Exploring Media Transparency With Multiple Views

Alexander Rind,[1,2] David Pfahler,[1] Christina Niederer,[2] and Wolfgang Aigner[1,2]

[1] TU Wien, Austria
[2] St. Poelten University of Applied Sciences, Austria

Abstract

Politically concerned citizens and data journalists want to investigate money flows from government to media, which are documented as open government data on 'media transparency'. This dataset can be characterized as a dynamic bipartite network with quantitative flows and a large number of vertices. Currently, there is no adequate visualization approach for data of this structure. We designed a visualization providing coordinated multiple views of aggregated attribute values as well as short tables of top sorted vertices that can be explored in detail by linked selection across multiple views. A derived attribute 'trend' allows selection of flows with increasing or decreasing volume. The design study concludes with directions for future work.

1 Introduction

Independent news and media are a cornerstone of modern democracy – often called the fourth power. However, governmental advertisement and sponsorships could influence news coverage limiting the media's independence. In Austria, the federal law on Transparency in Media Cooperation and Funding [Med15] makes it mandatory to disclose such flows of money from legal entities (e.g., federal ministries, cities, economic chambers, government-owned companies) to media institutions (e.g., newspaper, TV, radio, online). The Austrian Regulatory Authority for Broadcasting and Telecommunications [RTR] collects these data and makes them publicly available via the Austrian open government data portal [RTR16].

Copyright © by the paper's authors. Copying permitted for private and academic purposes.

In: W. Aigner, G. Schmiedl, K. Blumenstein, M. Zeppelzauer (eds.): Proceedings of the 9th Forum Media Technology 2016, St. Pölten, Austria, 24-11-2016, published at http://ceur-ws.org

This so-called media transparency (MT) dataset is a valuable resource for politically concerned citizens as well as for data journalists [Aus15, Lor10]. They are interested in exploring the available data independently looking for stories beyond prearranged summary statistics. However, the MT dataset is much too large to be browsed line by line. Neither is it sufficient to look only at the largest flows of money because many possible questions of interest focus on changes over time and the many-to-many relationship between legal entities and media [NRA+16]. For this purpose it is useful to conceptualize the MT dataset's money flows as time-dependent attributes on the edges of a bipartite network. Simple data analysis tools such as spreadsheets do not adequately support such a data structure.

Interactive visual representations of data [CMS99, Mun14] are a well-suited approach to explore complex datasets. Many visualization techniques have demonstrated their value in exploring time-oriented data [AMST11] and network data [BBDW16, HSS15]. However, the combination of time with quantitative flows in a bipartite networks is still an open challenge for visualization research [NAR15].

This paper contributes a visualization design study [SMM12] for time-oriented quantitative flows in a bipartite network. It uses the MT dataset as example and non-expert users such as citizens and journalists as target audience. After surveying related work in Section 2 and characterizing the domain problem in Section 3, we present the justified visualization design in Section 4. Next, a usage scenario demonstrates the design's utility in Section 5. The paper concludes with reflections for future development.

2 Related Work

The design space of network visualization has been mapped in some recent state-of-the-art reports: Hadlak et al. [HSS15] identified five facets of concern for visualizing a network: (i) its structure comprised of nodes and edges, (ii) partitions, (iii) the attributes of nodes and edges, (iv) dynamics, i.e., change over

Table 1: Raw format of the media transparency (MT) dataset (first four entries)

legal entity	time	law	medium	money amount
Abfallwirtschaft Tirol-Mitte GesmbH	Q4/2012	§2	Bezirksblätter Tirol	8,122.32 €
Agrarmarkt Austria Marketing GesmbH	Q4/2012	§2	Falstaff	26,418.00 €
Agrarmarkt Austria Marketing GesmbH	Q4/2012	§2	Connoisseur Circle	6,142.50 €
Agrarmarkt Austria Marketing GesmbH	Q4/2012	§2	bz-Wiener Bezirkszeitung	7,031.16 €
	⋮			

time, and (v) spatialization such as geographic context of nodes. Beck et al. [BBDW16] addressed in particular visual representations for dynamic networks such as animation and timeline. Von Landesberger et al. [vLKS+11] focused on large networks. Niederer et al. [NAR15] surveyed visualization of dynamic, weighted and directed networks, and thus, data of a structure similar to the MT dataset.

Examples of such related visualization designs are DOSA [vdEvW14], egoSlider [WPZ+16], egoLines [ZGC+16], Graph Comics [BKH+16], TimeArcTree [GBD09], and Visual Adjaceny List [HBW14]. However, none of these approaches explicitly considers the bipartite nature of the MT dataset, i.e., that there are distinct nodes for legal entities and for media.

