

graphVizdb: A Scalable Platform for Interactive Large Graph Visualization

Nikos Bikakis^{†‡} John Liagouris[#] Maria Krommyda[†]
George Papastefanatos[‡] Timos Sellis^{*}

[†]NTU Athens, Greece [‡]ATHENA Research Center, Greece

[#]Dept. of Computer Science, ETH Zürich, Switzerland ^{*}Swinburne University of Technology, Australia

Abstract—We present a novel platform for the interactive visualization of very large graphs. The platform enables the user to interact with the visualized graph in a way that is very similar to the exploration of maps at multiple levels. Our approach involves an offline preprocessing phase that builds the layout of the graph by assigning coordinates to its nodes with respect to a Euclidean plane. The respective points are indexed with a spatial data structure, i.e., an R-tree, and stored in a database. Multiple abstraction layers of the graph based on various criteria are also created offline, and they are indexed similarly so that the user can explore the dataset at different levels of granularity, depending on her particular needs. Then, our system translates user operations into simple and very efficient spatial operations (i.e., window queries) in the backend. This technique allows for a fine-grained access to very large graphs with extremely low latency and memory requirements and without compromising the functionality of the tool. Our web-based prototype supports three main operations: (1) interactive navigation, (2) multi-level exploration, and (3) keyword search on the graph metadata.

I. INTRODUCTION

Graph visualization is a core task in various applications such as scientific data management, social network analysis, and decision support systems. With the wide adoption of the RDF data model and the recent *Linked Open Data* initiative, graph data are almost everywhere [1]. Visualizing these data as graphs provides the non-experts with an intuitive means to explore the content of the data, identify interesting patterns, etc. Such operations require interactive visualizations (as opposed to a static image) in which graph elements are rendered as distinct visual objects; e.g., DOM objects in a web browser. This way, the user can manipulate the graph directly from the UI, e.g., click on a node or an edge to get additional information (metadata), highlight parts of the graph, etc. Given that graphs in many real-world scenarios are huge, the aforementioned visualizations pose significant technical challenges from a data management perspective.

First of all, the visualization must be feasible without the need to load the whole graph in main memory. These "holistic" approaches [2], [3] result in prohibitive memory requirements, and usually rely on dedicated client-server architectures which are not always affordable by enterprises, especially start-ups. Then, the visualization tool must ensure extremely low response time, even in multi-user environments built upon commodity machines with limited computational resources. Finally, the visualization must be flexible and meaningful to the user, allowing her to explore the graph in different ways

and at multiple levels of detail.

State-of-the-art works in the field [4]–[10] tackle with the previous problems through a hierarchical visualization approach. In a nutshell, hierarchical visualizations merge parts of the graph into abstract nodes (recursively) in order to create a tree-like structure of abstraction layers. This results in a decomposition of the graph into much smaller (nested) sub-graphs which can be separately visualized and explored in a "vertical" fashion, i.e., by clicking on an abstract node to retrieve the enclosed sub-graph of the lower layer. In most cases, the hierarchy is constructed by exploiting clustering and partitioning methods [2], [4], [6], [7], [9]. In other works, the hierarchy is defined with hub-based [8] and density-based [10] techniques. [5] supports ad-hoc hierarchies which are manually defined by the users. A different approach has been adopted in [11] where sampling techniques are exploited. Finally, in the context of the Web of Data, there is a large number of tools that visualize RDF graphs [12]; however, all these tools require the whole graph to be loaded on the UI. Although the hierarchical approaches provide fancy visualizations with low memory requirements, they do not support intuitive "horizontal" exploration (e.g., for following paths in the graph). Further, with hierarchical approaches it is not easy to explore dense parts of the graph in full detail (i.e., without using an abstract representation). Finally, the applicability of hierarchical approaches is heavily based on the particular characteristics of the dataset; for example, the existence of small and coherent clusters [4]–[7], [9] or the distribution of node degrees [8], [10].

Contribution. We introduce a generic platform for scalable multi-level visual exploration of large graphs. The proposed platform can easily support various visualizations, including all ad-hoc approaches in the literature, and bases its efficiency on a novel technique for indexing and storing the graph at multiple levels of abstraction. In particular, our approach involves an offline preprocessing phase that builds the layout of the input graph by assigning coordinates to its nodes with respect to a Euclidean plane. The same offline procedure is followed for all levels of abstraction. The respective points are then indexed with a spatial data structure (i.e., R-tree) and stored in a database. This way, our system maps user operations into efficient spatial operations (i.e., window queries) in the backend. The prototype we demonstrate here is a proof of concept that interactive visualizations can be effective on commodity hardware, still, allowing the user to perform intuitive navigations on the plane (e.g., follow paths in the graph) at




Fig. 1: Preprocessing Overview

any level of abstraction and regardless the size of the graph.

