree. We can now place each graph node at the centroid of its polygon and obtain a layout for the next step. But two problems arise: First, the shapes of the clusters do not look very appealing. In order to fix that, we apply some post processing to the boundary polygons after their computation, but before they are used as boundary polygons for their children. The idea is to round off sharp corners by cutting off small fractions, in our implementation about two percent of a polygon. Investing small amounts of area for a better shape is inspired by the partitions with slack in [3], here for general convex polygons rather than rectangles.

Fig. 1. Comparison of the two balancing approaches. The left picture (a) shows the original balancing. In (b), the result of our modified version is shown.

The second problem is that we can only use a binary tree as input for the partitioning, thus at least a temporary transformation of a the non binary cluster tree into a binary tree is required. We are in the lucky situation that most of the internal nodes of the tree generated by the aforementioned cluster algorithm are binary, but some of them are not. We replaced the original method of [3] by another variant proposed later by de Berg et al. [4]. The presented algorithm for creating a binary tree produces unpleasant drawings in some cases. The reason is that many real-world instances like, e.g., the graph drawing collaboration network [13] contain one big connected component and many small ones. The algorithm classifies the subtree containing the big connected component as a “heavy” child, putting it alone in one subtree. All the other smaller connected components are assigned to the second subtree. When this node is then partitioned, the large connected component ends up on one side, while all the small components are located on the opposite side (see Figure 1a).

In order to solve this problem, a slightly different transformation is used. The idea is to push the big child further down the tree. In the given order, we assign nodes to the first subtree until the weight ratio of that subtree exceeds c/k^2 for some constant c (our implementation uses $c = 2$). If no such node exists, the smallest element is used. Then we assign the rest to the second subtree, and recurse on these. As a result, the small pieces will fill up the corners first, and the large ones will be placed last, thus in the middle. Figure 1b displays the effect of this strategy.

Using the above algorithm results in a convex polygon for each node of the tree. These polygons are nested according to the hierarchy provided as input. For all leaves we calculate the centroids of their polygons and place the corresponding nodes there. This layout is then used as an input for the next step, the creation of the mesh that will be used to model the map.

3.3 Mesh Generation

As a first step, a conforming Delaunay triangulation for the nodes and the boundary polygon is constructed. Before we show how to obtain the triangle mesh that models

Fig. 2. Example for mapping the cluster tree nodes (a) onto the triangle edges (b), the corresponding subdivision (c) and subsequent lifting of the mesh to 2.5 dimensions (d)

the terrain features, we give some insights into the relationship of triangles and clusters. Consider a triangle of the Delaunay triangulation like displayed in Figure 2b and the *lowest common ancestors (LCA)* of its nodes in the cluster tree. We observe the following properties:

- We can associate with each edge $e = (u, v)$ a cluster C_e with $C_e = \text{LCA}(u, v)$.
- At least two of the three associated clusters are equal.

The first property is easy to see. Each vertex in the triangulation corresponds to a leaf in the cluster tree or is an additional point inserted for the conforming Delaunay triangulation. In the latter case, the vertex is part of the boundary polygon and we associate it with the root. Thus the lowest common ancestor is a cluster node.

The second property is based on the fact that for three nodes u , v , and w , the LCA of two pairs of three possible pairs must be the same. When considering the tree in Figure 2a, it becomes clear that without loss of generality

$$\text{LCA}(\text{LCA}(u, v), w) = \text{LCA}(u, w) = \text{LCA}(v, w)$$

holds. The idea of associating clusters with edges makes the subdivision of the triangles straightforward. We split a triangle into four subtriangles by splitting each edge in the middle as shown in Figure 2c. With each newly inserted node, we assign the cluster associated with the split edge.

All nodes of the triangle mesh are now either cluster or graph nodes. These mesh nodes are now “lifted” by computing their *elevation levels* that are simply chosen as their cluster hierarchy levels (see Figure 2d). The result is that each cluster node forms a valley because it is an inner node of the tree, thus has less distance to the root. The graph nodes of the mesh form peaks surrounded by cluster nodes which are all lower.

This mesh will serve as a basis for both drawing the map and routing the edges. Before we describe the edge routing, we make some further improvements to the mesh. Since the shape of the elevation model might be a bit coarse, we further subdivide the triangles and move the newly inserted vertices closer to the corner points. For aesthetic reasons, wide valleys should have a flat bottom instead of a very light slope. This results in hills with a more uniform slope and the flat area in valleys increases.

Fig. 3. In (a) the extra edges connecting the corners with the opposite vertex are displayed. On the right (b) an example of a routing network with the subdivisions and shortcut edges is shown. Allowed edges and nodes are drawn black, forbidden edges and nodes are drawn gray.

The result of the procedure described above is a 2.5D triangle mesh that covers seamlessly the complete boundary polygon. A software rasterizer is then used to interpolate the elevation levels given at the corners of a triangle to obtain a raster based representation of the elevation map.

