

Interactive 3D Force-Directed Edge Bundling

Daniel Zielasko, Benjamin Weyers, Bernd Hentschel, and Torsten W. Kuhlen

Visual Computing Institute, RWTH Aachen University JARA – High-Performance Computing

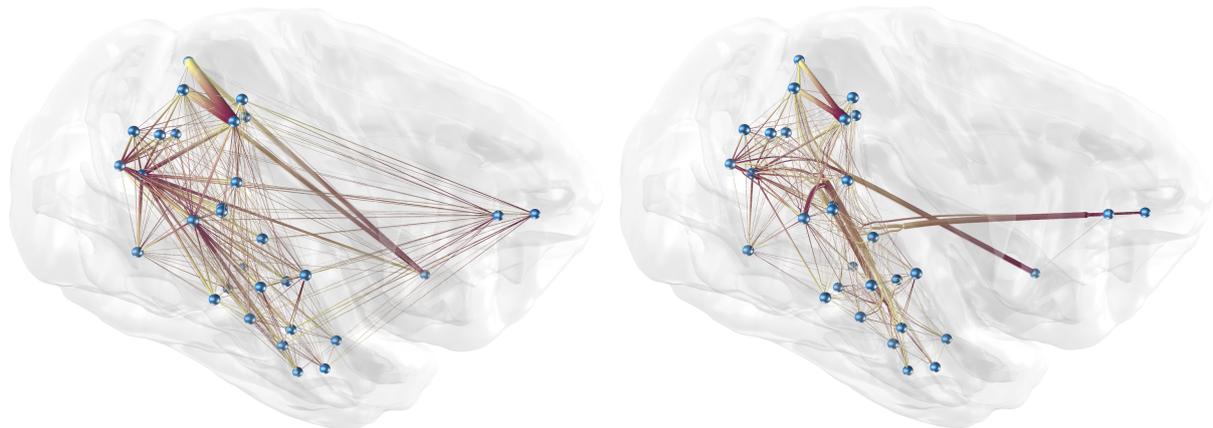


Figure 1: Node-link diagram of an almost fully connected, bidirectional graph, originating from a NEST simulation based on a macaque's brain [GD07]. These images depict 32 vertices each of which represents a brain region. The edges are the regions' interconnectivity. **Left:** original graph; **Right:** the same graph after edge bundling; the edges are directed from purple to yellow.

Abstract

Interactive analysis of 3D relational data is challenging. A common way of representing such data are node-link diagrams as they support analysts in achieving a mental model of the data. However, naïve 3D depictions of complex graphs tend to be visually cluttered, even more than in a 2D layout. This makes graph exploration and data analysis less efficient. This problem can be addressed by edge bundling. We introduce a 3D cluster-based edge bundling algorithm that is inspired by the force-directed edge bundling (FDEB) algorithm [HvW09b] and fulfills the requirements to be embedded in an interactive framework for spatial data analysis. It is parallelized and scales with the size of the graph regarding the runtime. Furthermore, it maintains the edge's model and thus supports rendering the graph in different structural styles. We demonstrate this with a graph originating from a simulation of the function of a macaque brain.

Categories and Subject Descriptors (according to ACM CCS): Data Structures [E.1]: Graphs and networks—
Computing Methodologies [I.3.8]: Computer Graphics—Applications

1. Introduction

Graphs are a ubiquitous type of data structure, which describes relational data and is often visually inspected in its representation as a node-link diagram. Ware and Mitchell showed [WM08] that when enriched with the appropriate depth cues, such as provided by immersive virtual environments, graphs laid out in 3D, can support data analysis. This is especially true for graphs with a natural spatial embedding, for example brain

region connectivity data [ABH*13, BSL13, BSJ*14], where a reduction to 2D results in a loss of information. However, when laid out in 3D, graphs are prone to visual clutter, even more than in a 2D layout. This makes graph exploration and data analysis less efficient. Aside from methods that change the positioning of vertices, like in [FR91], this problem can be addressed by methods changing the course of the edges. The latter can be classified into methods that perform local changes,

such as EdgeLense [WCG03] or Edge Plucking [WC07], and global methods such as edge bundling [Hol06, PX05]. Edge bundling is a method that combines geometrically close edges into bundles, which use much less screen space. This work introduces a 3D cluster-based edge bundling algorithm that is inspired by the force-directed edge bundling (FDEB) algorithm [HvW09b] and fulfills the requirements to be embedded in an interactive and immersive framework for spatial data analysis.

Interactivity, on the one hand, imposes the need of keeping the system's response time within 100ms [Nie93], while navigating and interacting with the visual representation of the graph and updating it at 30 frames per second or even more frequently when used with current 3D image projection devices. On the other hand, it requires algorithms to be fast enough so that they can be added to the workflow of an analyst. The requirement due to runtime is hard to define, since this is affected by several conditions. A duration in the order of a few seconds is acceptable [Nie93], until it has to be performed frequently.

Furthermore, the algorithm should be applicable to general graphs. This enables the possibility to be embedded into a general analysis framework. Moreover, special graphs like weighted and directed ones should be supported without restrictions. For this purpose, the presented approach is meant to maintain the graph's model, or rather an explicit geometry of it. On the one hand this allows applying the same visual representation to the graph when unbundled and on the other hand allows to perform semantic reasoning or standard graph analysis on the base of the bundled and thus changed graph. Additionally the inspection of a 3D graph on a 2D display can be difficult. This makes the use of immersive display technologies as CAVEs or HMDs promising. Here the explicit geometry again is useful as it is trivial to integrate into immersive virtual environments that then are accessible by those display technologies.

Our main contributions are a native 3D edge bundling algorithm that is optimized for interactive data analysis. Furthermore, it offers flexible rendering styles based on an explicit bundle topology. Finally, the algorithm's runtime is resolution independent and should in general scale with the graph size.

A preliminary work was published in the poster abstracts of IEEE InfoVis 2014 [ZWHK14].

2. Related Work

The representation of relational data as node-link diagram is prone to visual clutter. Holten et al. [Hol06, HvW09b] introduced the approach of merging geometrically close edges to bundles with the objective of reducing edge clutter for general graphs. For this purpose, their force-directed edge bundling algorithm (FDEB) attaches spring and electrostatic forces to segmentation points of the edges that then attract each other. This algorithm works very well for smaller graphs and is generalizable to 3D, but does not scale [Bon00], as it has a quadratic runtime in the number of edges. Selassie et al. [SHH11]

extend this algorithm to work better on directed graphs since antiparallel bundles are explicitly laid side by side.

