Eurographics Conference on Visualization (EuroVis) 2017
J. Heer, T. Ropinski and J. van Wijk
(Guest Editors)

Volume 36 (2017), Number 3

Visual Exploration of Global Trade Networks with Time-Dependent and Weighted Hierarchical Edge Bundles on GPU

J. Hofmann[1], M. Größler[1], M. Rubio-Sánchez[2], P.-P. Pichler[3], and D. J. Lehmann[1,3]

[1]University Magdeburg, Germany,[2]Rey Juan Carlos University Madrid, Spain
[3]Potsdam Institute for Climate Impact Research, Germany

Abstract

The UN Comtrade database is the world's largest repository of bilateral trade data. Their complexity poses a challenge to visualization systems, leading to issues such as scalability and visual clutter. Thus, we propose a radial layout-based visual exploration system to enable the user to smoothly explore the change over time and to explore different commodity classes at once by using a novel edge bundling concept. We evaluated our system with the aid of a group of domain experts.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Graphs and Networks, Hierarchy Data, Time Series Data, GPUs and Multi-core Architectures

1. Introduction

Over the previous two decades, global trade has consistently grown. Regarding this, in the 1960s the United Nations established a program to collect trade data continuously and worldwide, known as the United Nations Commodity Trade Statistics Database (UN Comtrade) [com15]. The UN Comtrade data encompasses the trade data for 170 countries, collected for more than the last 50 years. The data encodes important attributes, like the financial flow, the circulation of products and the exchange of services between countries, which makes them a multidimensional data set.

Data visualization techniques provide relatively easy and intuitive access to large and complex data sets for domain experts (e.g. economists, administration staff or climatologists) who might not always be familiar with other advanced methods for data analysis (e.g. statistics, graph theory). Thus, the analysis of the UN Comtrade data is predestined to be conducted within a visual analytics system. Regarding this, the complexity of the global trade networks poses a challenge on how to design such a visualization system. We propose to use a radial graph layout in order to present an initial graph visualization, which is thought to be the starting point of the visual search. In doing so, we use a scheme to appropriately place the visual representations of the countries along the layout's 1D periphery, where the angle range per representative is interactively parametrized by the current view of the time slice of the data that is supposed to be treated by the user. Note that an alternative approach would be to map the countries and their relations on a 2D spatial domain, e.g., in cartogram-based systems [Win11, KNPS02, KPN05]. This approach might quickly lead to clutter if the data complexity grows. Thus, our approach stresses the structural information of the trade relations more than the geographic position of a country, which can be done by radial graph layout schemes.

In addition, we present a novel weighting-based edge bundling concept that enables to reduce complexity on demand and to find trends quickly. For this, we establish individual edge bundling parameters. This way, our approach allows to optimally use the available screen space; it reduces edge overlapping and clutter effects. Furthermore, we enable the selection of single countries and their trade relations. By extending hierarchical edge bundles, our approach enables an interactive level of detail concept which strikes a balance by facilitating the presentation of the complete complexity, while also allowing to maintain or return to a high level overview (and vice versa). We combine our visualization with an interactive data-dependent animation in order to intuitively encode the direction and the amount of a product to ease the user's perception. We do so by exploiting GPU-based technologies, such as shaders, texture buffers, and vertex attributes, in order to store the complete data representation on the GPU. This enables an interactive scheme, since we also avoid the usual i/o overhead between CPU and GPU. With our concept, the users are able to visually interact with the data in real-time and thus they can reveal, disclose, and explore interesting patterns of the trade data on demand. Clearly, our visualization design targets to:

- present the connecting and time relationship,
- present the hierarchical information, and
- deal with large data

In the following, the related work is provided.

2. Related Work

Our approach relates to graph visualization and edge aggregation.

Graph Visualization: A traditional approach are *node-link diagrams* [GFC04, BBD09]. Here, nodes are depicted as points in a spatial domain and mutually linked, e.g., via line sets. Moreover,

rooted trees [BETT99, HMM00] allow a general hierarchical data visualization by the means of connections from parents to children. To improve the quality of space filling, Shneiderman originally proposed the *treemaps* [JS91, vHR08, SFL10], which recursively split the parent's space according to the children's relative size. While hierarchical subset relations can be figured out well, the relations between edges remain hidden. To ensure a more enclosing and space filling visualization of hierarchical relations and of the edges for trees and graphs, *radial layouts* were successfully established [Smi24, YFDH01, JLB*10, BKH*11]. The leafs are placed along concentric circles, while the edges (and inner nodes) are placed within the circle. Further information may be placed close to its periphery. Our approach is also based on a radial layout scheme.

Edge Aggregation is crucial in order to both reduce visual clutter and reveal structures of the data. For a larger ratio of edges to nodes, matrix-based presentations [GFC04] perform better than node-link diagram-based presentations, even if node-link diagrams show the structure more intuitively for sparse graphs. However, visualizing edges as curves is another way to face this problem [DB78, FWD*03]. Holten [Hol06] describes the *edge bundles* for the aggregation of edges: By using a linear blending scheme between a representative spline [DB78] and the original line between two adjacent nodes, edge clutter can be reduced [JGH09, KG06, SHH11]. See [vHR08] for details. We extend traditional bundling by using a non-uniform weighting. A geometry-based edge clustering is proposed by Weiwei et al. [CZQ*08]: A mesh is utilized as a classifier to disclose those links which should be bundled [CZQ*08]. Force-directed edge bundling has been proposed in [HVW09]. This self-organizing approach stresses the bundling strength according to distance properties. Our approach also relies on edge bundling techniques, which we extend based on the weight of the edge and the angle span between the endpoints.

In addition, we consider further established design approaches w.r.t. focus and context [SZ00], linking and brushing [BC87, TFH11, BVB*13] and interaction abilities [vdEvW11, TS08, CPMT07]. In doing so, we implement the Shneiderman mantra [Shn96] into our system to support the user best and to ease their analysis (and perception) process, which is elementary especially for data that change over time [CSP*06]. For reasons of practical usability, our approach requires top performance: thus, we rely on current shader, buffer, and texture technology by designing and implementing our approach to run completely on GPU [OHL*08].

3. UN Comtrade Data

The global trade networks (see Figure 1) were constructed using bilateral trade data in monetary (USD/yr) and physical (kg/yr) units from UN Comtrade, the largest repository for international trade data. UN Comtrade is maintained by the United Nations Statistics Division (UNSD) and contains reported data on traded goods and commodities of over 170 countries broken down into more than 6000 categories in annual time series since 1962. Data was used from 1995 to 2014 using the Standard International Trade Classification system (SITC, Revision 3) at the three-digit level.

Data quality in the UN Comtrade database varies substantially across different countries, commodities and years in terms of correctness and completeness. Possible causes for errors are plentiful and range from simple allocation mistakes and variation in national reporting standards to intentional distortions for political or economic reasons. Depending on purpose and scope of the question at hand, various methods have been proposed to address these problems, all of which come with their own set of assumptions, pitfalls and limitations. Our main goal for data pre-processing was to apply a conceptually simple and automated method which requires minimal additional information or heuristics and at the same time minimizes the statistical impact of large individual errors. The only additional assumption in the design of the method was to allow for systematic differences concerning the reporting quality of different countries regarding individual commodities. Data pre-processing was performed at the SITC 3-digit level and subsequently re-aggregated to the appropriate level.

Our method to derive estimated trade volumes and quantities from conflicting reports is a maximum likelihood estimation using weighted averages based on a measure of the statistical reliability of reporting countries regarding individual commodities.

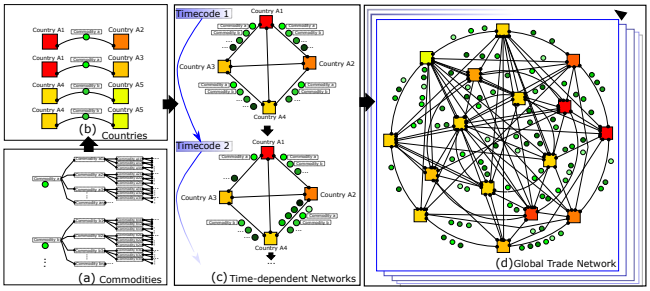

Figure 1: *Structure of UN Comtrade data: The data has three characteristics, the commodities (green), the countries involved in trading (red-yellow), and the level of time in years (blue).*

For a specific year and product, the data can be considered as a compound graph, consisting of a set of nodes and two kinds of relations (see Figure 2). There are relations for the node's hierarchy level and there are additional non-hierarchical components, called adjacency relations. In the context of our data set, the inclusion relations form a tree with the countries as its leaves. An adjacency relation depicts a trade connection between two countries and is therefore only defined between those leaves.