We could identify only one scholarly work focusing on visualizing the MT dataset in particular: Niederer et al. [NRA+16] investigated the visualization needs of data journalists based on four interviews that were anchored on the MT dataset as exemplary scenario. Besides that, there is some press coverage on the data and some articles are accompanied by interactive web infographics (e.g., derStandard.at [Ham], Paroli Magazin [Lan]). Yet, these infographics present a subset of the available data that has been aggregated and filtered to support their articles' story. Since they allow only minimal interactivity, further exploration is not possible. Furthermore, since 2013 the open source software project Medientransparenz Austria [SBSV] provides an interactive online tool that shows the complete MT dataset. It integrates several visual representations giving insight into the data, but its views require much scrolling and are distributed across multiple pages. In addition, changes of money flow over time are not explicitly represented.

3 Background

As a fundament for developing a novel visualization design for the MT dataset, we must first understand its background and characterize the domain problem.

The law [Med15] regulates three categories of money flows that need to be disclosed: §2 covers advertisement, §4 sponsorships, and §31 ORF programme fees. Each quarter, each legal entity is obligated to make a disclosure for both §2 and §4. Every media cooperation involving more than 5,000 € needs to be included with the recipient's name and the amount of money accumulated in the quarter. If a legal entity had no such media cooperation, it still has to submit a nil report.

The MT dataset is published on an open data portal [RTR16] each quarter of a year with data covering the preceding eight quarters. The raw data are formatted as semicolon-separated values in a text file. Table 1 shows the five relevant variables: name of the legal entity, time specified by year and quarter, category of legal background, name of the medium, amount of money (quantitative). Additionally, the raw data contains a variable that flags nil reports.

3.1 Data Abstraction

The MT dataset is comprised of the quarterly money transferred from legal entities to media. We can conceptualize these data as time-dependent flows in a bipartite network (Figure 1) [NRA+16]. The network's underlying graph is bipartite because its vertices can be divided in two disjoint sets – legal entities and media – and each edge connects vertices of different sets. These edges are directed and weighted representing the flow of money from legal entities to media. The network is dynamic both in terms of its structure (vertices and edges can appear or disappear over time) and its quantitative flows (weights changing over time) [vLKS+11]. The time-oriented aspect of the data can be characterized as instants on a discrete, interval-

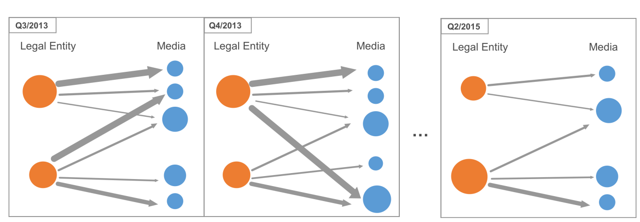

Figure 1: Conceptualizing the MT dataset as time-dependent flows in a bipartite network
Figure by [NRA+16] used with permission.

based, linear time domain with the granularities quarter and year [AMST11].

This abstract data structure has some benefits over the raw data's table structure: The central aspects of the problem domain (legal entities, media, and flows) are represented directly as data items, which can have properties from derived data such as aggregated money. Network metrics such as in-degree can be examined. They can also be manipulated by user interaction.

3.2 Preprocessing and Analysis of Data Scale

We perform some preprocessing to achieve a better data basis for our visualization:

(1) We substitute the original MT dataset with data from the Medientransparenz Austria project [SBSV], which have two benefits: First, they have included data for all quarters since the start of the MT dataset in Q3/2012. Second, they have preprocessed the data to clean different forms of writing the names of media and legal entities. Such inconsistencies could result either from typos or from the organization actually being renamed.

(2) Next, we discard nil reports from the data. Even though these nil reports make up about 80% of all records, they cannot add any insight to our design as they have missing values for media name and amount of money.

(3) Finally, we also discard programme fees (legal category §31) because on the one hand there are only one or two records per quarter and on the other hand their amount is much higher than any other record. The median §31 amount is about 80 times as much as the highest regular amount.

As of summer 2016, the preprocessed MT dataset encompasses 36,261 quarterly money flows over 15 quarters (Q3/2012–Q1/2016). So that one quarter has on average circa 2,400 flows. 34,717 flows (96%) have §2 as legal background and there are 1,544 flows for §4. (30 flows for §31 have been discarded.)

There are 993 distinct legal entities and 3,813 distinct media. Legal entities have between 1 and 1,782 outgoing flows (median = 8; average = 36.2), if we count each quarter as a separate flow. These flows connect them to between 1 and 618 distinct media (median = 3; average = 12.2). 71 legal entities maintain a continuous flow over all 15 quarters to between 1 and 24 media. Media have between 1 and 1,577 incoming flows (median = 1; average = 9.4). These flows connect them to between 1 and 285 distinct legal entities (median = 1; average = 3.2). 68 media maintain a continuous flow over all 15 quarters from between 1 and 18 legal entities.