Another class of edge bundling approaches implicitly bundles edges together by routing them through nearby, static control points [CZQ*08, LBA10b, PNBH11]. These points are obtained by a regular grid or a mesh geometry, which is generated from the graph's structure. These approaches avoid the costly edge-to-edge comparisons but require an appropriate geometry generation. If the density of points is too low, edges are sharply bent, or in the opposite case, are not bundled at all. Additional challenges arise with the extension to 3D. For example, there are a lot more possible routes for the edges to take and it is more difficult to aid them in taking a common one. However, for a limited three-dimensional case, i.e., with the vertices and edges bound to the surface of a sphere, Lambert et al. [LBA10a] showed that it is possible, in general.

A number of recent approaches take advantage of the pixel pipelines of today's GPUs [EHP*11, HET12, TE10]. This enables massively parallel processing in the pixel space. These image-based techniques are to some extent similar to the ones discussed before, as the result of the GPU accelerated calculations, e.g., a density field, can be described as a set of control points, even if equipped with a weight and an expansion to 2D. These approaches are very fast, but they are not naively adaptable to 3D, since when the view is changed, the complete edge bundling has to be recomputed. Depending on the size of the graph, this still could be possible in real time. However, without major changes in the algorithms, the resulting bundling depends on the view, which would be very confusing for the user, as it would be hard to obtain a consistent mental model of the graph. Simultaneously, view changes are very common in interactive applications. As mentioned in the introduction, graphs laid out in 3D support data analysis [WM08], such as provided by immersive virtual environments. But in projection setups like CAVEs [CNSD*92] the user usually does not look orthogonally at the projection screens, or even worse, simultaneously looks in different angles on different projection screens. This would lead to artifacts at the borders when using image based techniques. Another possibility to adapt the idea of pixel based methods to 3D is to transform them to the voxel space. This neutralizes a bunch of the mentioned issues, but simultaneously discards the main advantage of pixel based algorithms, namely, that for the full resolution not all voxels can be computed in parallel. In summary, it is very difficult to handle image space techniques for edge bundling in an interactive 3D application.

Böttger et al. [BSL13, BSJ*14] were the first to construct an algorithm explicitly adapted to full 3D. It combines methods from FDEB and kernel density estimation edge bundling (KDEEB) [HET12]. They compute all pairwise edge compatibilities and then move an increasing count of support points to a weighted mean of the surrounding points. However, the authors stated that their approach is not able to draw weighted graph bundles (e.g. by increasing line thickness), "... because

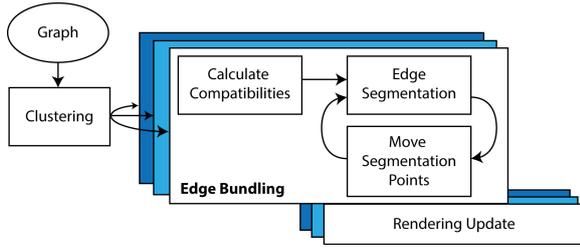


Figure 2: Edge bundling pipeline. The edges are clustered and then, per cluster task-parallel, bundled and drawn.

inclusion of all weighted connections would not be feasible due to current memory and computational limitations.” [BSJ*14]

In contrast to these approaches our keeps the graph’s model and thus there is no increase of memory consumption in general for any preferred bundle/edge representation, which includes different thicknesses for different weights.

Gansner et al. designed Multilevel Agglomerative Edge Bundling for Visualizing Large Graphs (MINGLE) consisting of two components [GHNS11]. First, they reduced the complexity of the problem by clustering the edges and building a proximity graph. Second, the decision to combine a pair of neighbored edges is made by calculating if there is a saving of “ink” in this case. They use the quantity of ink as a metaphor for the used pixel space and this again is a metric for display clutter. Our algorithm also follows a two-stage approach. But in the first step, it is based on the combination of a more drastic separation of edge clusters that are calculated in advance, and in the second step, we use a variation of the original FDEB algorithm for the edge bundling instead. This produces smoother edges, which is important as you want the user to be able to track the shape of connections without being distracted by sharp bends fetching the attention.

3. Method

In this paper, we define a Graph G as an ordered tuple (V, E) , where $V \in \mathbb{R}^3$ is a finite set of vertices and $E \subseteq V \times V$ a set of edges. Without loss of generality, we assume that an edge e is an ordered tuple (u, v) , with $u, v \in V$, as we can order u and v in case of an undirected graph, e.g., by the order of their occurrence in the underlying data structure. So, in case of a directed as well as of an undirected graph, we can identify an edge’s origin, destination, and direction. Thus, for a given $e = (u, v)$, we define $\vec{e} = v - u$. For the purpose of drawing, an edge is described as a set of points $e = \{e^{(i)}\}_{0 \leq i \leq k}$, with $k \in \mathbb{N}$. Furthermore let $n = |V|$ and $m = |E|$.

To avoid the complete edge-to-edge comparisons and reduce the runtime complexity, we take advantage of the fact, already stated by Holten et al. [HvW09b], that only very few edges really influence the positioning of any given one. Thus, we break down the edge bundling into a two-step process (see Figure 2). In Section 3.1 we describe how the edge population

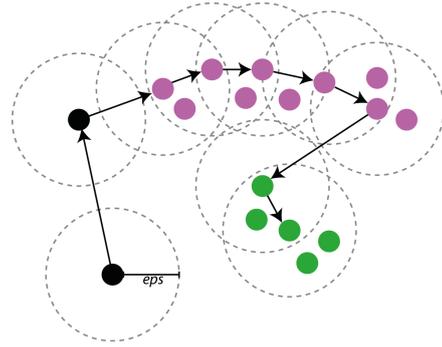


Figure 3: Example of a DBSCAN cluster calculation [EKSX96], with minimal cluster size is 2. The dashed circles, each framing one data point, depict the destiny parameter ϵ . Whenever a data point falls within the radius of any existing point of the current cluster, it is added recursively. The result is a green and a purple cluster. The two black data points are marked as unclustered.

is divided into clusters, without using a full pairwise comparison approach. From these edge clusters, subgraphs are created and a modified FDEB algorithm is calculated in parallel over these, which is described in Section 3.2. While these calculations are still running, the rendering takes place to provide the user with intermediate results (see Section 3.3). Here, various rendering styles are offered to the user. For example, a way of reasonably drawing weighted edges, or bundles, is presented.