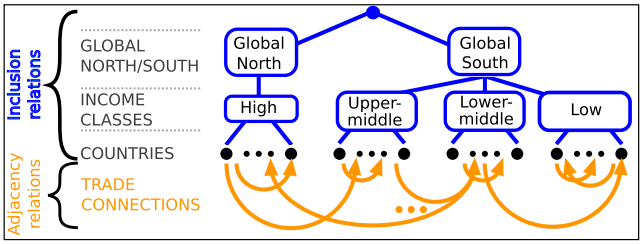

Figure 2: *One compound graph: A tree is build containing the countries at its lowest level. The next higher level groups the countries into classes based on their income. Those groups are then divided into global north and global south. The adjacency relations represent the global trade connections between countries and form a weighted, directed graph.*

Since the global trade network data set is quite complex, large, and even getting larger in the future, a scalable but easy to use interactive visual exploration system is required. For this, we present the design of our visual analysis system in detail in the next section.

4. Technical Background

Original edge rendering techniques rely on line representatives, as Figure 3 (a) illustrates, suffering from overplotting quickly with an increasing number of edges. Holten [Hol06] proposed edge bundles, an edge aggregation technique to resolve overplotting effects.

For this, the nodes $\mathbf{p}_i, i = 1,\ldots,n-1$ along the inner hierarchy path of start node $\mathbf{p}_0$ and end node $\mathbf{p}_n$ define a control polygon. See Figure 3 (b). Now, a spline representation $\mathbf{s}(t)$ is used instead of the initial line representation, illustrated in Figure 3 (c). In order to resolve ambiguities, a parameter β allows to blend between spline and line representation, which is exemplarily shown in Figure 3 (d). Then, the edge bundling model $\mathbf{c}(t), t \in [0,1]$ is given as

$$\mathbf{c}(t) = \beta \cdot \mathbf{s}(t) + (1-\beta)\cdot \mathbf{l}(t), \qquad (1)$$

where $\mathbf{s}(t)$ is a curve defined by the nodes $\mathbf{p}_0,\ldots,\mathbf{p}_n$, $\mathbf{l}(t)$ is a line $\mathbf{l}(t) = \mathbf{p}_0 + t\cdot(\mathbf{p}_n - \mathbf{p}_0)$, and $\beta \in [0,1]$ controls the blending.

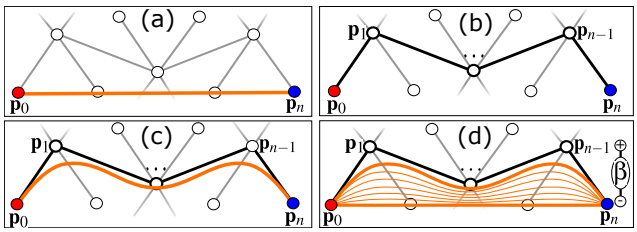

Figure 3: *Holton's hierarchical edge bundling scheme.*

Even though Holton's edge bundles successfully resolve overdrawing, they are accompanied with several issues:

- **Edge Allocation**: The start and end point can be obscured and thus endings cannot be revealed and allocated anymore.
- **Screen Space Fill Ratio:** The screen space might be unevenly filled, which is a waste of screen space and an inconvenient distribution of visual information. In fact, the ratio between filled and unfilled screen space is often poor. For instance, see Figure 4: some wide areas might be empty (orange) while smaller areas convey a large amount of information (blue), which is hard to recognize due to the small size and dense packing.
- **Outlier vs. Trend Analysis:** The bundles visually encode the general hierarchy trends. However, if the graph contains additional variables, such as scalar weighting attributes, then attribute-related trends and outliers cannot be visually analyzed.

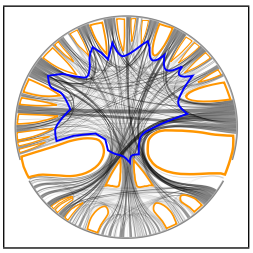

Figure 4: *Fill Ratio*

For our application of analyzing the UN Comtrade data, we extend the traditional edge bundles, called *weighted edge bundles*. We do so in a way that the screen space fill ratio is improved, an outlier vs. trend analysis will be supported and issues of edge allocations are reduced. Moreover, we enhance our concept with further interaction concepts, detailed views, and the option of a user-specific visual parametrization. This way, the aforementioned issues of edge bundling are addressed within one visual analytics scheme for the UN Comtrade data. The next section introduces the details of our visual design.

5. Visual Design for Analyzing Global Trade Networks

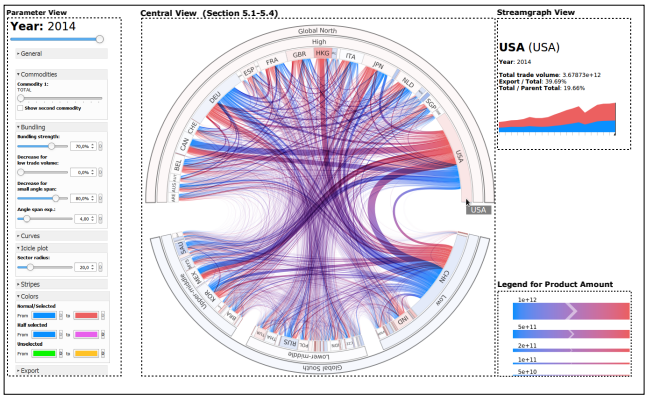

Figure 5: *Overview of our visual design.*

Figure 5 illustrates our visual exploration scheme for the UN Comtrade data. It consists of a set of three coordinated main views.

Streamgraph View: A streamgraph visualization [BW08] is supported in Figure 5 (right), which illustrates detailed information w.r.t. a selected element. It illustrates the evolution of the trade volume (in and out) of a selected country over time (see Figure 6) and thus supports a Focus and Context view to the user w.r.t. the Shneiderman mantra [Shn96].

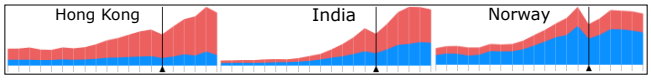

Figure 6: *Stream graphs to enable a detailed view on demand: total import (blue) and export (red) evolution of countries can be understood quickly (black line at* 2009*: impact of financial crisis).*

Parameter View/Interactions: A set of interactions is supported in our concept, which is seen, e.g., in Figure 5 (left). In order to enable a smooth exploration process of the UN Comtrade data, our prototype supports real-time interactions with the radial layer plot for reasons of data exploration, and for the parametrization of visualization properties to satisfy requirements for different groups of users. These interactions encompass: interactive selection of countries (cf. Figure 7), sliding through different years in time (either within the stream graph view – cf. Figure 5 (right) – or via a separating slider), choice of a commodity, interactive choice of the bundling parameters, and a set of parameters to control the visualization properties such as colors, animation speed, enabling flags ('animation','weighted radial tree', etc.), opacity, or line width.

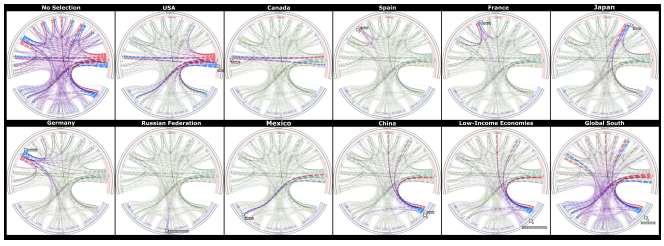

Figure 7: *Interactive selection of countries for 2006 (blue-red).*

Central View: The central element is a radial weighted edge bundling element, which consists of three parts: First, the *weighted*

radial layer plot (Section 5.1); second, the *weighted radial tree plots* (Section 5.2), and third, the *weighted edge bundles plot* (Section 5.3). All these elements are a novel extension for weighted trees w.r.t. the traditional concept from Holten [Hol06] for non-weighted trees. In the following Sections 5.1-5.3, these elements are substantially explained. Finally, we enhance our central view scheme with animations in order to further support the perception of relevant patterns. Section 5.4 explains the related details.