The quantitative values of quarterly flows vary between € 5,000 (the minimum to be reported) and € 1.929.533 (median = € 10,931; average = € 23,444).

3.3 Design Requirements

Based on our data analysis described above and the interviews with data journalists interested in the MT dataset as reported by Niederer et al. [NRA+16], we can identify five design requirements that a visualization design for the MT dataset should fulfill:

R1 Data scalability: The number of vertices for both legal entities and media is relatively large. Besides the institutions' names and their network relations, there are no further data that could be used for clustering vertices. While a majority of vertices is only sparsely connected, some central vertices have a large number of flows. Likewise, the weights representing amount of money can vary widely within the network. The time dimension adds additional scale.

R2 Development over time: The data journalists interviewed by Niederer et al. [NRA+16] expressed particular interest in patterns or abnormalities in the number and weight of flows over time.

R3 Data wrangling: For two reasons, users would need to refine the MT dataset by basic data wrangling functionality: First, they can add their implicit expert knowledge into the analysis. For example, they could group together the federal ministries run by politicians of the same party.
Second, data quality is still not sufficient for some data entries even though data quality measures have been taken by the RTR and the dataset has been pre-cleaned by the Medientransparenz Austria project. Table 2 shows some examples based on media from this dataset containing the string "standard". It should be possible to combine entries with different forms of writing or different media (print, online, app) of the same newspaper and to hide entries of poor quality.

R4 Ease of use: The target audience of the MT dataset such as interested citizens or data journalists will most likely have no expert knowledge of statistics or visualization. They will access the MT visualization as a spontaneous activity where no special training can be provided. Therefore, care should be taken that well-known visualization techniques are chosen and the user interface is self-explaining.

R5 Interactive exploration: Some users will approach the MT visualization trying to verify an existing hypothesis but we expect that a majority of usage session will consist of undirected exploration in search for patterns of interest. For this, interac-

Table 2: Media matching the query string “standard” ordered by the number of connected legal entities and showing the aggregated sum of transferred money. Three entries are different forms of writing the same website. The fifth entry contains the names of six separate newspapers and stands as example of inconsistent data collection when he MT dataset was started in 2012.

medium	#rel.	summed flows
Der Standard	189	18,905,741 €
derstandard.at	64	2,768,875 €
www.derstandard.at	19	312,242 €
Der Standard KOMPAKT	2	44,745 €
Standard Verlagsgesellschaft m.b.H.	1	3,099,082 €
Krone, Kurier, Presse, Salzburger Nachrichten, Standard, Kleine Zeitung	1	90,874 €
derstandard.at App	1	11,510 €
ES Evening Standard Magazine	1	10,884 €
http://www.derstandard.at	1	9,938 €

tive features are needed that are usable and help users maintain overview.

4 Visualization Design

Based on these design requirements, we developed a visualization design for the MT dataset (Figure 2). This section describes the design and explains its underlying rationale. In Subsection 4.1 the individual diagram views of the design are presented. How the user is able to interact with them is described in Subsection 4.2 and how the views are linked with each other is delineated in Subsection 4.3.

4.1 Attribute Visualization

The MT dataset contains 5 data attributes. The columns of Table 1 display these data attributes. It is not possible to visualize every single data record of this table in the dashboard, therefore the records are aggregated and the aggregated information is displayed. [Mun14, p. 305]

For example Figure 2.A shows aggregated data of money transferred over time. For this, the data attribute “money amount” is summed for all data records with the same value in the data attribute “time”. This reduces the data to only 2 data attributes and only 1 data record per quarter. The sum is a quantitative attribute and the quarters can be handled as ordinal attribute. A bar chart suites the task of looking up and comparing the values of the different quarters best [Mun14, p. 150]. A similar aggregation is visualized for sum of money transferred by legal background in Figure 2.B.

To visualize the distribution of a single quantitative attribute a histogram can be used [Mun14, p. 306]. Figure 2.C shows a histogram of the data attribute “money amount”.