II. SYSTEM OVERVIEW

The architecture of our prototype, graphVizdb, consists of three main parts: (1) the *Client*, (2) the graphVizdb *Core* module, and (3) the *Database*. The *Client* is the frontend UI that offers several functionalities to the users, e.g., an interactive canvas, search features, multi-level exploration, etc. The *Core* module contains the *Preprocessing* module and the *Query Manager* that is responsible for the communication between the *Client* and the *Database*. The Preprocessing module contains the following submodules: *Partitioning*, *Layout*, *Partition Organizer* and *Abstraction*. Finally, the *Database* contains all data needed for the visualization along with the necessary indexes. Details for each part are provided in the following.

A. Preprocessing

In our approach, the layout of the input graph is built on the server side once, during the preprocessing phase, and this can be done with any of the existing layout algorithms. The result of this process is the assignment of coordinates to the nodes of the graph with respect to a Euclidean plane. The state-of-art layout algorithms provide layouts of high quality, however, they require large amounts of memory in practice, even for graphs with few thousands of nodes and edges. In order to overcome this problem, we adopt a partition-based approach as shown in Fig. 1.

Initially, the graph is divided by the *Partitioning* module into a set of k distinct sub-graphs (*Step 1*), where k is proportional to the total graph size and the available memory of the machine. This is a k -way partitioning that aims at minimizing the number of edges between the different sub-graphs [13]. Then, the *Layout* module applies the layout algorithm to each partition independently, and assigns coordinates to the nodes of each sub-graph without considering the edges that cross different partitions (*Step 2*). Any layout algorithm can be used in this step, e.g., circle, star, hierarchical, etc. The edges between the different sub-graphs are taken into account by the *Partition Organizer* when arranging the partitions on the "global" plane at *Step 3*. Multiple abstraction layers of the input graph are constructed by the *Abstraction* module at *Step 4*. At *Step 5*, the input graph along with the abstract graphs are indexed and stored in the *Database*. In the

following, we provide more details on *Step 3*, *4*, and *5* of Fig. 1.

Organizing Partitions. Partitions are organized on the "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 edges between different partitions (crossing edges).

Initially, the algorithm counts the number of crossing edges for each partition. Then, it selects the partition with the largest number of crossing edges (to all other partitions), and places it at the center of the plane, i.e., it updates the coordinates of its nodes with respect to the "global" plane. This is the m -th partition in Fig. 1 which has 9 such edges (denoted with red color). The remaining partitions are kept in a priority queue, sorted on the number of the common crossing edges they have with the partitions that exist on the plane (in descending order). At each subsequent step, the algorithm assigns the first partition from the queue to an empty area on the plane so that the total length of the crossing edges between this partition and all other partitions on the plane is minimized. Then, the partition is removed from the queue and the coordinates of its nodes are updated with respect to the assigned area. The order of the partitions in the queue is also updated accordingly, and the algorithm proceeds to the next step. The above process terminates when the priority queue is empty. Intuitively, the efficiency of the algorithm is guaranteed by the small number of partitions (k), and also by the small size of the area we have to check for the best assignment at each step; this area lies around the non-empty areas from the previous steps.

Building Abstraction Layers. After arranging the partitions, a number of abstraction layers is constructed for the initial graph, as shown in Fig. 1. A layer i ($i > 0$) corresponds to a new graph that is produced by applying an abstraction method to the graph at layer $i - 1$. Hence, the overall hierarchy of layers is constructed in a bottom-up fashion, starting from the initial graph at layer 0. Each time we create a new graph at layer i , its layout is based on the layout of the graph at layer $i - 1$. The abstraction method can be any algorithm that produces a more condense form of the input graph, either by merging parts of the graph into single nodes (like the graph summarization methods we mentioned in the introduction) or by filtering parts of the graph according to a metric, e.g., a node ranking criterion like PageRank. We emphasize that our

approach does not pose any restrictions to the number of layers or the size of the graph at each layer. Finally, all layers are kept as separate graphs in the database as we explain below.

