

Towards Scalable Visual Exploration of Very Large RDF Graphs^{*}

Nikos Bikakis^{1,2}, John Liagouris³, Maria Kromida¹,
George Papastefanatos², and Timos Sellis⁴

¹ NTU Athens, Greece ² ATHENA Research Center, Greece
³ ETH Zürich, Switzerland ⁴ RMIT University, Australia

Abstract. In this paper, we outline our work on developing a disk-based infrastructure for efficient visualization and graph exploration operations over very large graphs. The proposed platform, called graphVizdb, is based on a novel technique for indexing and storing the graph. Particularly, the graph layout is indexed with a spatial data structure, i.e., an R-tree, and stored in a database. In runtime, user operations are translated into efficient spatial operations (i.e., window queries) in the backend.

Keywords: graphVizdb, graph data, disk based visualization tool, RDF graph visualization, spatial, visualizing linked data, partition based graph layout.

1 Introduction

Data visualisation provides intuitive ways for information analysis, allowing users to infer correlations and causalities that are not always possible with traditional data mining techniques. The wide availability of vast amounts of graph-structured data, RDF in the case of the Data Web, demands for user-friendly methods and tools for data exploration and knowledge uptake. We consider some core challenges related on the management and visualization of very large RDF graphs; e.g., the Wikidata RDF graph has more than 300M nodes and edges.

First, their size exceeds the capabilities of memory-based layout techniques and libraries, enforcing disk-based implementations. Then, graph rendering is a time consuming process; even drawing a small part of the graph (containing a few hundreds of nodes) requires considerable time when we assume real-time systems. The same holds for graph interaction and navigation. Most operations, such as zoom in/out and move, are not easily implemented to large dense graphs, as their implementations require redrawing and re-layout large parts of them.

Related works in the field handle very large graphs through hierarchical visualization approaches. Although hierarchical approaches provide fancy visualizations with low memory requirements, their applicability is heavily based on the particular characteristics of the input dataset. In most cases, the hierarchy is constructed by exploiting clustering and partitioning methods [1,4,5,14,19]. In other works, the hierarchy is defined with hub-based [15] and density-based [22]

^{*} This paper appears in 12th Extended Semantic Web Conference (ESWC 2015).

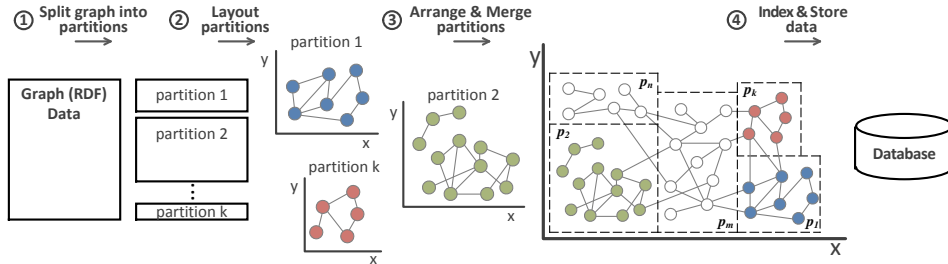


Fig. 1. Preprocessing Overview

techniques. [3] supports ad-hoc hierarchies which are manually defined by the users. Some of these systems offer a disk-based implementation [1,14,19] whereas others keep the whole graph in main memory [3,4,5,15,22].

In the context of the Web of Data [9,17,16,8,20,2,7], there is a large number of tools that visualize RDF graphs (adopting a node-link approach); the most notable ones are *ZoomRDF* [21], *Fenfire* [12], *LODWheel* [18], *RelFinder* [13] and *LODeX* [6]. All these tools require the whole graph to be loaded on the UI. Several tools that follow the same non-scalable approach have also been developed in the field of ontology visualization [10,11].

In contrast to all existing works, we introduce a generic platform, called *graphVizdb*, for scalable graph visualization that do not necessarily depend on the characteristics of the dataset. The efficiency of the proposed platform is based on a novel technique for indexing and storing the graph. The core idea is that in a preprocessing phase, the graph is drawn, using any of the existing graph layout algorithms. After drawing the graph, the coordinates assigned to its nodes (with respect to a Euclidean plane) are indexed with a spatial data structure, i.e., an R-tree, and stored in a database. In runtime, while the user is navigating over the graph, based on the coordinates, specific parts of the graph are retrieved and send to the user.

2 Platform Overview

The *graphVizdb* platform is built on top of two main concepts: (1) it is based on a “spatial-oriented” approach for graph visualization, similar to approaches followed in browsing maps; and (2) it adopts a disk-based implementation for supporting interaction with the graph, i.e., a database backend is used to index and store graph and visual information.

Partition-based graph layout. Here we outline the partition-based approach adopted by the *graphVizdb* in order to handle very large graph. Recall that, for graph layout, the graph is drawn once in a preprocessing phase, using any of the existing graph layout algorithms. However several graph layout algorithms require large amount of memory in order to draw very large graphs. In order to

overcome this problem, our partition-based approach (outlined in Figure 1) is described next.

(1) Initially, the graph (RDF) data is divided into a set of smaller sub-graphs (i.e., partitions) using a graph partitioning algorithm. At the same time, the graph partitioning algorithm tries to minimize the number of edges connecting nodes in different partitions. (2) Then, using a graph layout algorithm, each of the sub-graph resulted from the graph partitioning, is visualized into a Euclidean plane, excluding (i.e., not visualizing) the edges connecting nodes through different partitions (i.e., crossing edges). (3) The visualized partitions are organized and combined into a “global” plane using a greedy algorithm whose goal is twofold. First, it ensures that the distinct sub-graphs do not overlap on the plane, and at the same time it tries to minimize the total length of the crossing edges. (4) Based on the “global” plane, the coordinates for each node and edge are indexed and stored in the database.