5.1. Weighted Radial Layer Plot

A radial layer plot shapes the outer annulus of our visualization and is required to display the semantic tree elements defined by the amount of the inclusion relations. It consists of one annuli for each level of the tree, with the innermost annulus representing the leafs. Each node in the tree occupies a disjoint circular segment/sector of its corresponding annulus. We start by determining the angles of those circular segments for the innermost annulus. Each segment i is related to a sector angle α_i, and then rendered in a mathematical positive sense; i.e., segment 0 is rendered at 0 degree covering an angle of α_0, segment 1 is rendered at α_0 degree covering an angle of α_1, segment 3 is rendered at $\alpha_0+\alpha_1$ degree covering an angle of α_2, etc. The sum of all angles α_i gives 2π. In detail, the angle α_i of such a segment i consists of a sector angle α_{s_i} of the segment and a gap angle α_{g_i} to the neighbored segment:

$$\alpha_i = \alpha_{s_i} + \alpha_{g_i}.$$

Sector Angle α_{s_i}: Given a sequence of n leafs $(\mathbf{l}_1,\ldots,\mathbf{l}_n)$ representing and indexing the countries: an adjacency $n \times n$ matrix $\mathbf{A}$ encodes pairwise adjacency relations of trade volumes between each pair of leafs $(\mathbf{l}_i\ \mathbf{l}_j)$ – and thus between each country – represented by a weight $a_{ij} \in \mathbf{A}$, with $a_{ii} = 0$, since countries do not trade with themselves per definition. a_{ij} is the total import volume and a_{ji} is the total export volume between country i and j from the perspective of country i. Furthermore, let $w(\mathbf{l}_i)$ be the total trade import and export volume per year for leaf $\mathbf{l}_i$ w.r.t. the country i, and let w_Ω be the total trade volume worldwide with

$$w(\mathbf{l}_i) = \sum_{j=1}^{n} (a_{ij} + a_{ji}) \quad \text{and} \quad w_\Omega = \sum_{i=1}^{n}\sum_{j=1}^{n} a_{ij}.$$

Then, the sector angle $\alpha_{s_i} \in [0, 2\pi]$ of an annular sector for this country i is proportional to the sum of weights of all its adjacency relations $w(\mathbf{l}_i)$, given by:

$$\alpha_{s_i} = \frac{w(\mathbf{l}_i)}{w_\Omega} \cdot (1-g) \cdot \pi,$$

with g being the fraction parameter of the circle that all gaps between the sectors should occupy, $g \in [0,1]$.

Gap Angle α_{g_i}: A gap g_i between two neighboring sectors depends on the distance of its neighbored leafs in the tree, given with $g_i = (LEN(SP(\mathbf{l}_i,\mathbf{l}_j)) - 3) \cdot 0.5$, where $j = ((i \mod n) + 1)$ and $LEN(*)$ being the number of elements of the sequence within the shortest path $SP(\mathbf{l}_i,\mathbf{l}_j)$ from leaf $\mathbf{l}_i$ to $\mathbf{l}_j$. Finally, let g_Ω be the sum of all gaps and α_{g_i} the span angle for the gap g_i with

$$g_\Omega = \sum_{i=1}^{n} g_i \quad \text{and} \quad \alpha_{g_i} = \frac{2\pi \cdot g_i \cdot g}{g_\Omega}.$$

Sectors that represent non-leaf nodes are defined by the outermost angles of their children sectors. Figure 8 (top) shows exemplarily how the sector angle and gab angle for a set of $n = 3$ leaf nodes

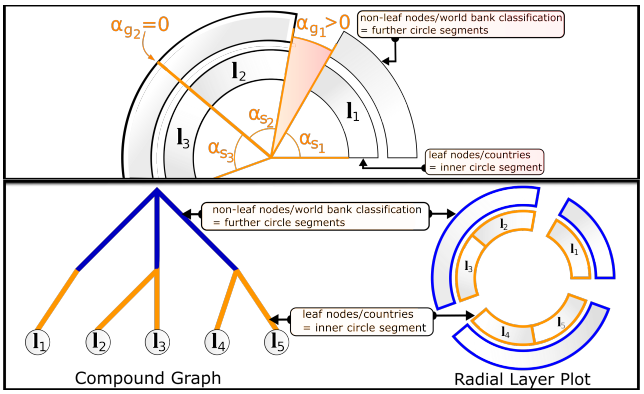

Figure 8: *The construction scheme for our radial layer plot.*

$\mathbf{l}_1,\ldots,\mathbf{l}_3$ would look like, and Figure 8 (bottom) shows how non-leaf nodes are further integrated.

The idea beyond the segment angle α_{s_i} is that larger circle segments represent a larger total amount of trade volume regarding the country that is related to this segment than smaller segments do. This gives an intuitive way to perceive which segments/countries in the visualization participate more in the trade activities.

Since further hierarchical and structural information have to be visually conveyed, gaps and openings between the sectors are an appropriate techniques to do so, expressed by the gap angle α_{g_i}. Due to the fact that our domain experts rely on the World Bank country classification – emphasized in Figure 2, which groups countries to classes such as *Global North*, *Global South* etc. – countries are grouped and thus the gap angle α_{g_i} might often vanish, as Figure 8 (right) reveals. Regarding this, Figure 9 (left) illustrates how our radial layer plot visually encodes the sector information of the leaf nodes aka countries and how non-leaf nodes, which are related to the World Bank country classification, are emphasized by a gap angle α_{g_i} larger than zero.

As an additional source of information, the export excess and the import surplus are optionally visualized (red to blue gradient), making the user quickly aware of which direction of trading prevails, e.g., for a country of interest.

In Figure 9 (right) the resulting radial layer plots for the years 1990, 2002 and 2014 are presented. It can be seen well how the data-driven plots adopt themselves to the related weightings, which relate to the specific year in time. Here, the economical strengthening in total over the last years of the countries belonging to the *Global South* can be figured out easily in this visualization.

In total, our radial layer plot visually conveys (i) information about the amount of trade activities, (ii) the trend or main direction of trading, and (iii) the evolution of these relative relations over the years. Subsequently, we discuss our weighted radial tree concept in order to construct the control polygons later (cf. Section 4).

5.2. Weighted Radial Tree

Before we are able to draw the path of a curve in order to represent the hierarchical structure of the adjacencies between two leafs, we need to map the nodes of our weighted graph onto a 2D space configuration, named radial tree visualization. There, the 2D representations define our curve's control polygons.

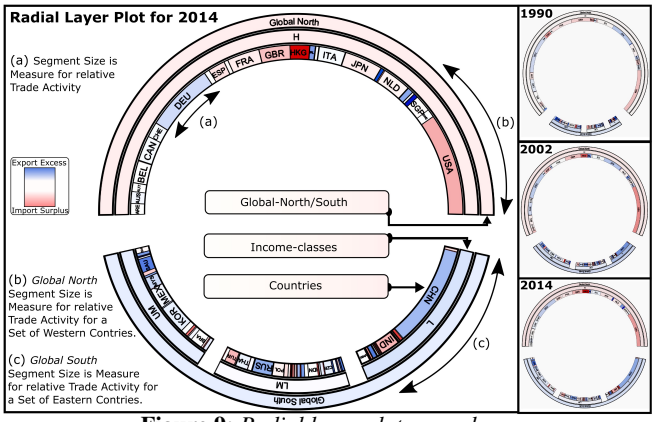

Figure 9: *Radial layer plot examples.*

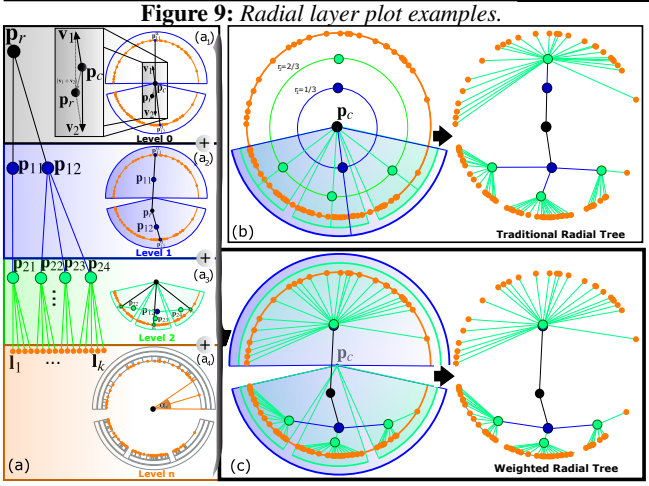

Figure 10: *Construction of the weighted radial tree.*

A traditional radial tree is designed by assuming a circle of radius r: The root is placed in the circle center point $\mathbf{p}_c$, and the inner nodes of each level $l = 0,...,n$ are placed onto circles of radius $r_l = (r \cdot l)/n$, with n being the maximum level number. See [Hol06] for details. Figure 10 (b) illustrates a traditional radial tree for the tree with $n = 3$ of Figure 10 (a). In this traditional concept, the weights of a weighted graph are not considered. Thus, we subsequently describe how our weighted radial tree is constructed.