Figure 2.D is a histogram of the data attribute “trend”. This attribute is derived from all amounts of money e_i flowing from a legal entity to a medium over time i. The trend T quantifies the relative difference of money transferred between the first half of the quarters $|Q_m|$ and the second half.

$$|Q_m| = \left\lfloor \frac{|Q|}{2} \right\rfloor \tag{1}$$

$$T = \frac{\sum_{i=|Q_m|}^{|Q|} e_i - \sum_{i=1}^{|Q_m|-1} e_i}{\sum_{i=1}^{|Q|} e_i} \tag{2}$$

The categorical data attributes legal entity and medium both have a large number of categories (see Subsection 3.2), which are too many to visualize them in a bar chart. Neither is it possible to aggregate the categories in a reasonable way. But the entries can be sorted by another aggregated quantitative data attribute so that only the most relevant ones are displayed For example, it is possible to sort legal entities by the sum of money transferred from them to media. Figure 2.E shows the details for the first 10 sorted legal entities as a table with 4 columns. The first column shows the name of the legal entity. The second column shows the sum of the transferred money to various media over time. Additionally a sparkline sized bar chart represents the distribution of the transferred money over time. This enables the user to see the trend over time [Tuf06, AMST11]. The third column displays the number of relations, i.e. the count of media receiving money from the legal entity. The forth column displays the “trend” as calculated by Equation 1. This data table visualization enables the user to receive detailed aggregated information for a few entities. Figure 2.F applies the same visual representation to the first 10 sorted media.

The last visualization in Figure 2.G shows the flow of money from legal entities to media using a chord diagram [KSB+09]. The aggregated amount of money is encoded with the length of an arc of a circle segment of the diagram. This allows the user to see from which legal entity how much money is transferred to which medium. Because there are too many different categories, placeholder segments are generated for legal entities and media, which contain all not displayed entries and aggregate the money for all of them.

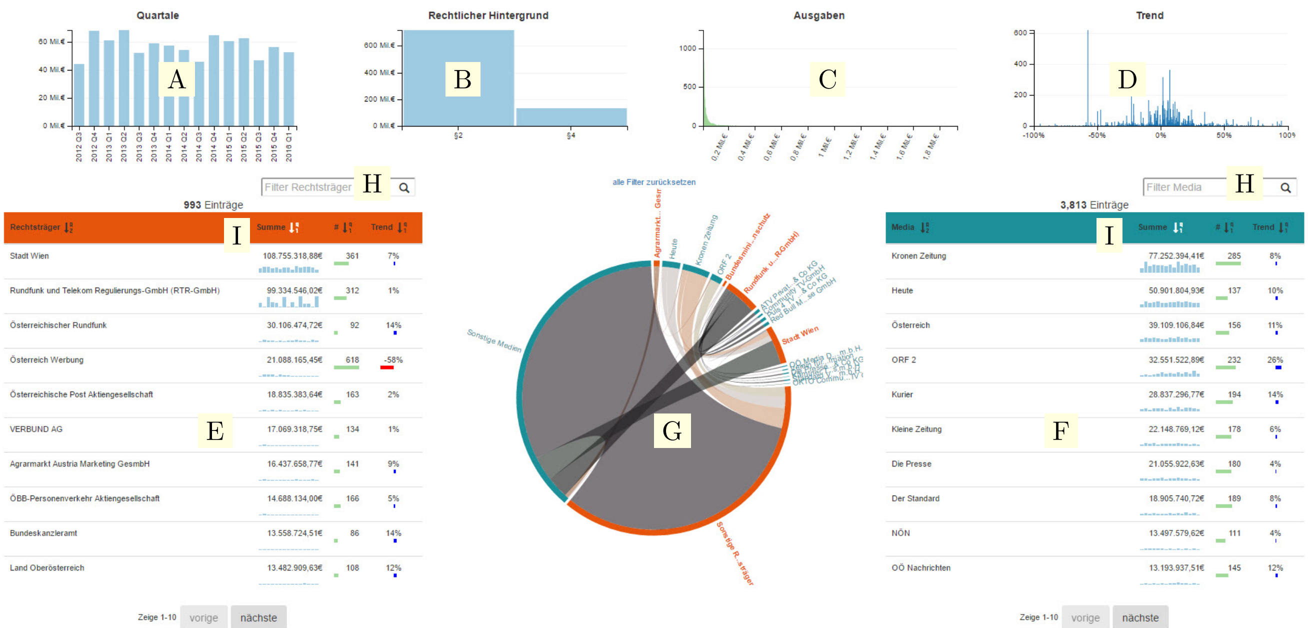

Figure 2: The visualization design for the media transparency (MT) dataset is comprised of seven views: (A) bar chart of aggregated money by time, (B) second bar chart for money by legal category, (C) histogram of money by flow, (D) histogram of increasing/decreasing trend, (E) table of 10 legal entities with total money, sparkline of money, number of connected media, and trend, (F) second table of top 10 media, and (G) chord diagram of flows. Both tables can be (H) searched and (I) sorted.

4.2 Interaction Components

The interaction with the diagrams is essential for the user to explore the data and to verify or refute an initial hypothesis (R5).