3.1. Edge Clustering

Most of the runtime of FDEB originates from the calculations of the compatibility for every pair of edges and of the force every single edge exerts on the current one. Notably, the compatibility and consequently the forces for most of the pairs are in most cases almost zero. Therefore, we cluster edges with high compatibility in advance and only do further calculations within these clusters. To measure the edge similarity, we define some basic edge metrics. These metrics are inspired by the ones used to calculate the pairwise edge compatibility in FDEB. This first is the edge length corresponding to the scale from one edge to the other

$$\|v - u\| \in \mathbb{R},$$

the edge’s gradient corresponding to the angle between two edges in FDEB

$$\left(\frac{(v_i - u_i)}{\|v - u\|} \right)_{1 \leq i \leq 3} \in \mathbb{R}^3, \text{ with } v = (v_1, v_2, v_3)^T$$

and the edge’s position

$$\frac{v + u}{2} \in \mathbb{R}^3.$$

These metrics are composed of very basic geometric measures and can be complemented or replaced by more application-specific ones, e.g., metrics describing graph topology or edge

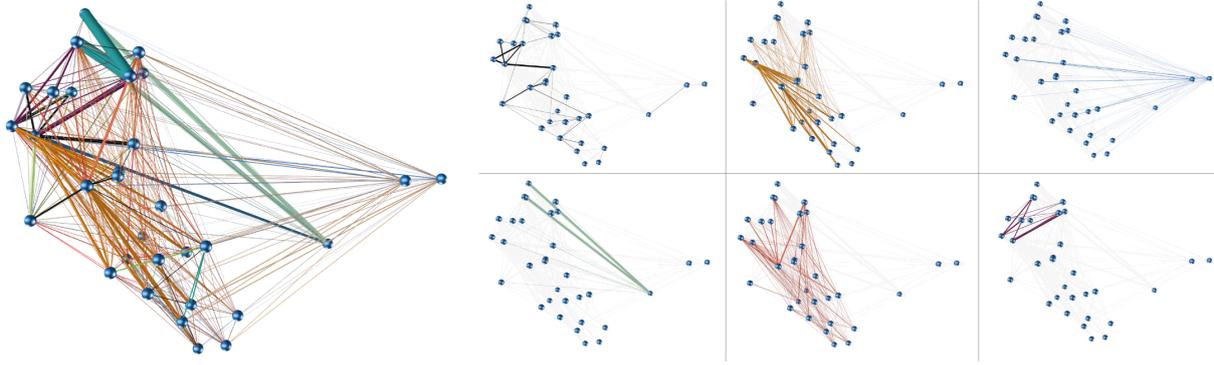


Figure 4: Graph from Figure 1 shown with color-coded clusters consisting of similar edges; black edges are unclustered, i.e., not similar to any other. **Left:** Complete graph; **Right:** 6 example clusters, where the upper left depicts not a cluster, but all unclustered edges.

weights if it is desired that bundled edges share, e.g., common origin/destination vertices or only similar-weighted edges should be laid together. It is important to note that the chosen metrics should not refer to pairs of edges but instead are computable per edge. Otherwise the quadratic complexity of FDEB just would be moved to a preprocessing step. In most cases they do not have the same expressiveness as their compatibility counterparts, but it is completely sufficient to find a preselection. For example, the angle between two edges is more meaningful than comparing their gradient. This means that the found clusters do not have to be perfect, but a superset of similar edges, as the connected edge bundling algorithm will do the precise work in the following.

This basic set of the three metrics finally specifies a seven-dimensional feature vector for every edge. After normalizing every component to the interval $[0, 1]$, these vectors are inserted into an R^* -tree [BKSS90], which enables an efficient access to this multi-dimensional data with spatial queries. Compared with an R -tree, an R^* -tree guarantees a better aligned indexing within a longer setup time but faster access during runtime. For the clustering we use DBSCAN [EKX96], a density-based clustering algorithm, whose parameter for the minimal cluster size is set to a value of 2. This is because we want to bundle sets of similar edges starting with a size of 2. The number of clusters, or the similarity of edges, necessary to form a cluster is determined by the density parameter eps . It is a threshold for deciding if two data points, according to their Euclidean distance in the parameter space, are close enough to be assigned to a common cluster (see Figure 3). A default eps value is chosen by first precomputing an interval of reasonable values $[l, u]$, starting with a lower bound where most of the edges first become assigned to any cluster and an upper bound where the result is only one cluster. Finally the value is set to $l + (u - l) / 3$ as this often has turned out to be a good choice in our experiments, which means there is no loss of bundling quality while a good runtime performance is achieved. However, always computing the “right” density parameter for the clustering algorithm is not feasible. On the one hand, it could be very different for two graphs and on the other, a good

parameter selection strongly depends on the current analysis task. But since the clustering is fast, it allows us to put the user in the loop of interactively choosing or changing the value on demand within the precalculated interval (see Section 5). After selecting an eps value, a color-coded visual representation of the resulting clusters is instantly shown (cf. Figure 4).

Out of the resulting clusters we create new subgraphs, whereas from the implementation view they only hold references to the edges of the original one, such that changes in the subgraphs propagate to the root graph. Note that the subgraphs are not necessarily connected, even if the parent graph was. They are processed by a modified FDEB algorithm, described in the next section, below.

3.2. Edge Bundling

Our bundling strategy is based on the force directed edge bundling algorithm described by Holten and van Wijk [HvW09b]. This approach merges geometrically similar edges to common bundles. For this purpose, it attaches spring and electrostatic forces to segmentation points upon the edges that mutually attract each other (see Fig. 5). We made major changes and extensions to the algorithm to achieve especially three goals:

(G1) Reducing parameters

Of course it could be asked why especially an approach with the aim to be interactive wants to reduce parameters, as user involvement with direct feedback is a good method to find suitable parameters. However, it should be kept in mind that edge bundling is a tool for supporting data analysis and the user should need as little time as possible to tweak algorithm parameters.

(G2) More stable force system

The simulation of forces is often very fragile, e.g., to numerical instability. We want to clear out the problems we identified.

(G3) Increase computational efficiency

This goal addresses the desired interactivity of the edge bundling.