To prepare the construction, we define the following terms:

Definition 1: We define $\mathbf{p}_{lj}$ as node j of the level l in a tree with n levels ($l = 0,\ldots,n$); a leaf $\mathbf{l}_j$ as $\mathbf{l}_j = \mathbf{p}_{nj}$, and the root $\mathbf{p}_r$ as $\mathbf{p}_r = \mathbf{p}_{00}$. Figure 10 (a) illustrates such a tree with $n = 3$ levels.

Definition 2: Let the *sector angle* $\alpha_s(\mathbf{p}_{lj})$ be the total sum of the sector angles α_i of the leafs contained in the sub-graph of the inner node $\mathbf{p}_{lj}$. If the node $\mathbf{p}_{lj}$ is a leaf $\mathbf{p}_{lj} = \mathbf{l}_j$ then the *sector angle* is $\alpha_s(\mathbf{p}_{lj}) = \alpha_s(\mathbf{l}_j) = \alpha_{s_i}$.

Definition 3: Let the *sector intersection point* $\mathbf{p}^{\alpha}_{lj}$ be a point given by the intersection with the periphery of the unit circle and the ray starting by the circle center $\mathbf{p}_c$ and going into the direction of the angle bisector of $\alpha_s(\mathbf{p}_{lj})$ of the sector within the weighted radial layer plot that is related to the tree node $\mathbf{p}_{lj}$.

Definition 4: Let a *weighting sum* $w_\omega(\mathbf{p}_{lj})$ of a node $\mathbf{p}_{lj}$ be the total sum of all weights $a_i^{\mathbf{p}_{lj}}$ that belong to all of the remaining child nodes of the sub-graph of $\mathbf{p}_{lj}$, given by $w_\omega(\mathbf{p}_{lj}) = \sum_i a_i^{\mathbf{p}_{lj}}$.

Definition 5: Let $\mathbf{p}_{l-k}(\mathbf{p}_{lj})$ be the k^{th} parent node of node $\mathbf{p}_{lj}$, i.e., $\mathbf{p}_{l-1}(\mathbf{p}_{lj})$ is the parent node of node $\mathbf{p}_{lj}$, $\mathbf{p}_{l-2}(\mathbf{p}_{lj})$ is the parent node of the parent node of node $\mathbf{p}_{lj}$, etc.

In order to construct our weighted radial tree placement based on a weighted tree with n levels, we handle root level 0 and level n in a specific way, and apply then a certain scheme to the remaining nodes of the levels 1 to $n-1$.

Root Node Placement (Level 0): The root node $\mathbf{p}_r$ of level 0 is placed on a position that is a centroid of the weightings – see Figure 10 (a_1) – with the aid of shift vectors $\mathbf{v}_j$, based on the first level nodes $\mathbf{p}_{1j}$, given by

$$\mathbf{p}_r = \mathbf{p}_c + \sum_j \mathbf{v}_j \quad \text{with} \quad \mathbf{v}_j(\mathbf{p}_{1j}) = -\frac{\mathbf{p}^{\alpha}_{1j} \cdot w_\omega(\mathbf{p}_{1j})}{||\mathbf{p}^{\alpha}_{1j}|| \cdot w_\omega}. \tag{2}$$

Leaf Nodes Placement (Level n): The nodes $\mathbf{p}_{nj}$ of last level n, i.e., the leafs $\mathbf{l}_j = \mathbf{p}_{nj}$ are placed at their *sector intersection points* $\mathbf{p}^{\alpha}_{nj} = \mathbf{l}^{\alpha}_j$. See Figure 10 ($a_4$).

Nodes Placement (Level 1,...,n-1): A node $\mathbf{p}_{lj}$ of level 1 to $n-1$ is placed – in a top-down scheme, i.e., level 1 nodes at first, then level 2 nodes, then level 3 nodes, etc. – to

$$\mathbf{p}_{lj} = \mathbf{p}^{\alpha}_{lj} + \left[\frac{w_\omega(\mathbf{p}_{lj})}{w_\omega(\mathbf{p}_{l-1}(\mathbf{p}_{lj}))}\right] \cdot (\mathbf{p}_{l-1}(\mathbf{p}_{lj}) - \mathbf{p}^{\alpha}_{lj}). \tag{3}$$

If it applies that $w_\omega(\mathbf{p}_{lj}) = w_\omega(\mathbf{p}_{l-1}(\mathbf{p}_{lj}))$, the node $\mathbf{p}_{lj}$ is placed onto the position of its parent node; and if $w_\omega(\mathbf{p}_{lj}) = 0$ the node $\mathbf{p}_{lj}$ is placed onto the periphery (since then, the weights of $\mathbf{p}_{lj}$'s children are 0, too). This reflects the distribution of the weightings of the graph geometrically. See Figure 10 (a_2-a_3).

Figure 10 (c) illustrates our data-driven weighted radial tree of the graph shown in Figure 10 (a) compared to the traditional radial tree in Figure 10 (b). Note that the distribution of the weighting is represented by the geometrical distribution of the node representatives. Moreover, Figure 11 illustrates the evolution of our data-driven weighted radial tree for different years in time of our data.

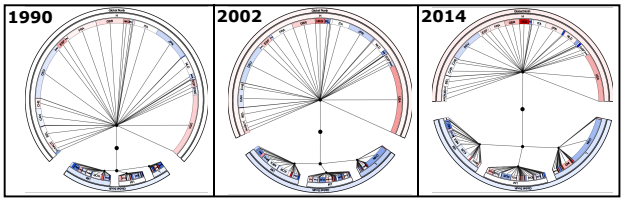

Figure 11: *Our data-driven weighted radial tree visualizations for different years in time of the UN Comtrade data.*

In order to draw an edge bundle later on, we consider a typical control polygon $\mathbf{c}(k_1,k_2)$ for this w.r.t. two leafs $\mathbf{l}_{k_1}$ and $\mathbf{l}_{k_2}$ as to be given by the placement representatives of the shortest path between these two leafs in our weighted radial tree. Figure 12 (top-left/middle) illustrates the control polygon given as the sequence:

$$\mathbf{c}(k_1,k_2) = (\mathbf{l}_{k_1}\ \mathbf{p}_{2j_1}\ \mathbf{p}_{1j_2}\ \mathbf{p}_r\ \mathbf{p}_{1j_3}\ \mathbf{p}_{2j_4}\ \mathbf{l}_{k_2}). \tag{4}$$

To minimize overplotting at the inner circle of the radial tree, we apply a shift operator s_h to the points $\mathbf{p}$ in $\mathbf{c}$, given by

$$\bar{\mathbf{p}} = s_h(\mathbf{p}) = (1 - f_m) \cdot \mathbf{p} + f_m \cdot \mathbf{p}_c, \tag{5}$$

in order to transform the polygon $\mathbf{c}(k_1, k_2)$ into our control polygon

$$\bar{\mathbf{c}}(k_1, k_2) = (\mathbf{l}_{k_1}\ \bar{\mathbf{l}}_{k_1}\ \bar{\mathbf{p}}_{2j_1}\ \bar{\mathbf{p}}_{1j_2}\ \bar{\mathbf{p}}_r\ \bar{\mathbf{p}}_{1j_3}\ \bar{\mathbf{p}}_{2j_4}\ \bar{\mathbf{l}}_{k_2}\ \mathbf{l}_{k_2}). \tag{6}$$

Figure 12 (top-right) illustrates the polygon $\bar{\mathbf{c}}(k_1, k_2)$. In fact, the relative relations between the control points in $\mathbf{c}(k_1, k_2)$ and $\bar{\mathbf{c}}(k_1, k_2)$ are preserved. However, by shifting the control points and thus doubling the start/end points $\mathbf{l}_{k_i}$ and $\bar{\mathbf{l}}_{k_i}$, it can be enforced that edges are perpendicular at the start/end points to the inner circle of the radial tree, as shown in Figure 12 (down-left), and even better in the detailed view in Figure 12 (down-right). For the rest of the work, we consider $\bar{\mathbf{c}}(k_1, k_2)$ as basis to draw our weighted edge bundles, by setting the shift parameter $f_m = 0.2$.