In Figure 2 the data is visualized without any manipulations by the user. The first three diagrams (Figure 2.A,B,C&G) give the user an overview of the underlying data. To analyze the data further the user is able to manipulate the view of the data.

Details on Demand The visualization design enables the user to receive details of a visual encoding of an aggregation of a data attribute of a chart. The visual encoding of a number, for example the height of a bar in a bar chart or the length of an arc in a flow visualization, supports the user to compare the encoding with the same encoded data attributes. To receive exact numbers the user is able to hover over each visual encoded element and receive in place information with a tool-tip [Dix09].

Select Elements To receive even more detailed information about the highlighted visual element, the user is able to select it. The data is then filtered by the selection and all other visualizations are updated with the newly filtered data. This interaction method is implemented in the visualization design as simple left-mouse-click and works for every visualization. By clicking onto a data table row the entity is selected. The histograms (Figure 2.C&D) do not support a simple click operation, but a click and drag operation to select a one dimensional range of the attribute in the histogram [Mun14, ch. 11.4].

Highlight Elements To visualize which visual elements are selected, the color saturation of the visual element is increased and for the filtered elements the saturation is decreased. Figure 3 shows this difference in saturation in contrast to the not selected visual elements in Figure 2.A&B. The coloring of the small-multiple bar charts in the data table is also linked with the highlighting of the time bar chart. The used colors are selected using the ColorBrewer2 tool, which is based on evaluation of “385 unique colour schemes [...] across different computer platforms and monitors, [...] for possible colour-blind confusions, as well as in printed formats.” [HB03]

Sort To explore detail information for the trend over time, money, and the number of relations from one entity to another, the user is able to sort the data table along the data attribute of her/his interest (see Figure 2.I).

Search To support users’ who want to analyze the data for a specific entity, full-text search is integrated. In our visual design this is implemented as a simple form text fields for the legal entities and media (see Figure 2.H).

Combine and Remove Like already mentioned in Section 3.1 the data quality might not be optimal. As modifying the underlying data cannot be expected by the target user group, interactive visual editing should be possible (R3). In our prototype, users may remove entries and combine multiple entries into a single entry. With a click onto the labels above a data table the selected rows of that table are combined or removed.

4.3 Coordinating Multiple Views

The designed interface connects the different visualizations and widgets and organizes them. The views are arranged on fixed positions, but the user is able to filter the data [EB11, Rob07]. Because all visualizations of the media transparency database use the same data set it is possible to link the selection between all views and thus use each view for dynamic query [AS94, ST98]. Additionally the color of the visual elements indicate which aggregation is used. This helps the user to see the connection between the visualizations and it enables the user to understand the connection of a data attribute in one to the distribution of a data attribute in another diagram [Mun14, ch. 12].

5 Implementation

The visualization design has been implemented as a web-based software using JavaScript with the libraries D3.js [BOH11], Crossfilter [cro], and dc.js Dimensional Charting [dc].

The implementation is available from `https://github.com/VALIDproject/mtdb2` as free and open source software under a BSD-2-clause license and can be tested at `http://medientransparenz.validproject.at/dashboard/`. For iterative refinement an informal usability test with two subjects was conducted.

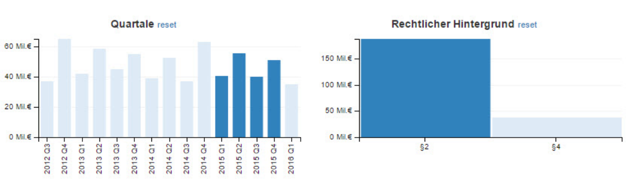

Figure 3: The elements of the visualizations adapt color saturation upon selection changes.

6 Usage Scenario

This sections presents a usage scenario to understand how the visualization of the MT dataset enables users to obtain a deeper insight into the data. The steps of the scenario can be followed in a video located at `https://vimeo.com/188278798`.

In most cases a user that is interested in a data set has some a-priori knowledge and a hypothesis that she/he wants to verify or falsify. In this scenario the user is interested in which legal entities spend money on online advertisement with Google.

Entering "google" into the full text search, the list returns 57 entries (e.g.: google, google.at, www.google.at, ...) due to data quality issues. By interactively manipulating the data the user is able to obtain a deeper insight. For example by combining the 57 categorical entries of the data attribute media to one entry named "Google". The flow visualization is now easier to read because the number of visual elements was reduced.

The user is able to filter data which she/he is not interested in to obtain new information. For example by selecting only the entries of universities. This results in 3 legal entities, which the user is now able to compare in more detail (see Figure 4).