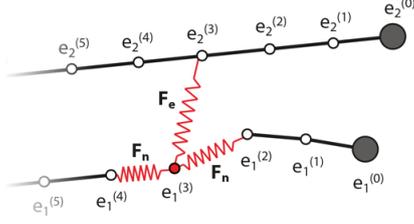


Figure 5: Segmented edges e_1 and e_2 . Example force calculation for the 3rd segmentation point of e_1 , assumed there is only one other relevant edge e_2 : $F(e_1^{(3)}) = F_e(e_1^{(3)}) + F_n(e_1^{(3)})$. All other points are static. Figure based on [HvW09b].

In the initialization step, the compatibilities for all edge pairs are calculated with more precise metrics than within the clustering and stored (**G3**, cf. Algorithm 1, line 1). They are needed later, as it depends on the compatibility of two edges to which extent they attract each other. It is possible to compute them now, as the compatibility measurements only depend on the origin and destination point of the edges, as described below, which are assumed not to change during the algorithm. If it is required to change single vertex positions during the simulation, a first naïve solution to this would be to restart the bundling for the affected clusters. If the changes provoke changes even in the clustering, the computations have to be restarted.

The compatibility criteria starts with the angle between two edges $C_a(e_1, e_2)$, with

$$C_a(e_1, e_2) = \left| \frac{\vec{e}_1 \cdot \vec{e}_2}{\|\vec{e}_1\| \cdot \|\vec{e}_2\|} \right| \in [0, 1].$$

The second one describes the scale between two edges $C_s(\vec{e}_1, \vec{e}_2)$ with,

$$C_s(e_1, e_2) = \frac{\min(\|\vec{e}_1\|, \|\vec{e}_2\|)}{\max(\|\vec{e}_1\|, \|\vec{e}_2\|)} \in (0, 1].$$

At last, we rate the positioning to each other $C_p(\vec{e}_1, \vec{e}_2)$ with,

$$C_p(e_1, e_2) = \frac{\frac{\|\vec{e}_1\| + \|\vec{e}_2\|}{2}}{\frac{\|\vec{e}_1\| + \|\vec{e}_2\|}{2} + \|\text{mid}(e_2) - \text{mid}(e_1)\|} \in [0, 1],$$

$$\text{with } \text{mid}(e) = \frac{e^{(0)} + e^{(k)}}{2}.$$

We intentionally leave out the so called visibility criterion, described by Holten et al. [HvW09b], as we argue that it is only a weighted combination of the other criteria. The reason for adding this fourth metric are edges that are comparable with respect to the other metrics, but not comparable in the intended sense. As an example, the opposite edges of a skewed parallelogram are mentioned. However, in our opinion this already is covered by the position metric.

Now, the total compatibility $C(e_1, e_2)$ is calculated as follows:

$$C(e_1, e_2) = C_a(e_1, e_2) \cdot C_s(e_1, e_2) \cdot C_p(e_1, e_2) \in [0, 1]$$

Algorithm 1: Edge Bundling

Data: $G = (\text{Vertices}, \text{Edges})$,
 numOfCycles,
 maxDisp // maximum Displacement

```

1 comp = CalcCompatibilityMap(Edges);
2 for numOfCycles do
3   DoEdgeSegmentation(Edges);
4   while newEnergy < oldEnergy do
5     oldEnergy = newEnergy;
6     newEnergy = 0;
7     foreach Edge e do
8       foreach Edge e_x do
9         foreach segmentation point i do
10          // e^{(i)} is the ith segmentation point of e
11          force = CalcForces(e^{(i)}, e_x^{(i)}) \cdot comp[e, e_x];
12          e^{(i)} = e^{(i)} + \min(\text{maxDisp}, \text{force});
13          newEnergy = newEnergy + \text{force};
14        end
15      end
16    end
17  end
18 end
    
```

After the initialization, a number of cycles c are simulated (see Algorithm 1, line 2). At the beginning of any cycle, each edge is subdivided by adding new segmentation points, refining the resulting curves (line 3). With every cycle, the number of points is doubled by evenly distributing the new ones along the risen edge line segments and simultaneously dropping the old ones. Within these cycles, we avoid an explicit number of iterations but instead iterate until the total energy in the system does not further decrease. This eliminates a parameter (**G1**) and in addition assures that edges are not bundled with different densities (**G2**), a problem already described by Böttger et al. [BSL13]. Within every iteration, for every segmentation point of any edge, the total force it is affected by is calculated. This consists of two components (see Figure 5). First, the force F_n holding the segmentation point $e^{(i)}$ in between its neighbors, with spring constant k

$$F_n(e^{(i)}) = k \cdot (\|e^{(i-1)} - e^{(i)}\| + \|e^{(i)} - e^{(i+1)}\|)^2.$$

We changed the force from a linear to a quadratic behavior, which leads to more stable results (**G2**), as we want to better imitate the behavior of a physical spring and penalize the points in running away much harder. The second force component F_e is the sum over all forces to the corresponding segmentation points of the other edges

$$F_e(e^{(i)}) = \sum_{e_x \in E} (C(e, e_x) \cdot \|e_x^{(i)} - e^{(i)}\|)^2.$$

Here again we take a squared distance, for the same reasons as before. The total force on an edge segmentation point then is

$$F(e^{(i)}) = F_n(e^{(i)}) + F_e(e^{(i)}).$$

Finally, we move the given point along the force vector by the step size $s_c = s/2^c$, with c the number of the current cycle run.



Figure 6: Two bundles coming from the left become one bundle, with their widths stacked. (a) without an additional smoothing step. (b) with an additional smoothing.

As an additional difference to FDEB, we cap this value by the length of the total force vector (**G2**). This prevents an edge being attracted by another edge when they are not similar, as a point originally was moved along the combined force with a defined step size, as long as the force was not exactly 0. From the increasing number of segmentation points, Holten et al. also expect smoothed edges. We think that the bundling with respect to **G3** is too expensive for this task and therefore recommend to choose as few cycles as possible and as many as necessary, and to instead smooth the edges in the following rendering step (see Section 3.3).

It should be mentioned that we added the option to perform an ambiguity free variant of the algorithm as described in [LLCM12], where only edges with either a common origin or a common destination are bundled. This would guaranty that there is no loss of information in the bundled representation, but in general reduces less clutter.

3.3. Rendering

We included in our framework a graph view, which renders a graph and consists of a vertex view and an edge view. A single vertex view draws all vertices as spheres, due to the framework's flexibility, it would be possible to place different vertex representations here. Every edge cluster is drawn by an independent edge view, so less data has to be updated on the graphics card when one parallelized edge bundling task (see Section 3.4) has finished and triggers an update of the corresponding edge segments. In addition, keeping the separation helps to reduce computational costs of all following steps.