Spatial operations for graph exploration. In `graphVizdb`, most of the user’s requests are translated into simple spatial operations evaluated over the database. In this context, *window queries* (i.e., spatial range queries that retrieve the information contained within a specific spatial region) are the core operation for most user’s requests. The user navigates on the graph by moving the viewing window. When the window is moved, its new coordinates with respect to the whole canvas are tracked on the client side, and a window query is sent to the server. The query is evaluated on the server using the R-tree indexes. This way, for each user request, `graphVizdb` efficiently renders only visible parts of the graph, minimizing in this way both backend-frontend communication cost as well as rendering and layout time. Additionally, more sophisticated operations, e.g., abstraction/enrichment zoom operations are also implemented using spatial operations.

Implementation. We have implemented a `graphVizdb` prototype¹ which provides interactive visualization over large graphs. The prototype offers three main operations: (1) interactive navigation, (2) multi-level exploration, and (3) keyword search. We use MySQL for data storing and indexing, the Jena framework for RDF data handling, Metis² for graph partitioning, and Graphviz³ for drawing the graph partitions. In the front-end, we use `mxgraph`⁴, a client-side JavaScript visualization library. A video presenting the basic functionality of our prototype is available at: vimeo.com/117547871.

References

1. J. Abello, F. van Ham, and N. Krishnan. ASK-GraphView: A Large Scale Graph Visualization System. *IEEE Trans. Vis. Comput. Graph.*, 12(5), 2006.

¹ graphvizdb.imis.athena-innovation.gr

² glaros.dtc.umn.edu/gkhome/views/metis

³ www.graphviz.org

⁴ www.jgraph.com

2. M. Alonen, T. Kauppinen, O. Suominen, and E. Hyvönen. Exploring the linked university data with visualization tools. In *ESWC*, 2013.
3. D. Archambault, T. Munzner, and D. Auber. GrouseFlocks: Steerable Exploration of Graph Hierarchy Space. *IEEE Trans. Vis. Comput. Graph.*, 14(4), 2008.
4. D. Auber. Tulip - A Huge Graph Visualization Framework. In *Graph Drawing Software*. 2004.
5. M. Bastian, S. Heymann, and M. Jacomy. Gephi: An Open Source Software for Exploring and Manipulating Networks. In *ICWSM*, 2009.
6. F. Benedetti, L. Po, and S. Bergamaschi. A Visual Summary for Linked Open Data sources. In *ISWC*, 2014.
7. N. Bikakis, M. Skourla, and G. Papastefanatos. rdf:SynopsViz - A Framework for Hierarchical Linked Data Visual Exploration and Analysis. In *ESWC*, 2014.
8. J. M. Brunetti, S. Auer, R. García, J. Klímek, and M. Necaský. Formal Linked Data Visualization Model. In *I1WAS*, 2013.
9. A. Dadzie and M. Rowe. Approaches to visualising Linked Data: A survey. *Semantic Web*, 2(2), 2011.
10. M. Dudás, O. Zamazal, and V. Svátek. Roadmapping and Navigating in the Ontology Visualization Landscape. In *EKAU*, 2014.
11. B. Fu, N. F. Noy, and M.-A. Storey. Eye Tracking the User Experience - An Evaluation of Ontology Visualization Techniques. *Semantic Web Journal*, 2015.
12. T. Hastrup, R. Cyganiak, and U. Bojars. Browsing Linked Data with Fenfire. In *WWW*, 2008.
13. P. Heim, S. Lohmann, and T. Stegemann. Interactive Relationship Discovery via the Semantic Web. In *ESWC*, 2010.
14. J. F. R. Jr., H. Tong, A. J. M. Traina, C. Faloutsos, and J. Leskovec. GMine: A System for Scalable, Interactive Graph Visualization and Mining. In *VLDB*, 2006.
15. Z. Lin, N. Cao, H. Tong, F. Wang, U. Kang, and D. H. P. Chau. Demonstrating Interactive Multi-resolution Large Graph Exploration. In *ICDM*, 2013.
16. S. Mazumdar, D. Petrelli, and F. Ciravegna. Exploring user and system requirements of linked data visualization through a visual dashboard approach. *Semantic Web*, 5(3), 2014.
17. S. Mazumdar, D. Petrelli, K. Elbedweihy, V. Lanfranchi, and F. Ciravegna. Affective graphs: The visual appeal of Linked Data. *Semantic Web*, 6(3), 2015.
18. M. Stühr, D. Roman, and D. Norheim. LODWheel - JavaScript-based Visualization of RDF Data. In *Workshop on Consuming Linked Data*, 2011.
19. C. Tominski, J. Abello, and H. Schumann. CGV - An interactive graph visualization system. *Computers & Graphics*, 33(6), 2009.
20. L. D. Vocht, A. Dimou, J. Breuer, M. V. Compennolle, R. Verborgh, E. Mannens, P. Mechant, and R. V. de Walle. A visual exploration workflow as enabler for the exploitation of linked open data. In *IESD*, 2014.
21. K. Zhang, H. Wang, D. T. Tran, and Y. Yu. ZoomRDF: semantic fisheye zooming on RDF data. In *WWW*, 2010.
22. M. Zinsmaier, U. Brandes, O. Deussen, and H. Strobel. Interactive Level-of-Detail Rendering of Large Graphs. *IEEE Trans. Vis. Comput. Graph.*, 18(12), 2012.

Acknowledgement. This research has been co-financed by the European Union (European Social Fund - ESF) and Greek national funds through the Operational Programs “Education and Lifelong Learning” - Funding Program: THALIS and “Competitiveness and Entrepreneurship” (OPCE II) - Funding Program: KRIPIS of the National Strategic Reference Framework (NSRF).