7 Conclusions

This paper presented a visualization design to explore the MT dataset, a large open government data asset reporting on the flows of money from government to media. We implemented the design as a web-based prototype, made it publicly available, and showcased it on science communications events like Lange Nacht der Forschung. Based upon these experiences and informal feedback we received, we can now reflect how well the visualization design addresses its design requirements and provide directions for future research:

R1 Data scalability: The various views of aggregated attributes are useful to provide a big-picture overview of the dataset. Subsequently, the interaction concept of linked selection, sorting tables, and showing the first results works to learn about the details. Some users criticized the chord diagram as being too cluttered and hard to read. A Sankey diagram is being considered as alternative. In future work, two additional proposals from the preceding problem characterization study by Niederer et al. [NRA+16] can be adopted: The large number of legal entities and media could be automatically clustered into hierarchical groups using text or network analytics. Alternatively, supplementary data could be loaded to provide additional properties such as geographic area for legal entities and/or media. These properties

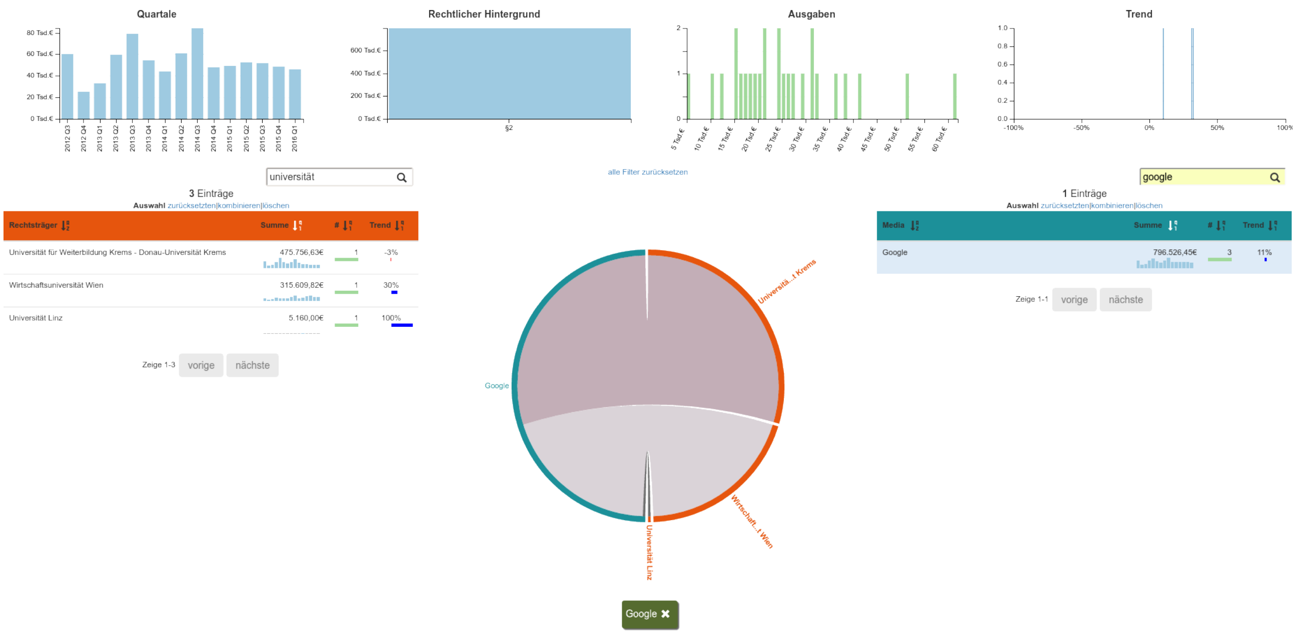

Figure 4: Snapshot of the visualization after following the steps of the usage scenario.

would subsequently be used for filtering and aggregation.

R2 Development over time: Both the bar chart view showing aggregated money flow over time, and the sparkline sized bar chart for each legal entity/medium work well to show distribution, abnormalities and other temporal patterns for the currently selected respectively visible items. The derived attribute "trend" was added to allow overview and direct manipulation of one concrete temporal pattern. While being a powerful feature, it is hard to grasp for novice users of the MT dataset visualization. Further design experiments are necessary to provide user-friendly exploration of temporal dynamic flows in bipartite networks.

R3 Data wrangling: The interactions to combine legal entities and/or media offers some benefits. The views are less cluttered by different entries for related institutions. In some cases data wrangling can eliminate a perceived false patterns such as abrupt end of flow to one medium that is in fact continued to a medium of a slightly different name.

Further work on data wrangling is indicated: On the one hand, we found the current functionality too limiting in several exploration sessions and desired more flexibility such as hierarchical groups and/or multi-group assignment like tags. On the other hand, the two functions 'combine' and 'remove' introduced more confusion for first time users. Possibly, a dedicated data wrangling mode should be provided so that these features are not visible by default.