3.3.1. Edge Rendering Style

During all computations, the graph and the changing edge segments are redrawn every frame with a so-called edge rendering style (see Figure 7a). Thus, the user is able to interactively explore a largely cleaned graph, while the bundling is being processed. This is positively affected by the observation that the main structural changes, specifically the reduction of edge clutter, take place during the first cycles of the bundling algorithm, whereas the following ones mainly just refine the result.

The edge segmentation points are supplied to the OpenGL pipeline as line segments, where they are rendered as polylines using a geometry shader. In the case of a directed graph, the direction is encoded by the intensity, from dark-to-light according to a study of Holten et al. [HvW09a]. Although, a tapered representation is even higher rated, this would interfere with the edge weight coding of the other views (see Section 3.3.2) and for reasons of usability, we prefer to keep

Algorithm 2: Segmentation to explicit Bundles

```

Data:  $G = (\text{Vertices}, \text{Edges})$ ,
//  $e^{(i)}$  is the  $i$ th segmentation point of  $e$  and
//  $k = 2^{(\text{cycles} - 1)} - 1$ 
 $\forall$  Edges  $e: \exists_1 (e^{(0)}, \dots, e^{(k)})$ ,
 $\epsilon ps$  // density parameter
Result: Bundles // every bundle is a strip of segments
Topology // tracks for every edge the bundles it participates in

new segment = (Vector3D, Vector3D);
foreach Edge  $e$  do
  segment[0] =  $e^{(0)}$ ;
  segment[1] =  $e^{(1)}$ ;
   $R^*$ Tree.Add(segment);
end
end
for  $i = 1..k$  do
  Cluster = DoClustering( $R^*$ Tree,  $\epsilon ps$ );
   $R^*$ Tree.Clear();
  // consider unclustered segments as a set with size  $l$ 
  foreach set of segments  $\{s_0, \dots, s_r\}$  in Cluster do
    //  $\{e_0, \dots, e_r\}$  corresponding set of Edges
    foreach  $e \in \{e_0, \dots, e_r\}$  do
      UpdateTopology( $e$ );
      segment[0] = GetMeanSegEndPoint( $\{s_0, \dots, s_r\}$ );
      segment[1] =  $e^{(i+1)}$ ;
       $R^*$ Tree.Add(segment);
    end
    if  $\{e_0, \dots, e_r\}$  were a cluster before then
      | ExtendBundle( $\{s_0, \dots, s_r\}$ );
    else
      | AddBundle( $\{s_0, \dots, s_r\}$ );
    end
  end
end

```

it consistent. All together, the chosen representations are just examples and the framework allows to change the representations and even interactively switch between a set of them, further demonstrated in the following (also see Section 5).

3.3.2. Bundle Rendering Styles

The result of most edge bundling approaches are lines, or in our case tubes, laying over each other. However, as bundles are only implicit, this means there is neither explicit knowledge about which edges become part of a bundle, nor when or where this happens. However, this is not sufficient for all applications. For instance, the user could be interested in the number of edges forming a bundle, but usually this information is lost. Of course, there are work-arounds, by, for example, drawing edges with a stacking alpha [ZBDS12] or just beside each other [PNBH11]. But as work-arounds they all come with their limitations, as bundles stay implicit and the representation is not exchangeable. Another prominent example are weighted edges. To overcome this issue, we extract the explicit bundling topology based on the segmented edges. For this purpose, we have to decide when sets of edges form bundles and when not. As a result of the quasi-continuous physical simulation, the edge segmentation points do not necessarily lie precisely on top of each other. However, considering single points is usually not a good idea, as co-located points do not clearly imply partially co-located edges. The edges could cross each other in this point or could even be antiparallel.

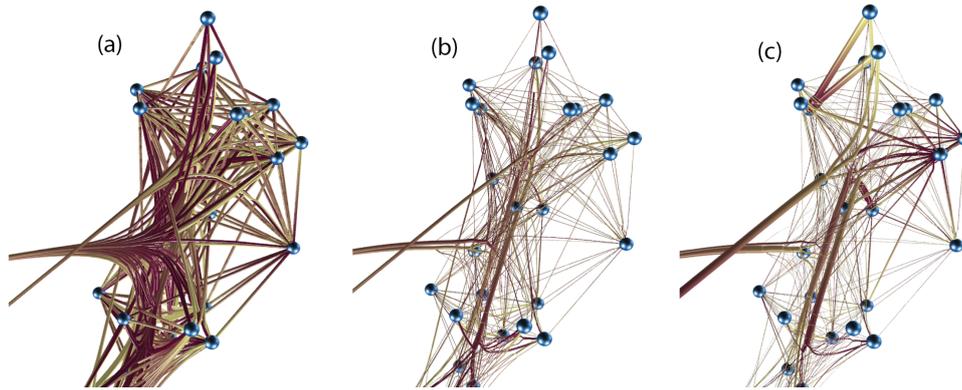


Figure 7: A section of the graph shown in Figure 1 depicting different rendering styles: (a) edge rendering style, (b) bundle rendering style, line width represents the number of edges combined in this bundle, (c) bundle rendering style, line width represents the combined edge weight in this bundle.

Therefore, we introduce an algorithm that identifies co-located segments, each consisting of a pair of segmentation points (see Algorithm 2). If two segments lie on top of each other, we can be sure that they should be part of a bundle. To find these segments, it is useful to use spatial clustering, so we use the DBSCAN algorithm [EKSX96]. Again, we have to decide on the eps parameter. For this purpose, we could evenly resample the edges and then calculate the eps value so that two directly successive segments are just not detected as lying upon each other. But as long as we do not resample with the length of the smallest segment, we could possibly lose existing smoothness. Therefore, we leave that step optional but recommend to use it when there are artifacts in the results. Hence, the algorithm by default just chooses an eps value close to the smallest segment. Nevertheless, because of two other conditions described in the following, we should not get too confused with the clustering. First, we are still processing the edge clusters independently, which reduces the number of simultaneously processed edges, but even more important, due to the clustering conditions all included edges are very similar and it is more difficult to get confused with just crossing segments, etc. Second, the previously performed bundling algorithm constrains that only corresponding segmentation points, i.e., with the same index, attract each other. This is favored by the fact that the segmentation points of implicitly bundled edges are very close together. Furthermore, this allows to systematically cluster the segments in groups from the origin of all edges to the destination (cf. Algorithm 2). The latter additionally enables the possibility to keep track of the edges' topology. Thus, the results are not only explicit bundles, but also do not lose the edge semantic, as it is still known that a specific edge consists of these bundles in that sequence.