3.4 Edge Routing

Now we use the mesh as a routing network that depends on the terrain features. The idea is to apply a shortest path based edge routing algorithm that is aware of the clustering by using the terrain mesh as a routing network. We follow a very general framework that consists of two major steps, namely, the construction of the routing network and the computation of a shortest path for each edge to obtain the control points for the curve.

Recall from the mesh generation section that we have subdivided each triangle by inserting vertices that correspond to clusters in the tree. Each graph node is associated with one node in the routing network. However, for a cluster, there are usually many nodes in the routing network. After some experiments it turned out that the generated curves look more pleasant when we allow shortcuts. Extra edges connect the graph nodes with the opposite cluster node, allowing a direct connection without making unnecessary turns when using one of the adjacent cluster nodes. Figure 3a displays a triangle of the routing network with the shortcut edges.

The costs for the routing edges used by the shortest path algorithm are chosen as a mixture of the Euclidean distance in the plane and the elevation difference. The extra costs for the elevation difference makes the edge routing aware of the terrain features. For some constant c (we set $c = 1/4$ after some experiments) the distance function $d(e)$ for an edge $e = (u, v)$ is defined as:

$$d(e) = c \cdot (\text{level}(u) + \text{level}(v) - 2 \cdot \text{level}(\text{LCA}(u, v))) + (1 - c) \cdot \|p_u - p_v\|,$$

where the level and coordinates of a node are normalized.

Given an edge $e = (u, v)$, we are looking for a shortest path starting at the routing node that corresponds to u and ends at the routing node that corresponds to v . Only cluster nodes are allowed in-between, because we do not want an edge to be drawn

through another node. Notice that the number of nodes and edges of the routing network are linear in the number of graph nodes.

An example is given in Figure 3b where the path is used as control points for a curve representing the edge from a to e . The black edges are those that allowed for a shortest path. The gray edges are forbidden in order to avoid drawing through a node.

4 Performance and Implementation Details

All graph and tree map related code has been written in C++ using the *OGDF - Open Graph Drawing Framework* [21]. The conforming Delaunay triangulation for the mesh is created with CGAL [6]. For rendering, we wrote a custom rasterizer that transforms the 2.5 dimensional mesh to a *digital elevation model (DEM)* for use in *GIS - geographic information systems*. Furthermore, data required for drawing nodes and edges are converted into in a more GIS-friendly format. These are then used as input for *Mapnik* [19], a tool for drawing “real” maps and developing mapping applications. This enables us to use other tools from the GIS tool chain like the *GDAL - Geospatial Data Abstraction Library* [12] that is used for extracting contour lines from the DEM and to calculate the shading for the resulting image. The performance of our approach is

Table 1. Runtime in seconds for the three phases for instances of different size

Instance	$ V $	$ E $	$ C $	Clustering	Fat Polygon	Edge Routing
gdea_2011	819	1 880	1 366	3.02	0.14	0.09
netscience	1 589	2 742	2 363	0.1	0.23	0.11
composers	2 815	11 429	4 708	4 888.42	1.22	5.50
PGPgiantcompo	10 680	24 316	16 640	⁻¹	2.83	21.55
cond-mat	16 726	47 594	26 235	⁻¹	4.03	70.18

mainly governed by a few steps. Table 1 shows the runtime² of the clustering algorithm, the fat polygon partitioning and the edge routing on a few instances.

The used clustering algorithm is not practical for large dense graphs. The algorithm suggested by Brandes [5] takes time $O(|V||E|)$ for calculating the betweenness, and this results in a total running time of $O(|V||E|^2)$. For the fat polygon partitioning, to our knowledge, no bound is known. But the time required suggests that it scales well, compared to the edge routing, even for larger instances with 16 726 nodes. The running time for the edge routing can be bounded by $O(|V|^2 \log |V|)$ when running a single source shortest path query for all edges incident to one node at the same time. For the instances presented here, this is acceptable. But with increasing graph size, the algorithm takes much more time and leaves room for improvement.

5 Results

In Figure 4, the co-authorship in the *Proceedings of the International Symposium on Automatic Graph Drawing* is visualized. The network is taken from the *Graph*

¹ Instances have been clustered in parallel on a different machine due to the input size.

² Machine with Core i7 2.7 GHz and 8 GB RAM.

Fig. 4. Graph drawing collaboration graph [13] with 819 nodes and 1880 edges

Drawing E-Print Archive [13] and covers articles from 1995 to 2011. The graph has been constructed by creating edges between all authors of an article with an edge weight of $1/(\text{number of authors} - 1)$. The map displays the different communities in graph drawing. Usually a mountain or hill contains one or two authors who have published not only many papers, but also for a long time and can be considered backbones of the research area. We have developed a prototype web interface for this network to offer a convenient way to explore authors, publications and their relations [13]. The graph in Figure 5 displays the co-authorship of scientists working in network science. Unlike the previous instance, this graph contains, besides one big connected component, a few medium sized components. The effect of the modified tree transformation on their shape is clearly visible.