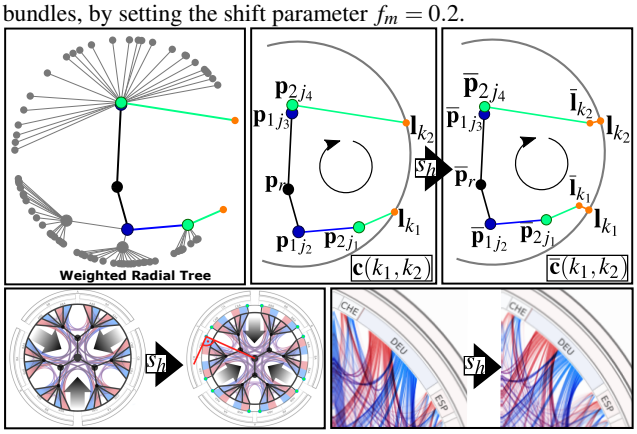

Figure 12: *Construction of our control polygon* $\bar{\mathbf{c}}(k_1, k_2)$.

Fanning out the Start/End Points: Finally, since in each sector that represents a leaf s different edges start/end there, we introduce a parameter $e_b \in [0, 1]$ to fan out these edges along a certain portion of this sector. This eases to recognize which edges belong to each other and to different pairs of leafs as well. For this, we assume that a sector is x-aligned (Figure 13 (a-b)), and spread the edges along the sector as follows: Be $\Delta\alpha = (\alpha_{s_i} \cdot e_b)/2$, then the s start/end points for the edges are placed by $d = 1, \ldots, s$

$$(x_d\ y_d) = (\cos([d-1] \cdot (\Delta\alpha/2))\ \sin([d-1] \cdot (\Delta\alpha/2))). \tag{7}$$

Figure 13 (c) illustrates this for different e_b. For the rest of the work, we consider e_b to be $e_b = 0.5$. Finally, we choose the width of curve $\mathbf{c}(t)_{ij}$ as $\mathbf{c}(t)_{ij} = e_b \cdot \alpha_{s_i}$. Subsequently, we explain how our edge bundles are defined based on our control polygon.

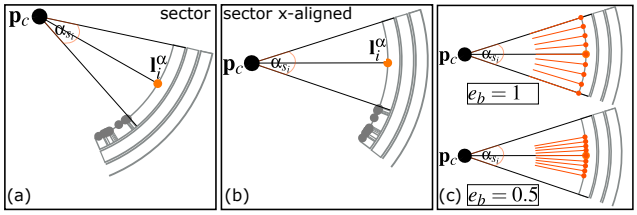

Figure 13: *Start/end point of edges within the sectors.*

5.3. Weighted Edge Bundling

We describe our weighted edge bundling scheme in this section. In our approach, we consider the bundle strength as an individual parameter β_{ij} for each curve $\mathbf{c}(t)_{i,j}$, and thus as a function of the trade volume/weighting a_{ij} between the leafs $\mathbf{l}_i$ and $\mathbf{l}_j$. This enables to parametrize our bundling per curve depending on the weighting factors of the connected leafs as well as of certain geometrical properties of the leafs in the radial layer plot. In detail, we define our individual bundling strength β_{ij} as:

$$\beta_{ij} = \beta_g \cdot \underbrace{\left(1 - \beta_\alpha \cdot \left(\frac{\pi - \Delta\alpha_{ij}}{\pi}\right)\right)^{e_\alpha}}_{\text{Geometrical Influence}} \cdot \underbrace{\left(1 - \beta_o \cdot \left(\frac{w_\omega - a_{ij}}{w_\omega}\right)\right)^{e_o}}_{\text{Weighting Influence}}, \tag{8}$$

where $\beta_\alpha \in [0, 1]$ controls the influence of geometrical properties to the bundling and $\beta_o \in [0, 1]$ controls the influence of the weighting to the bundling. Parameter $\beta_g \in [0, 1]$ steers how strong the two further parameters β_α and β_o shall be considered for the bundling. Then, replacing β with β_{ij} and $\mathbf{c}(t)$ with $\mathbf{c}_{ij}(t)$ in Equation 1 gives the path of the bundled curves that is used in our concept. Note that if $\beta_\alpha = \beta_o = 0$, the curves equate to the traditional concept from Holten, which is thus a boundary case of our individually weighted edge bundles. In the following, we describe in detail how the individual geometrical and weighting properties influence the bundling to enhance the visual exploration of our weighted trees.

Geometrical Adjacent-based Bundling of the Edges with β_α: A geometrical adjacent pair of leafs are leafs that are placed closely together within the radial layer plot (cf. Figure 14 (top)). Curves that connect geometrical adjacent leafs may run along the root node and back (cf. Figure 14 (a)) – depending on the bundling strength – which may confuse the user: the user follows the pathway across the radial tree and back, only to realize then that adjacent leafs are connected. To avoid this kind of confusing detour, the bundling of geometrical adjacent leafs should behave differently.

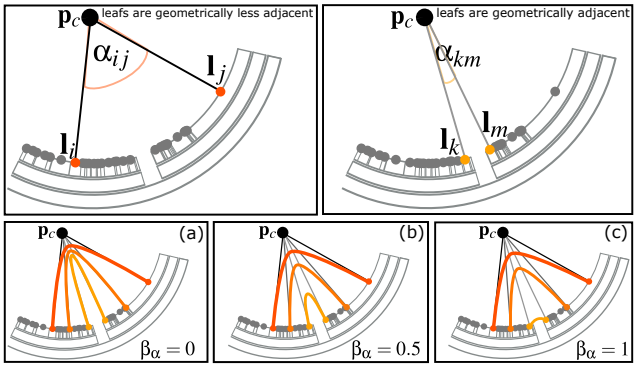

Figure 14: *Geometrical influence to the bundling of parameter* β_α *in order to reduce overlapping of neighbored leafs (with* $\beta_g = 1$*).*

We consider this and introduce an individual bundling factor $\beta_\alpha \in [0, 1]$ to control the bundling of adjacent leafs. Therefore, we rely on the following adjacency criteria:

Definition 6: Be the *span angle* α_{ij} between two leafs $\mathbf{l}_i$ and $\mathbf{l}_j$ given as $\alpha_{ij} = \min(\angle \mathbf{l}_i \mathbf{p}_c \mathbf{l}_j, \angle \mathbf{l}_j \mathbf{p}_c \mathbf{l}_i)$ in arc length, i.e., $\alpha_{ij} \in [0, \pi]$.

Please note that a span angle α_{ij} might cover different sectors and that this term is not equivalent to the term *sector angle* α_{si} (cf. Section 5.1). Figure 14 (top) illustrates how the span angle and the geometrical adjacency are related: Here, $\alpha_{ij} > \alpha_{km}$ means that the leafs $\mathbf{l}_i$ and $\mathbf{l}_j$ are less adjacent than the leafs $\mathbf{l}_k$ and $\mathbf{l}_m$. In our approach, β_α parametrizes how strong the curve is bundled based on

the individual span angle α_{ij}: If β_α is large, the bundling does not apply to edges/curves between pairs of leafs that are adjacent, but strongly affects edges of less adjacent leafs. This way, the detour for edges of adjacent leafs is avoided and overlapping at the inner circle of the radial plot is minimized. Figure 14 (a-c) illustrates how the bundling is reduced for adjacent leafs depending on β_α.

Weighting-based Bundling of the Edges with β_o: The weighting a_{ij} between two leafs l_i and l_j is the most relevant information that shall be visualized for data exploration purposes. Especially information about the trends and outliers are of interest. In order to enable a weighting-based visual data analysis, we consider this by coupling the weighting values a_{ij} together with the bundling strength by using an individual bundling factor $\beta_o \in [0,1]$ to control the weighting-based bundling. The idea is that weighting outliers shall be visually stressed. For this, if β_o is chosen to be close to 1, the edges/curves between two leafs that share a small value of the weightings are not or only barely bundled, and large weighting values cause a strong bundling of the curves between the related leafs. In fact, lower weights are less bundled and vice versa. Figure 15 illustrates the influence of the weightings to the bundling for the choice of different β_o.