During the process of replacing sets of segments by bundles, segmentation points have to be averaged to single points in a bundle and as everything needs to be kept connected, although previous points have to be shifted. This could cause unaesthetic bends as illustrated in Figure 6. Hence, we pipe the results

through a Laplacian smoothing. The resulting bundles finally are described by Catmull-Rom splines [LAM10] and supplied to the OpenGL pipeline, where they are rendered using a geometry shader. Due to the fact that we have calculated an explicit representation and kept the edge topology, we are now able to add various data to our representation. We implemented two examples, edge density and weight (see Figure 7). Beyond that, it would be even possible to represent time varying data.

In summary, the user has the opportunity to switch between different views depending on what she wants to analyze (see Section 5) or highlight.

3.4. Parallelization

Up to now, the edge clustering was only used to partition the data before it is processed by a non-scaling algorithm. But as the available edge clusters (see Section 3.1) are independently processed by the edge bundling (see Section 3.2), it is possible to create a new task for every cluster (referring to task-based parallelization). Thus, the bundling is parallelized. The same applies for an update of the drawing (see Section 3.3). However, a large variance in cluster sizes causes load-balancing issues in static parallelization.

This leads to a second layer of parallelization within the bundling. Two outer *for* loops of the edge bundling algorithm, each iterating over all edges in a cluster, are parallelized, too. The first one is hidden in the edge segmentation function (cf. Algorithm 1, line 3) and does not need any adjustments as this happens completely independent and therefore no data races are possible. The second one is the outer one iterating over all edges within a cycle (cf. Algorithm 1, line 7). To solve data races here, we buffer the position updates and swap them outside of the loop. Finally, we synchronize the update of the total energy.

Both layers together parallelize the edge bundling and simultaneously solve possible load balancing issues.

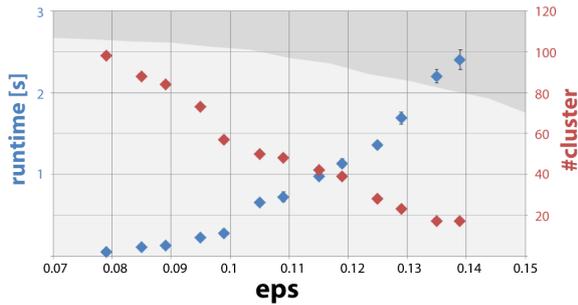


Figure 8: Dependency of the density parameter ϵ to runtime and number of clusters, within the useful range of ϵ values for this particular graph. This means from, "there is first time a real clustering" to, "clusters start getting so small that there is no more bundling worth mentioning". Additionally in the background the ratio of unclustered edges (dark) to clustered ones (bright) is depicted.

4. Results – Runtime

As proof of concept we bundled an example graph with only 32 vertices but almost 600 edges (see Figure 1), our algorithm took about 1.36 seconds (SD = ± 0.036 , in 10 measurements) with an ϵ value of 0.125 and 28 identified clusters. Thus, it was about 150 times faster than our basic 3D implementation of the FDEB algorithm, which took about 220 seconds. In this example that makes the difference in aiming for an interactive usage. Without any parallelization and the same parameters our algorithm took about 5.03 seconds (SD = ± 0.037 , in 10 measurements). Measurements were performed on an Intel Xeon E5540 2.53GHz quadcore-processor running Windows 7 with 12 GB of RAM and a GeForce GTX480 graphics card. The whole measurement series was performed on the graph with different density parameters and in relation to the number of resulting clusters, which is shown in Figure 8. As expected, the plot shows a decreasing runtime with an increasing number of clusters, although the distribution of cluster sizes was not examined here. All bundling results for the chosen ϵ interval are well-shaped and just differ in the degree of bundling.

Much more interesting than these raw numbers, as they heavily depend on the number and distribution of the clusters (see Fig. 8) that again are chosen by the user, is the question concerning runtime scalability [Bon00] regarding graph size. Following the processing pipeline (see Figure 2), first, setting up the R*-tree and clustering all edges costs $O(m \log m)$, with m number of the edges. Unfortunately, due to the changed termination condition within an edge bundling cycle (see Section 3.2), we are no longer able to guarantee any meaningful theoretical runtime complexity. This is because it is theoretically possible that there is an epsilon-small decrease in the total energy with every iteration so the algorithm does not terminate. This is not an issue in practice as the algorithm stops when the total energy has stayed the same or has increased in comparison to the last step of the last cycle. Due to numerical reasons the total energy cannot infinitely

decrease without getting zero. Nevertheless, to get an idea of the runtime complexity we could assume a constant number of iterations or a threshold. This would lead to a quadratic runtime with respect to the number of edges in a cluster. When considered together, we get an overall complexity of $O(|C_{max}|^2)$, with C_{max} the largest edge cluster, which in the worst case again is $O(m^2)$. Based on the runtime analysis, the proposed algorithm does not scale in the worst case.

But indeed, with examples that are not specially designed, the algorithm behaves very well. This is because $|C_{max}|$ does not necessarily increase with the size of the input as many graphs and networks in real world applications are sparse and have a low connectivity. Thus, with increasing size, graphs can be broken down to an increasing number of components. Edges in different components are very likely not similar, regarding the used metrics (see Section 3.1 and 3.2), and therefore would be assigned to different clusters.

In conclusion, graphs with the mentioned properties will not generate larger edge clusters and therefore support a scaling of the proposed algorithm. Of course, these considerations do not hold for all graphs. In Figure 4 a nearly fully connected, bi-directed graph is shown. With every vertex added to such a graph, the mean size of the clusters will grow, so that the runtime is quadratic, as within the clusters we still operate for every edge on every other one. Nevertheless, the clustering still reduces the problem, as observable in Figure 4.

5. Application

The need for the presented edge bundling technique results from applications like the one described in the following. VisNEST [NSA*13] visualizes various data originated from a NEST simulation [GD07], i.e. simulations of large scale neuronal networks, in immersive environments. For this purpose there are different views on the data that are linked. One of these is the connectivity view, which "shows how each brain area is connected to others and how it may spread activity along its pathways" [NSA*13]. The vertices have static positions, each in the center of the brain region they represent. The application is designed to run in a CAVE-like environment, but it is possible to use it with other VR devices as well as on a desktop PC. A basic set of interaction techniques is given by the application context. First, there is basic navigation and selection. Second, there are some options to directly manipulate the graph, as limiting the depicted edges by their weight, change colors, or move vertices.