R4 Ease of use: The visualization design is built using simple visual representation techniques that are well known to the general public. Still, the multiple views in composition were described as slightly overwhelming at first impression. In addition, novice users were not aware of direct manipulation so they did not expect they could filter the data e.g. by clicking on a bar.

R5 Interactive exploration: As demonstrated in the usage scenario, the visualization design allows free exploration of the MT dataset. While doing so, users can maintain overview of system state, i.e. which selections are active and also reset selections.

As further support for exploration, the data journalists interviewed by Niederer et al. [NRA+16] suggested documentation of the research path in order to provide analytic provenance [NCE+11].

Thus, our design study yielded not only a possible visualization design but also a range of directions for future work on exploring flows in dynamic bipartite networks.

Acknowledgements

This work was supported by the Austrian Ministry for Transport, Innovation and Technology (BMVIT) under the ICT of the future program via the VALiD project (no. 845598) and by the Austrian Science Fund (FWF) via the KAVA-Time project (no. P25489).

References

[AMST11] Wolfgang Aigner, Silvia Miksch, Heidrun Schumann, and Christian Tominski. *Visualization of Time-Oriented Data.* Springer, London, 2011.

[AS94] Christopher Ahlberg and Ben Shneiderman. Visual information seeking: Tight coupling of dynamic query filters with starfield displays. In *Proceedings of the SIGCHI conference on Human factors in computing systems: celebrating interdependence*, pages 313–317. ACM, 1994.

[Aus15] Julian Ausserhofer. "Die Methode liegt im Code": Routinen und digitale Methoden im Datenjournalismus. In Axel Maireder, Julian Ausserhofer, Christina Schumann, and Monika Taddicken, editors, *Digitale Methoden in der Kommunikationswissenschaft*, Digital Communication Research, pages 87–111. Berlin, 2015.

[BBDW16] Fabian Beck, Michael Burch, Stephan Diehl, and Daniel Weiskopf. A taxonomy and survey of dynamic graph visualization. *Computer Graphics Forum*, published online before print: 25 January, 2016. doi: 10.1111/cgf.12791.

[BKH+16] Benjamin Bach, Natalie Kerracher, Kyle Wm. Hall, Sheelagh Carpendale, Jessie Kennedy, and Nathalie Henry Riche. Telling stories about dynamic networks with graph comics. In *Proc. CHI Conf. Human Factors in Computing Systems*, pages 3670–3682. ACM, 2016.

[BOH11] M. Bostock, V. Ogievetsky, and Jeffrey Heer. D3: Data-driven documents. *IEEE Trans. Visualization and Computer Graphics*, 17(12):2301–2309, 2011.

[CMS99] Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman, editors. *Readings in Information Visualization: Using Vision to Think.* Morgan Kaufmann, San Francisco, 1999.

[cro] Crossfilter – fast multidimensional filtering for coordinated views. `http://square.github.io/crossfilter/`. Accessed: 2015-11-16.

[dc] dc.js – dimensional charting Javascript library. `https://dc-js.github.io/dc.js/`. Accessed: 2015-11-16.

[Dix09] Alan Dix. *Human-Computer Interaction.* Springer, 2009.

[EB11] Micheline Elias and Anastasia Bezerianos. Exploration views: understanding dashboard creation and customization for visualization novices. In *Proc. 13th IFIP TC 13 Int. Conf. Human-Computer Interaction, INTERACT*, pages 274–291. Springer, 2011.

[GBD09] Martin Greilich, Michael Burch, and Stephan Diehl. Visualizing the evolution of compound digraphs with TimeArcTrees. *Computer Graphics Forum*, 28(3):975–982, 2009.

[Ham] Markus Hametner. Inserate: 40,6 Millionen im ersten Quartal. `http://derstandard.at/2000017464403/`. Web-Standard: 2015-06-15.

[HB03] Mark Harrower and Cynthia A Brewer. ColorBrewer.org: An online tool for selecting colour schemes for maps. *The Cartographic Journal*, 40(1):27–37, 2003.

[HBW14] M. Hlawatsch, M. Burch, and D. Weiskopf. Visual adjacency lists for dynamic graphs. *IEEE Trans. Visualization and Computer Graphics*, 20(11):1590–1603, 2014.

[HSS15] Steffen Hadlak, Heidrun Schumann, and Hans-Jörg Schulz. A survey of multifaceted graph visualization. In Rita Borgo, F. Ganovelli, and Ivan Viola, editors, *Proc. Eurographics Conf. Visualization – State of The Art Report, EuroVis STAR*, pages 1–20. Eurographics, 2015.