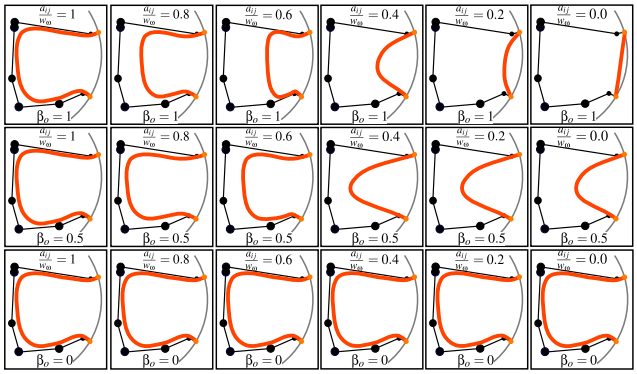

Figure 15: *Weighting influence to bundling of β_o (with $\beta_g = 1$).*

Discussion of the Individual Parameters: Figure 16 illustrates some outcomes of our weighted bundling. On top, Holton's approach is seen by going from 0 to 1 for β_g. Below, our new parameters β_α, β_0 vary (while setting $\beta_g = 1$). When β_α grows, then the adjacent edges almost go to the inner circle and become more clearly this way. Especially the available screen space is used in a better way. When β_o grows, connections with low weights go to the inner circle and become more even, i.e., they appear rather as lines now, while the edges that represents large weights remain curvy. Here, the idea is to increase the user's readability for outliers and trends. This is realized because line-esque patterns of edges more likely represent small weights (= outlier for trade volume) and curvy patterns more likely represent large weights (= trends for trade volume). This way, visual clutter w.r.t. the weighting distribution is reduced and distribution properties are visually distinguishable in a qualitative sense. Let us clarify this with an example in Figure 16: The parameters are $(\beta_g, \beta_\alpha, \beta_o)$. The unbundled layout (0,0,0) in the top left is cluttered. Holton's bundled layout (1,0,0) in the top right is (more or less) uncluttered, but edges with small weighting factors cannot be recognized anymore. Our weight-based bundled layout (1,0,[0.5-1]) in the top row of the parameter matrix shows the main trends still as bundles by curvy edges, but the small weights (more or less) as lines. This gives a trade-off between distribution estimation and clutter reduction. In addition, from our experience, the "green" parameter space is optimal for this analysis.

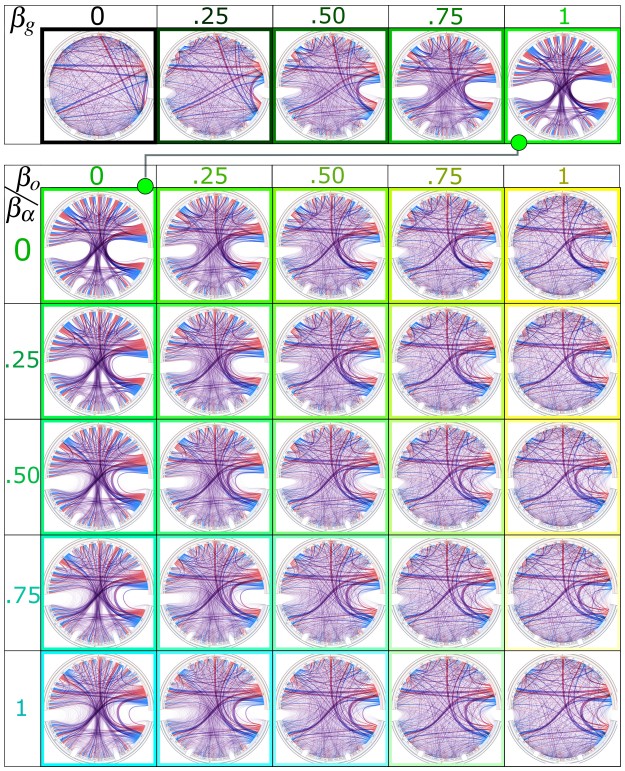

Figure 16: *Weighted Bundling: varying of $\beta_g, \beta_\alpha, \beta_o$.*

5.4. Flowing Stripe Animation

In order to resolve the issue of edge allocation (cf. Sec. 4), and inspired by [HIvWF11], our approach animates stripes along the curves which look like arrows that are going in the direction of the trade volume flow. The animated stripes are realized in the GPUs fragment shader as time-dependent procedural texture where the stripe $s(\textbf{pixel})$ is

$$s(\textbf{pixel}) = 1 + \sin\left(T \cdot g(a_{ij}) - d + f\right) \text{ with } g(a_{ij}) = 1 + (e_b \cdot v_s \cdot 10^{-2}),$$

with T being the running system time, parameter f controlling the stripe density, the product volume function g causing a differential velocity for the stripes on different curves parametrizable by $v_s \in [0,1]$. The local opacity op is given by the Heaviside step function: $op(\textbf{pixel}) = 1$ if $s(\textbf{pixel}) > 1$, 0 else. Figure 17 (top-left) illustrates this scheme and (bottom) shows two stripe outcomes. The faster the stripes move, the larger the related product volume is. Thus, the endpoints along the paths can be better recognized by the user. This is due to the fact that the human perception is sensitive w.r.t. movements and differences in movements [AQS05, EN99]. We exploit this to tackle the edge allocation issue.

5.5. GPU Implementation

This section describes our prototype implementation. Since our application partners usually do not run edge-cutting hardware, neither Nvidia GPUs (i.e., no CUDA) need to be available nor up-to-date DirectX or OpenGL support (i.e., no Geometry or Tessellation shader). Thus, we assume an ordinary GPU system that performs

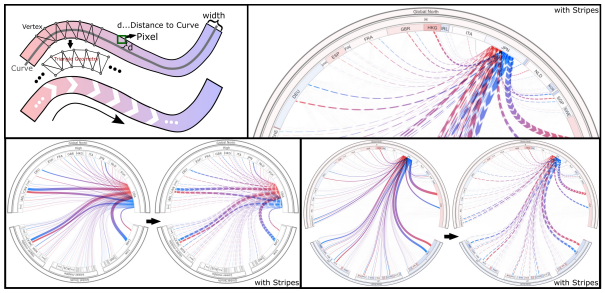

Figure 17: *Flowing Stripes: (top-left) Scheme and (top-right) detailed result, (bottom) two examples from USA and Japan in 2008.*

averagely for our implementation. We restrict ourselves to widely available OpenGL ES 2.0 support (i.e., at least Fragment and Vertex shader are available) to render the curves on GPU in real-time. In addition, our prototype application was implemented in C/C++ and was tested on up-to-date Linux and Windows OS, using the Qt framework for the user interface.

A set of precomputed information is stored in GPU textures and buffers. For instance, information about the geometry, such as the Bézier point position of each Bézier curve $\mathbf{s}(t)$ per triangle vertex and year, the start/end point of line $\mathbf{l}(t)$ on a unit circle per triangle vertex and year, the span angle per leaf of curve $\mathbf{c}_{i\ j}(t)$ and per year, and a set of references, i.e., IDs for curves, triangles, vertices, etc. for texture lookups. In addition, also information about the data are stored on the GPU, such as the trade volumes per available commodity and year. Moreover, a buffer object is stored in GPU that contains all curve segments as triangle strips: Preprocessing of the curves and uploading these data to the GPU on program start ensures that the application only needs to transfer a minimal amount of data afterwards while exploring the data set. The dynamic and interactive bundling function is realized in the vertex shader following Section 5.3, and the flowing stripe animation in the fragment shader following Section 5.4, by adopting the vertex positions in the triangle strips buffer within the fragment shader. Exploration of the data set is realized internally by activating textures and setting uniform variables. The only exception is selecting nodes; a comparatively small texture containing the selection state of all leafs is then changed. Figure 18 illustrates this. With our test system, a notebook with an Intel Core i3-2367M processor (1.40 GHz) with integrated GPU (HD Graphics 3000) and 4 GB of RAM, we were able to consistently provide more than 20 frames per second exploring the global monetary trade data set. This data set contains 201 nodes, up to 10000 visible adjacency edges, 102 triangles per connection (i.e., more than 10^6 triangles) for each year and commodity class and includes data from 25 years (1990-2014) and 11 commodity classes. Due to the low level of required hardware power and GPU storage (only 1 GB right now), the approach scales well for both larger UN Comtrade Data (which can be expected to be available in the future) and a large amount of workstations, laptops, and PCs spread over the world.