We added the possibility to initiate a bundling of the current graph (see Figure 1 and 7). In addition, it is possible to enable a bundling mode, where with every change in the graph structure, the bundling is automatically refreshed. This method and all options are accessible via an extended pie menu [GPL*13]. The usual work flow starts with the clustering of the graph with the precalculated ϵ value (see Section 3.1). The result is color coded as depicted in Figure 4. The colors are evenly distributed within the CIE L*a*b* color space, which is

perceptually uniform [Hun58, McL76]. However, currently we are not optimizing the cluster coloring with respect to the clusters' distance to each other. Altogether, the visualization is not sufficient to analyze the clusters in detail, but it is for obtaining a quick overview of whether the calculated clusters are meaningful or not. If this is not the case, the user is able to change the *eps* value via a slider within the precalculated interval (see Section 3.1) and the cluster visualization will be updated accordingly. When satisfied with the clustering, the user starts the edge bundling and simultaneously gets feedback, as the updated positions of the edge's segmentation points are directly drawn. The user usually starts inspecting the graph by rotating and translating the view point, while the bundling is finished. Now it is possible to switch between the rendering styles (see Section 3.3). Depending on the active style, additional parameters can be adjusted, as for example the minimal, maximal, or both line widths relating to an edge parameter. Finally, it is possible to continuously fade between the Edge Rendering Style (see Section 3.3.1) and the unbundled graph, which can be interesting for finding just the right ratio between clutter reduction and information loss through edge bundling.

6. Discussion & Future Work

In the following, single aspects of the presented approach are discussed in sequence of their appearance in the pipeline, according to Figure 2, and start with the clustering of the edges.

For the calculation of edge clusters, we have used a density-based approach, which expects a density parameter. To support the user in choosing the right parameter, we precalculate an interval of meaningful parameters. To avoid patronizing the user, we used a very conservative method for this purpose, starting with a lower bound where most of the edges first become assigned to any cluster and an upper bound where the result is only one cluster. This can lead to the situation that still large parts of the interval are not of interest, e.g., because there is only one big cluster for larger parts of the interval. It may be sufficient to be more restrictive for this part and further analyze the clustering behavior in the future. This and other user oriented questions have to be addressed by a subsequent user study, but are out of the scope of this work.

Then, we took the calculated edge clusters and started an optimized, force-directed edge bundling algorithm on each cluster independently and in parallel.

Furthermore, we described that the clusters are drawn independently. This can be an issue with directed graphs and antiparallel bundles, as it is possible that they are drawn at the same position and so cover each other. Although we could change the clustering metrics in a way that these edges fall in one common cluster, there is no applicable solution yet to place those bundles relative to each other in 3D. In a general 2D case, Selassie et al. [SHH11] placed bundles side by side with respect to a highway metaphor. Unfortunately, this metaphor does not work in a three dimensional space, as there is no left and right anymore.

We embedded the proposed method in a data analysis tool and added user interaction, such as the possibility to fade in and out between the bundled and unbundled graph. Some of this interaction could be even more beneficial if it would be possible to perform them only on parts of the graph. These parts can often be expected to semantically correspond to the calculated clusters. The only missing component for a realization is finding a working selection metaphor for the clusters, which is not straightforward as the clusters are usually interwoven or hidden. A first idea to achieve this is to combine a standard selection strategy, with the possibility to select a specific cluster from a list of small multiples, as depicted in Figure 4 (right). This additionally would support the pure cluster visualization.

We argued that the proposed algorithm should scale for most graphs appearing in data analysis. This statement in future could be supported by an empirical user study. Nevertheless, this represents a challenge due to the sample size. On the one hand, graphs would have to be synthetically created, and on the other, they need to be meaningfully distributed with respect to their properties.

Finally, future work includes investigating the relevance of edge bundling as a misleading factor in inferring paths or structures in the data that are actually not there. For example, axon paths in the brain do not have to follow the paths taken by the bundles, but can be interesting for neuroscientists and for this reason are noticed from another view point as for a specialist from a different domain. Of course, this is not only a risk, but an opportunity too. If, for instance, the real pathways are known, the bundles could be attracted by these.

7. Conclusion

Interactive analysis of 3D relational data is challenging. A common way of representing such data are node-link diagrams as they support analysts in achieving a mental model of the data. However, naive 3D depictions of complex graphs tend to be visually cluttered. This makes graph exploration and data analysis less efficient. This problem can be addressed by edge bundling, which combines geometrically close edges into bundles. We have presented a native 3D, edge cluster-based and parallel edge bundling algorithm that fulfills the requirements necessary to be embedded in an interactive framework for spatial data analysis. Furthermore, it maintains the topology of the edge bundles and thus supports rendering of the graph in different structural styles. Finally, the proposed algorithm scales in runtime with the number of edges for most of the graphs.

Acknowledgments

The authors would like to acknowledge the support by the Helmholtz portfolio theme "Supercomputing and Modeling for the Human Brain". The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 604102 (HBP).