Storage Scheme. Our database includes a single relational table per abstraction layer that stores all information about the graph of this layer. All these tables have the same schema as depicted in Fig. 2. Intuitively, each graph is stored as a set of triples of the form $(node_1, edge, node_2)$. A row in the table of Fig. 2 contains the following attributes: (1) the unique ID of the first node (Node₁ ID), (2) the label of the first node (Node₁ Label), (3) the geometry of the connecting edge (Edge Geometry) which is a binary object that represents the line between node₁ and node₂ on the plane, (4) the label of the edge (Edge Label), (5) the unique ID of the second node (Node₂ ID), and (6) the label of the second node (Node₂ Label). When the edge is directed, node₁ is always the source node whereas node₂ is the target node. This information is encoded in the binary object that represents the geometry of the edge.

B⁺-trees are built on attributes (1) and (5) to retrieve all information about a node efficiently. The full text indexes shown in Fig. 2 correspond to tries, and they are used to support fast keyword search on the graph metadata. Finally, an R-tree is used to index the geometries of the edges on the plane. Note that each such geometry is internally defined by the coordinates of the first and the second node whose IDs and labels are stored in the same row of the table.

index	B-tree	fulltext	R-tree	fulltext	B-tree	fulltext
attribute	Node ₁ ID	Node ₁ Label	Edge Geometry	Edge Label	Node ₂ ID	Node ₂ Label
type	int	text	geometry	text	int	text

Fig. 2: Storage Scheme

B. Online Operations

On the client side, our system provides three main visual operations:

Interactive Navigation. The user navigates on the graph by moving the viewing window ("horizontal" navigation). When the window is moved, its new coordinates with respect to the whole canvas are tracked on the client side, and a spatial range query (i.e., a window query) is sent to the server. This query retrieves all elements of the graph (nodes and edges) that overlap with the current window. The query is evaluated with a lookup in the R-tree of Fig. 2, and the respective part of the graph is fetched from the database and sent to the client. After the part of the graph is rendered on the canvas, the user can start the exploration.

Multi-level Exploration. The user moves up or down at different abstraction layers of the graph through a *Layer Panel* ("vertical" navigation). When changing a level of abstraction, the graph elements are fetched through spatial range queries on the appropriate table that corresponds to the selected layer. Vertical navigation can be combined with traditional zoom in/out operations in order to give the impression of a lower/higher perspective. In this case, the size of the window (rectangle) that is sent to the server is decreased/increased proportionally according to the zoom level.

Keyword-based Exploration. Finally, the user searches the graph using keywords through a *Search Panel*. In this case,

a keyword query is sent to the server and it is evaluated on the whole set of node labels which are indexed with tries. The result of this query is a list of nodes whose labels contain the given keyword. By clicking on a node from the list, the user's window focuses on the position of this node. In this case, the spatial query sent to the server uses as window the rectangle whose size is equal to the size of the client's window and whose center has the same coordinates with the selected node from the list.

III. IMPLEMENTATION & EVALUATION

This section provides information on the implementation of our prototype. It also includes the results of our experimental evaluation with two real graph datasets.

Implementation. graphVizdb¹ is implemented on top of several open-source tools and libraries. The *Core* module of our system is developed in Java 1.7, and the database we use is MySQL 5.6.12. The partitioning of the graph, during the preprocessing phase, is done with Metis 5.1.0 whereas the layout of each partition is built with Graphviz 2.38.0. The web-based frontend is entirely based on HTML and JavaScript. For the interactive visualization of the graph on the client side, we use mxGraph 3.1.2.1.

Web UI. The user interface consists of the following panels: (1) *Visualization*, i.e., the interactive canvas, (2) *Information* that provides information about a selected node (metadata), (3) *Control* that offers the basic functionality (e.g., select dataset, zoom, abstraction level/criterion), (4) *Birdview*, i.e., a large-scale image of the whole graph on the plane, (5) *Search* that provides keyword search functionalities, (6) *Statistics* that offers basic statistics for the graph (e.g., average node degree, density, etc.), (7) *Filter* that provides filtering operations on the canvas (i.e., hide edges/nodes), and (8) *Edit* that allows the user to store in the database the graph modifications made through the canvas.

Performance Evaluation. The experiments we present here were conducted on the Okeanos cloud using a VM with a quad-core CPU at 2GHz and 8GB of RAM running Linux. For the client application, we used Google Chrome on a laptop with an i7 CPU at 1.8GHz and 4GB of RAM. The cache size of MySQL on the server side was set to 6GB.