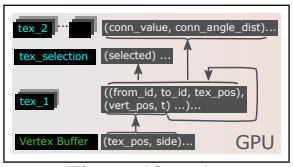

Figure 18: *Selecting*

6. Evaluation

In order to evaluate our application to visually analyze the UN Comtrade data, in Section 6.1 we present user feedback in order to judge the quality of our concept. Following this, the visual analysis capacities are illustrated on a set of economic relations. Please find more evaluation in the additional material.

6.1. User Feedback

To gain user feedback, three participants from the Potsdam Institute for Climate Impact Research tested our application in order to explore the UN Comtrade Network Data.

Conduct: First, we gave an introduction to our system. It required less than 10 minutes, since it was necessary to explain the different parameters in advance. We rated the level of expertise for our participants, whereas 1 is "poor/low" and 5 is "excellent/high". In addition, from the feedback, we derived and rated criteria to judge our concept, such as *Intuitiveness*, *Clarity*, and *Usability* for the visual exploration. Moreover, further criteria are how *Misleading* our approach is, i.e., are the insights correct/false, and how much *Required Training* is needed before the approach can be used (see the embedded table).

Participant	#1	#2	#3
Level of Experience	5	5	3
Intuitiveness	4	4	4
Clarity	4	4	4
Usability	4	4	4
Misleading	1	1	2
Required Training	1	1	1

Analysis: The participants had minor to no problems navigating and exploring the data. It was also regarded as visually pleasing due to the sensible use of animation and color. Most of the users started by exploring the data on a global scale before analyzing specific countries or commodity classes and their evolution over time. By providing three options/parameters for changing the bundling strength, the participant was enabled to emphasize certain aspects of the visualizations as they were needed for further exploration. The uniform bundling strength β_g is useful in terms of finding trends in the world market. The angle-dependent bundling parameter β_α was found to be more useful when it comes to detecting short curves or outliers. No additional clutter was produced, since the number of crossing curves remained constant. With the trade volume-weighted parameter β_o it was easier to follow the curves' evolution when they moved over time. To explore single countries regarding small amounts of traded goods, the bundling was loosened. In conclusion, our participants were enabled to both acquire insight in global trends and reveal hidden structures. Additional things demanded by the users were an easier comparison of two countries; and, depending on the task (global or local exploration), it would have been helpful to start with suitable parameters. Either way, in total, the participants appreciated our visual design. Subsequently, we present examples for the visual exploration of economic relations, as they partly turned out from our user feedback.

6.2. Exploring Skills: Investigate Known Economic Facts

Our system facilitates to visually detect known economic facts.

Financial World Crisis in 2007/08: Fig. 19(a) shows the impact of the global financial crisis. A significant drop (thinner curves) in trade volume is noticed in 2008/09, reflecting the crisis' footprint.

China's Growth: In Fig. 19(b) three commodity classes illustrate this development: From 1990, China started to grow from hardly noticeable to the world's leading importer in 2014.

Great Britain's Restitution of Hong Kong: Fig. 19(c) shows the economic impact when Hong Kong was returned to China in 1997. From 1992 to 1997 Hong Kong's import from China was constant, but significantly grew between 1997 and 2014. Interestingly, now its main source of income has changed to the area of service.

Influence of Free Trade Agreement NAFTA: In 1994, the North

American Free Trade Agreement (NAFTA) was signed by Mexico, Canada and the United States. In Fig. 19(d), the growing trade of Canada (top row) and Mexico (bottom row) is shown from 1994 to 2014, illustrating the impact of the agreement (US, red circle).

Big Player of Economy: Fig. 19(e) depicts the world's largest economies in 2014: China is the biggest exporter and the US the largest importer. In general, the world market is dominated by Western Europe, US, China & Japan. Our system shows Germany's extreme economic openness in terms of its large trade volume relative to its small population when compared to the US or China.

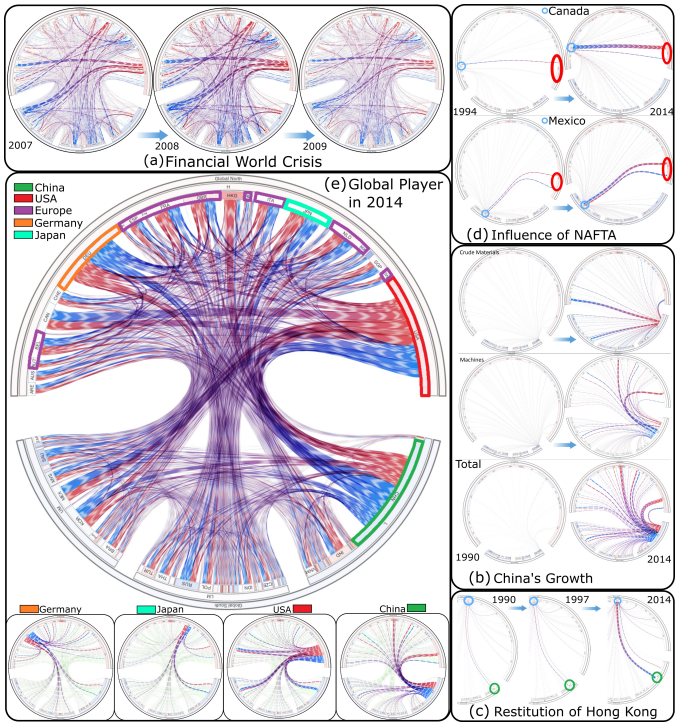

Figure 19: *Investigate known economic facts.*

Figure 20: *Reveal unknown economic hypotheses.*

6.3. Exploring Skills: Discover Unknown Phenomena

Our system aids the discovery of phenomena previously unknown to the experimenter.

Uncared Export Nations: Fig. 20(a) shows that Brazil and Australia started growing their crude material exports alongside China's growth. This is likely causally connected as increasing prices from Chinese demand and geographic proximity made this profitable.

Unrecognized Oil Champions: Fig. 20(c) shows that Malaysia and Indonesia dominate the export of animal and vegetable oils, fats and waxes for which increasing global demand exists. This was surprising to the experimenter who previously assumed those products to originate from glasshouses in the US.

Unrecognized Fuels Relations: Fig. 20(c) shows an interesting pattern of the fuel export. The US and Japan import energy fuels primarily from Canada, Saudi Arabia and the Middle East. Europe, on the other hand, imports mostly from Norway and Russia and is less dependent on the Middle East than expected.

Unrecognized African Poverty Circle: Fig. 20(d) shows that in 2010 energy fuel exports from a number of African countries was comparable to that of the Middle East while the same countries depended on imports for more advanced manufactures (e.g. machines). The oil richness of these countries has not led to substantial development and most remain poor.

Unrecognized World Crisis Impact: Fig. 20(e) shows that the global financial crisis has disproportionately affected trade in rich countries. While trade volumes of the US, Germany and Japan have markedly declined, the crisis had little effect on Bangladesh, Afghanistan and Greece.

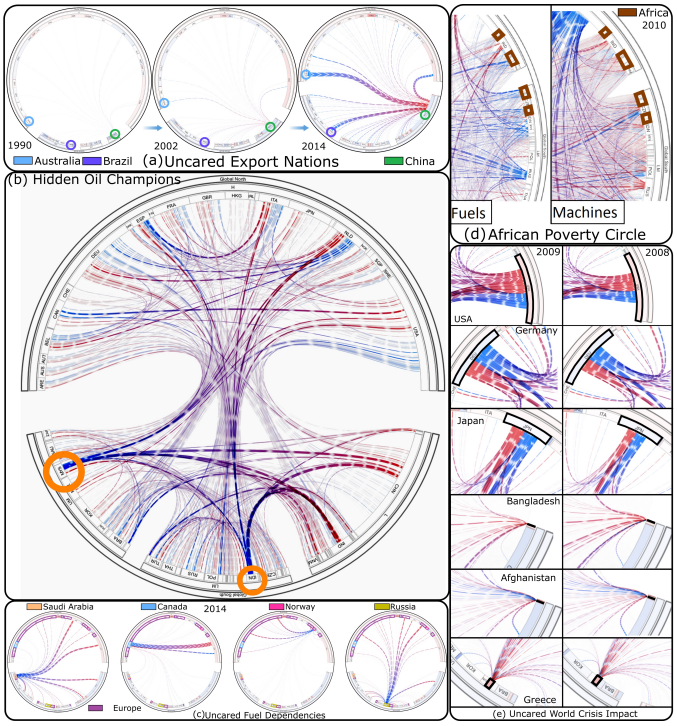

7. Discussion

We presented a novel system for the visual search in the UN Comtrade data base with radial network visualization schemes. For this, we extend Holton's traditional hierarchical edge bundling scheme for (time-dependent) weighted compound trees, which reflect the change of weightings over time in the shape of our novel weighted radial layer plot (cf. Sec 5.1), the control polygons of our novel weighted radial tree (cf. Sec 5.2), and in the parameter space of the exploration by our novel weighting-based edge bundling parameter β_o (cf. Sec. 5.3). This way, we contribute an extension of the traditional edge bundling concept. Moreover, our concept addresses some visual shortcomings from Holton's pioneer work: In order to allow to minimize **Edge Allocation** issues (cf. Sec. 4), we introduce the concept to fan out start/endpoints (cf. Sec 5.2) and the use of stripe flow animation with different flow speed (cf. Sec 5.4). To improve the **Screen Space Fill Ratio** (cf. Sec. 4), we introduce the geometrically bundling relations (cf. Sec. 5.3); and to enable a more intuitive **Outlier vs. Trend Analysis:** (cf. Sec. 4), the weighted-bundling concept has been established (cf. Sec. 5.3).