[KSB+09] Martin Krzywinski, Jacqueline Schein, İnanç Birol, Joseph Connors, Randy Gascoyne, Doug Horsman, Steven J. Jones, and Marco A. Marra. Circos: An information aesthetic for comparative genomics. *Genome Research*, 19(9):1639–1645, 2009.

[Lan] Fabian Lang. Medientransparenz - die zweite. `http://www.paroli-magazin.at/555/`. Paroli-Magazin: 2013-03-18.

[Lor10] Mirko Lorenz. Status and outlook for data-driven journalism. In *Data-driven* *journalism: what is there to learn?*, pages 8–17. European Journalism Centre, 2010.

[Med15] Medienkooperations- und -förderungs-Transparenzgesetz, MedKF-TG. Austrian Legal Information System, 2015. `https://www.ris.bka.gv.at/Dokumente/Erv/ERV_2011_1_125/ERV_2011_1_125.html`. Accessed: 2016-08-31.

[Mun14] Tamara Munzner. *Visualization Analysis and Design*. CRC Press, 2014.

[NAR15] Christina Niederer, Wolfgang Aigner, and Alexander Rind. Survey on visualizing dynamic, weighted, and directed graphs in the context of data-driven journalism. In Hans-Jörg Schulz, Bodo Urban, and Uwe Freiherr von Lukas, editors, *Proc. Int. Summer School on Visual Computing*, pages 49–58. Frauenhofer Verlag, 2015.

[NCE+11] Chris North, Remco Chang, Alex Endert, Wenwen Dou, Richard May, Bill Pike, and Glenn Fink. Analytic provenance: process+interaction+insight. In *Proc. 2011 Ann. Conf. Ext. Abstracts Human Factors in Computing Systems, CHI EA*, pages 33–36. ACM, 2011.

[NRA+16] Christina Niederer, Alexander Rind, Wolfgang Aigner, Julian Ausserhofer, Robert Gutounig, and Michael Sedlmaier. Visual exploration of media transparency for data journalists: Problem characterization and abstraction. In *Proc. 10th Forschungsforum der österreichischen Fachhochschulen*. FH des BFI Wien, 2016.

[Rob07] Jonathan C. Roberts. State of the art: Coordinated & multiple views in exploratory visualization. In *Proc. Conf. Coordinated and Multiple Views in Exploratory Visualization, CMV*, pages 61–71, 2007.

[RTR] Austrian regulatory authority for broadcasting and telecommunications. Bekanntgegebene Daten. `https://www.rtr.at/de/m/veroeffentl_medkftg_daten`. Accessed: 2016-08-31.

[RTR16] RTR GmbH. Katalog Medientransparenz. Offene Daten Österreich, 2016. `https://www.data.gv.at/katalog/dataset/58e02823-2bd2-4db7-9e2f-72a9ea7c7ffd`. Accessed: 2016-08-31.

[SBSV] Peter Salhofer, Amir Basyouni, Mercedes Stibler, and Stephan Vrecer. Medientransparenz Austria. `http://www.medien-transparenz.at/`. Accessed: 2016-09-09.

[SMM12] Michael Sedlmair, Miriah Meyer, and Tamara Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Trans. Visualization and Computer Graphics*, 18(12):2431–2440, 2012.

[ST98] Robert Spence and Lisa Tweedie. The Attribute Explorer: Information synthesis via exploration. *Interacting with Computers*, 11(2):137–146, 1998.

[Tuf06] Edward Rolfe Tufte. *Beautiful Evidence*. Graphics Press, Cheshire, CT, 2006.

[vdEvW14] Stef van den Elzen and Jarke J. van Wijk. Multivariate network exploration and presentation: From detail to overview via selections and aggregations. *IEEE Trans. Visualization and Computer Graphics*, 20(12):2310–2319, 2014.

[vLKS+11] Tatiana von Landesberger, Arjan Kuijper, Tobias Schreck, Jörn Kohlhammer, Jarke J. van Wijk, Jean-Daniel Fekete, and Dieter W. Fellner. Visual analysis of large graphs: State-of-the-art and future research challenges. *Computer Graphics Forum*, 30(6):1719–1749, 2011.

[WPZ+16] Yanhong Wu, Naveen Pitipornvivat, Jian Zhao, Sixiao Yang, Guowei Huang, and Huamin Qu. egoSlider: Visual analysis of egocentric network evolution. *IEEE Trans. Visualization and Computer Graphics*, 22(1):260–269, 2016.

[ZGC+16] Jian Zhao, Michael Glueck, Fanny Chevalier, Yanhong Wu, and Azam Khan. Egocentric analysis of dynamic networks with EgoLines. In *Proc. CHI Conf. Human Factors in Computing Systems*, pages 5003–5014. ACM, 2016.