Fig. 5. Network science collaboration graph [20] with 589 nodes and 2742 edges

6 Conclusion and Future Work

In this paper we have presented our approach for drawing clustered graphs as topographic maps. Besides the use of fat polygon partitioning, we showed how to extract a mesh from the layout that models the terrain features based on the clustering. As mentioned in the previous section, the clustering algorithm should be replaced by a more scalable approach for large instances. Furthermore, we believe that the visual appearance can be improved by modifying the treemap approach so that it produces more natural looking shapes. In addition, the placement by fat polygon partitioning is not directly aware of the adjacencies induced by the edges. As a result, the placement relies heavily on the clustering method used and does not take edge length into account.

While the proposed method works well for the applications in Section 5, it may fail when the edges and the clustering are largely unrelated. Such an example is discussed in [15], and it is demonstrated how the degrees of freedom of the partitioning algorithm can be exploited with moderate extra effort in order to obtain pleasing maps due to increased “edge-awareness”. We are currently working on improving the edge routing, because the current version is impractical for larger/denser instances, and a more sophisticated way to display coincident edges in the valleys is needed.

References

1. Balzer, M., Deussen, O.: Level-of-detail visualization of clustered graph layouts. In: Asia-Pacific Symposium on Visualization, pp. 133–140 (2007)
2. Balzer, M., Deussen, O., Lewerentz, C.: Voronoi treemaps for the visualization of software metrics. In: Proceedings of the 2005 ACM Symposium on Software Visualization, pp. 165–172 (2005)
3. De Berg, M., Onak, K., Sidiropoulos, A.: Fat polygonal partitions with applications to visualization and embeddings. CoRR abs/1009.1866 (2010)
4. de Berg, M., Speckmann, B., van der Weele, V.: Treemaps with Bounded Aspect Ratio. In: Asano, T., Nakano, S.-i., Okamoto, Y., Watanabe, O. (eds.) ISAAC 2011. LNCS, vol. 7074, pp. 260–270. Springer, Heidelberg (2011)
5. Brandes, U.: On variants of shortest-path betweenness centrality and their generic computation. *Social Networks* 30(2), 136–145 (2008)
6. CGAL - Computational Geometry Algorithms Library, <http://www.cgal.org/>
7. Cortese, P.F., Battista, G.D., Moneta, A., Patrignani, M., Pizzonia, M.: Topographic visualization of prefix propagation in the internet. *IEEE Trans. Vis. Comput. Graph.* 12(5), 725–732 (2006)
8. Di Battista, G., Eades, P., Tamassia, R., Tollis, I.G.: *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice Hall (1999)
9. Fabrikant, S.I., Montello, D.R., Mark, D.M.: The natural landscape metaphor in information visualization: The role of commonsense geomorphology. *J. Am. Soc. Inf. Sci. Technol.* 61(2), 253–270 (2010)
10. Gansner, E.R., Hu, Y., Kobourov, S.G.: GMap: Visualizing graphs and clusters as maps. In: *PacificVis*, pp. 201–208. IEEE (2010)
11. Garland, M., Kumar, G.: Visual exploration of complex time-varying graphs. *IEEE Transactions on Visualization and Computer Graphics* 12, 805–812 (2006)
12. GDAL - Geospatial Data Abstraction Library, <http://gdal.org/>
13. GDEA - Graph Drawing E-Print Archive, <http://gdea.informatik.uni-koeln.de/>
14. Girvan, M., Newman, M.E.J.: Community structure in social and biological networks. *Proceedings of the National Academy of Sciences* 99, 7821–7826 (2002)
15. Gronemann, M., Jünger, M., Kriege, N., Mutzel, P.: MolMap: Visualizing molecule libraries as topographic maps. Tech. rep (2012)
16. Holten, D.: Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE Transactions on Visualization and Computer Graphics* 12(5), 741–748 (2006)
17. Johnson, B., Shneiderman, B.: Tree-maps: a space-filling approach to the visualization of hierarchical information structures. In: *Proceedings of the 2nd Conference on Visualization* 1991, pp. 284–291 (1991)
18. Lambert, A., Bourqui, R., Auber, D.: Winding roads: Routing edges into bundles. *Computer Graphics Forum* 29(3), 853–862 (2010)

19. Mapnik, <http://mapnik.org/>
20. Newman, M.E.J.: Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E* 74, 036104 (2006)
21. OGDF - Open Graph Drawing Framework, <http://www.ogdf.net/>
22. Wise, J., Thomas, J., Pennock, K., Lantry, D., Pottier, M., Schur, A., Crow, V.: Visualizing the non-visual: Spatial analysis and interaction with information from text documents. *IEEE Trans. Vis. Comput. Graph.*, 51–58 (1995)