However, we do so for a specific application, namely for interactively analyzing the UN Comtrade data. In fact, our approach enables the use case of analyzing different commodities for hundreds of countries and dozens of years at once in real-time by our scheme, GPU implementation, and a set of reasonable interaction approaches. Thus, our approach allows an intuitive visual analytics process for those data. In Sec. 6, as proof of concept, the evaluation illustrates the usability of our concept. Nevertheless, some improvements ought to be done in the future, such as enabling a selection of a larger set of countries at once or an automatic parameter detection. We leave this to future work.

References

[AQS05] AZNAR J., QUEVEDO L., SINNETT S.: The effects of drift and displacement motion on dynamic visual acuity. *Psicologica*, 26 (2005), 105–119. 7

[BBD09] BECK F., BURCH M., DIEHL S.: Towards an aesthetic dimensions framework for dynamic graph visualisations. In *Information Visualisation, 13th International Conference* (2009), pp. 592–597. 1

[BC87] BECKER R., CLEVELAND W.: Brushing scatterplots. *Technometrics 29, 2* (1987), 127–142. 2

[BETT99] BATTISTA D., EADES P., TOLLIS I. G., TAMASSIA R.: *Graph drawing: algorithms for the visualization of graphs*. 1999. 2

[BKH*11] BURCH M., KONEVTSOVA N., HEINRICH J., HÖFERLIN M., WEISKOPF D.: Evaluation of traditional, orthogonal, and radial tree diagrams by an eye tracking study. *IEEE TVCG 17*, 12 (2011), 2440ff. 2

[BVB*13] BORKIN M. A., VO A. A., BYLINSKII Z., ISOLA P., SUNKAVALLI S., OLIVA A., PFISTER H.: What makes a visualization memorable? *IEEE InfoVis* (2013). 2

[BW08] BYRON L., WATTENBERG M.: Stacked graphs – geometry & aesthetics. *IEEE TVCG 14*, 6 (Nov. 2008), 1245–1252. 3

[com15] UN Comtrade, 2015. Data retrieved 2015-12-01, http://comtrade.un.org/. 1

[CPMT07] CUADROS A. M., PAULOVICH F. V., MINGHIM R., TELLES G. P.: Point placement by phylogenetic trees and its application to visual analysis of document collections. In *IEEE VAST* (2007), pp. 99–106. 2

[CSP*06] CARD S. K., SUH B., PENDLETON B. A., HEER J., BODNAR J. W.: TimeTree: Exploring time changing hierarchies. In *IEEE VAST* (2006), pp. 3–10. 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), 1277–1284. 2

[DB78] DE BOOR C.: *A practical guide to splines*, vol. 27. Springer-Verlag New York, 1978. 2

[EN99] EDWARDS M., NISHIDA S.: Global-motion detection with transparent-motion signals. *Vision Research 39*, 13 (1999), 2239ff. 7

[FWD*03] FEKETE J., WANG D., DANG N., ARIS A., PLAISANT C.: Interactive poster: Overlaying graph links on treemaps. In *IEEE InfoVis* (2003), pp. 82–83. 2

[GFC04] GHONIEM M., FEKETE J.-D., CASTAGLIOLA P.: A comparison of the readability of graphs using node-link and matrix-based representations. In *IEEE InfoVis* (2004), pp. 17–24. 1, 2

[HIvWF11] HOLTEN D., ISENBERG P., VAN WIJK J. J., FEKETE J.-D.: An extended evaluation of the readability of tapered, animated, and textured directed-edge representations in node-link graphs. In *PacificVis* (2011), pp. 195–202. 7

[HMM00] HERMAN I., MELANCON G., MARSHALL M. S.: Graph visualization and navigation in information visualization: A survey. *IEEE TVCG 6*, 1 (2000), 24–43. 2

[Hol06] HOLTEN D.: Hierarchical Edge Bundles: Visualization of Adjacency Relations in Hierarchical Data. *IEEE TVCG 12*, 5 (2006), 741–748. 2, 3, 4, 5

[HVW09] HOLTEN D., VAN WIJK J. J.: Force-directed edge bundling for graph visualization. *Computer Graphics Forum 28*, 3 (2009), 983–990. 2

[JGH09] JIA Y., GARLAND M., HART J. C.: Hierarchical edge bundles for general graphs. *Technical Report, Illinois Digitial Enviroment for Access to Learning and Scholarship* (2009). 2

[JLB*10] JIA M., LI L., BOGGESS E., WURTELE E. S., DICKERSON J. A.: Visualizing multivariate hierarchic data using enhanced radial space-filling layout. In *ISVC* (2010), pp. 350–360. 2

[JS91] 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. 2

[KG06] KUMAR G., GARLAND M.: Visual exploration of complex time-varying graphs. *IEEE TVCG 12*, 5 (2006), 805–812. 2

[KNPS02] KEIM D. A., NORTH S. C., PANSE C., SCHNEIDEWIND J.: Efficient cartogram generation: a comparison. In *IEEE InfoVis* (2002), vol. 0, pp. 33–36. 1

[KPN05] KEIM D. A., PANSE C., NORTH S. C.: Medial-Axis-Based Cartograms. *IEEE Computer Graphics and Applications 25*, 3 (2005), 60–68. 1

[OHL*08] OWENS J. D., HOUSTON M., LUEBKE D., GREEN S., STONE J. E., PHILLIPS J. C.: GPU computing. *Proceedings of the IEEE 96*, 5 (2008), 879–899. 2

[SFL10] SUD A., FISHER D., LEE H.-P.: Fast dynamic voronoi treemaps. In *International Symposium on Voronoi Diagrams in Science and Engineering (ISVD)* (2010), pp. 85–94. 2

[SHH11] SELASSIE D., HELLER B., HEER J.: Divided edge bundling for directional network data. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)* (2011). 2

[Shn96] SHNEIDERMAN B.: The eyes have it: A task by data type taxonomy for information visualizations. *IEEE Symposium on Visual Languages* (1996), 336–343. 2, 3

[Smi24] SMITH W. H.: Graphic statistics in management. *McGraw-Hill Book Company, New York* (1924). 2

[SZ00] STASKO J., ZHANG E.: Focus+context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. In *IEEE InfoVis* (2000), pp. 57–65. 2

[TFH11] TURKAY C., FILZMOSER P., HAUSER H.: Brushing dimensions - a dual visual analysis model for high-dimensional data. *IEEE Trans. Vis. Comput. Graph. 17*, 12 (2011), 2591–2599. 2

[TS08] TEKUSOVA T., SCHRECK T.: Visualizing time-dependent data in multivariate hierarchic plots – design and evaluation of an economic application. In *IEEE InfoVis* (2008), pp. 143–150. 2

[vdEvW11] VAN DEN ELZEN S., VAN WIJK J. J.: BaobabView: Interactive construction and analysis of decision trees. In *IEEE VAST* (2011), pp. 151–160. 2

[vHR08] VAN HAM F., ROGOWITZ B.: Perceptual organization in user-generated graph layouts. *IEEE TVCG 14*, 6 (2008), 1333–1339. 2

[Win11] WINTER M. J.: Diffusion Cartograms for the Display of Periodic Table Data. *J. Chem. Educ. 88*, 11 (July 2011), 1507–1510. 1

[YFDH01] YEE K.-P., FISHER D., DHAMIJA R., HEARST M.: Animated exploration of dynamic graphs with radial layout. In *IEEE InfoVis* (2001), p. 43ff. 2