References

- [ABH*13] ALPER B., BACH B., HENRY RICHE N., ISENBERG T., FEKETE J.-D.: Weighted Graph Comparison Techniques for Brain Connectivity Analysis. *Proc. Conference on Human Factors in Computing Systems* (2013), 483–492. 1
- [BKSS90] BECKMANN N., KRIEGEL H., SCHNEIDER R., SEEGER B.: The R*-Tree: An Efficient and Robust Access Method for Points and Rectangles. *Proc. ACM SIGMOD International Conference on Management of Data 19* (1990), 322–331. 4
- [Bon00] BONDI A. B.: Characteristics of Scalability and Their Impact on Performance. *Proc. ACM International Workshop on Software and Performance* (2000), 195–203. 2, 8
- [BSJ*14] BÖTTGER J., SCHURADE R., JAKOBSEN E., SCHÄFER A., MARGULIES D. S.: Connexel Visualization: A Software Implementation of Glyphs and Edge-Bundling for Dense Connectivity Data Using BrainGL. *Frontiers in Neuroscience* 8 (2014), 15. 1, 2, 3
- [BSL13] BÖTTGER J., SCHÄFER A., LOHMANN G.: Three-Dimensional Mean-Shift Edge Bundling for the Visualization of Functional Connectivity in the Brain. *IEEE TVCG* 20, 3 (2013), 471–480. 1, 2, 5
- [CNSD*92] CRUZ-NEIRA C., SANDIN D. J., DEFANTI T. A., KENYON R. V., HART J. C.: The CAVE: Audio Visual Experience Automatic Virtual Environment. *Communications of the ACM* 35, 6 (1992), 64–72. 2
- [CZQ*08] CUI W., ZHOU H., QU H., WONG P. C., LI X.: Geometry-Based Edge Clustering for Graph Visualization. *IEEE TVCG* 14, 6 (2008), 1277–84. 2
- [EHP*11] ERSOY O., HURTER C., PAULOVICH F. V., CANTAREIRO G., TELEA A.: Skeleton-Based Edge Bundling for Graph Visualization. *IEEE TVCG* 17, 12 (2011), 2364–73. 2
- [EKXS96] ESTER M., KRIEGEL H., SANDER J., XU X.: A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. *Proc. ACM Knowledge Discovery and Data Mining* (1996), 226–231. 3, 4, 7
- [FR91] FRUCHTERMAN T., REINGOLD E.: Graph Drawing by Force Directed Placement. *Software – Practice and Experience* 21 (1991), 1129–1164. 1
- [GD07] GEWALTIG M., DIESMANN M.: NEST (NEural Simulation Tool). *Scholarpedia* 2, 4 (2007), 1430. 1, 8
- [GHNS11] GANSNER E. R., HU Y., NORTH S., SCHEIDEGGER C.: Multilevel Agglomerative Edge Bundling for Visualizing Large Graphs. *Proc. IEEE Pacific Visualization Symposium* (2011), 187–194. 3
- [GPL*13] GEBHARDT S., PICK S., LEITHOLD F., HENTSCHEL B., KUHLEN T.: Extended Pie Menus for Immersive Virtual Environments. *IEEE TVCG* 19, 4 (2013), 644–51. 8
- [HET12] HURTER C., ERSOY O., TELEA A.: Graph Bundling by Kernel Density Estimation. *Computer Graphics Forum* 31, 3pt1 (2012), 865–874. 2
- [Hol06] HOLTEN D.: Hierarchical Edge Bundles: Visualization of Adjacency Relations in Hierarchical Data. *IEEE TVCG* 12, 5 (2006), 741–8. 2
- [Hun58] HUNTER R. S.: Photoelectric Color Difference Meter. *Journal of the Optical Society of America* 48, 12 (1958), 985. 9
- [HvW09a] HOLTEN D., VAN WIJK J. J.: A User Study on Visualizing Directed Edges in Graphs. *Proc. International Conference on Human Factors in Computing Systems* (2009), 2299–2308. 6
- [HvW09b] HOLTEN D., VAN WIJK J. J.: Force-Directed Edge Bundling for Graph Visualization. *Computer Graphics Forum* 28, 3 (2009), 983–990. 1, 2, 3, 4, 5
- [LAM10] LAMBERT A., AUBUR D., MÉLANÇON G.: Living Flows: Enhanced Exploration of Edge-Bundled Graphs Based on GPU-Intensive Edge Rendering. *Proc. International Conference on Information Visualization* (2010), 523–530. 7
- [LBA10a] LAMBERT A., BOURQUI R., AUBER D.: 3D Edge Bundling for Geographical Data Visualization. *Proc. International Conference Information Visualisation* (2010), 329–335. 2
- [LBA10b] LAMBERT A., BOURQUI R., AUBER D.: Winding Roads: Routing Edges into Bundles. *Computer Graphics Forum* 29, 3 (2010), 853–862. 2
- [LLCM12] LUO S.-J., LIU C.-L., CHEN B.-Y., MA K.-L.: Ambiguity-Free Edge-Bundling for Interactive Graph Visualization. *IEEE TVCG* 18, 5 (2012), 810–21. 6
- [McL76] MCLAREN K.: XIII-The Development of the CIE 1976 (L* a* B*) Uniform Colour Space and Colour-Difference Formula. *Journal of the Society of Dyers and Colourists* 92 (1976), 338–341. 9
- [Nie93] NIELSEN J.: *Usability Engineering*. Morgan Kaufmann Publishers Inc., 1993. 2
- [NSA*13] NOWKE C., SCHMIDT M., ALBADA S. J. V., EPPLER J. M., BAKKER R., DIESMANN M., HENTSCHEL B., KUHLEN T.: VisNEST – Interactive Analysis of Neural Activity Data. *Proc. IEEE Symposium on Biological Data Visualization* (2013), 65–72. 8
- [PNBH11] PUPYREV S., NACHMANSON L., BEREG S., HOLROYD A.: Edge Routing with Ordered Bundles. *Proc. International Conference on Graph Drawing* (2011), 136–147. 2, 6
- [PX05] PHAN D., XIAO L.: Flow Map Layout. *Proc. IEEE InfoVis* (2005), 219–224. 2
- [SHH11] SELASSIE D., HELLER B., HEER J.: Divided Edge Bundling for Directional Network Data. *IEEE TVCG* 17, 12 (2011), 2354–63. 2, 9
- [TE10] TELEA A., ERSOY O.: Image-Based Edge Bundles: Simplified Visualization of Large Graphs. *Computer Graphics Forum* 29, 3 (2010), 843–852. 2
- [WC07] WONG N., CARPENDALE S.: Supporting Interactive Graph Exploration Using Edge Plucking. *Visualization and Data Analysis 6495* (2007). 2
- [WCG03] WONG N., CARPENDALE S., GREENBERG S.: Edge-Lens: An Interactive Method for Managing Edge Congestion in Graphs. *Proc. IEEE InfoVis 2003* (2003), 51–58. 2
- [WM08] WARE C., MITCHELL P.: Visualizing Graphs in Three Dimensions. *ACM Transactions on Applied Perception* 5, 1 (2008), 1–15. 1, 2
- [ZBDS12] ZINSMAIER M., BRANDES U., DEUSSEN O., STROBELT H.: Interactive Level-of-Detail Rendering of Large Graphs. *IEEE TVCG* 18, 12 (2012), 2486–2495. 6
- [ZWHK14] ZIELASKO D., WEYERS B., HENTSCHEL B., KUHLEN T. W.: Interactive 3D Force-Directed Edge Bundling on Clustered Edges. *In Poster Abstracts of IEEE InfoVis* (2014). 2