To evaluate the response time of our system, we used several real graph datasets with rather different characteristics. Due to lack of space, here we present only the results for two datasets: the *Wikidata* RDF dataset, and the *Patent* citation graph. The first one is an RDF export of Wikidata having 151M edges and 146M nodes. The second dataset is taken from the SNAP repository of large network datasets; it contains 16.5M edges and 3.8M nodes.

Table I presents the preprocessing time for each step of Fig. 1. These times are higher for Wikidata since it is much bigger than the Patent dataset. The only exception is the time spent in Step 1 for applying the *k-way* partitioning; this process takes longer for Patent due to the higher average node degree. Note that the most expensive part of the preprocessing is the indexing step; however, the presented times correspond to the total time spent in indexing 5 layers of each dataset, one after

¹graphvizdb.imis.athena-innovation.gr

TABLE I: Time for each Preprocessing Step (min)

Dataset	#Edges	#Nodes	Step 1	Step 2	Step 3	Step 4	Step 5
Wikidata	151M	146M	1.8	4.5	25.5	16.5	670.1
Patent	16.5M	3.8M	5.1	2.8	9.7	8.2	41.2

the other. In practice, we can speed up this step by distributing the layers to different nodes of the cluster and perform the indexing in parallel. In this case, the time spent in Step 5 equals the time for indexing the input graph (layer 0), that is, 274.5 and 17.4 minutes for Wikidata and Patent respectively.

Our experimental scenario includes the evaluation of window queries with different sizes. These queries are evaluated by the server and sent to the client for visualization. In particular, we used window queries whose size varies from 200^2 to 3000^2 pixels, and we evaluated them on the initial graph of each dataset, i.e., on the bottom layer of abstraction (layer 0). For each window size, we generated 100 random queries. The results we present in Fig. 3 correspond to the following average times per query (msec): (1) *DB Query Execution*: the time spent to evaluate the query in the database, (2) *Build JSON Objects*: the time required for the server to process the query result and build the JSON objects that are sent to the client, (3) *Communication + Rendering*: the time spent in the client-server communication plus the time needed to render the graph on the browser, and (4) *Total Time*: the sum of the above times. The *Nodes + Edges* in Fig. 3 refer to the average number of nodes and edges included in the 1K random windows of each size.

The first observation is that the performance of our approach scales linearly with the window size and the total number of objects in it. This behaviour is similar for both datasets. As we can see in Fig. 3, the overall response time of the system is dominated by the time spent in *Communication + Rendering*. We do not present these two operations separately because the part of the graph included in the window of the user is sent from the server to the client in small pieces, i.e., in a streaming fashion; hence, the respective times cannot be easily distinguished. As a final comment, the time spent to evaluate the query in the database is negligible and increases slightly as the size of the window increases.

IV. DEMONSTRATION OUTLINE

In this section, we outline our demonstration scenario. The attendees will first select a dataset from a number of real-world datasets (e.g., ACM, DBLP, DBpedia). Then, they will be able to have a quick glance on the graph using various navigation methods such as panning, selecting a specific part of the graph in the birdview panel, etc. Finally, the attendees will be able to filter (i.e., hide) edges and/or nodes of specific types (e.g., RDF literals), as well as to zoom in/out over the graph. For example, in the ACM dataset, a user interested in exploring the citations between articles will be able to filter out irrelevant edges (e.g., *has-author*, *has-title*, etc.) and visualize only the *cite* edges.

Additionally, the attendees will be able to explore the “Focus on node” mode, which is suitable for pathway navigation, as well as for helping users to further understand the relations amongst the nodes of interest. In this mode, only

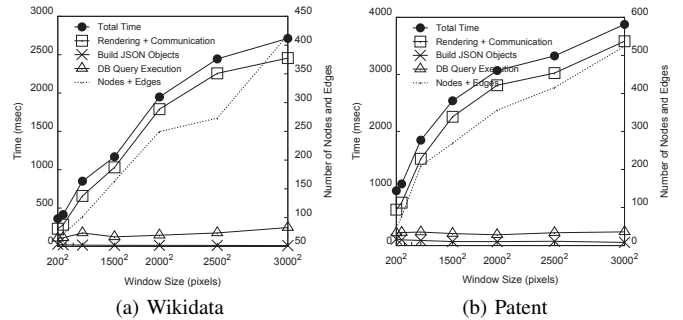


Fig. 3: Time vs. Window Size

the selected node and its neighbours are visible. The user interested in exploring the scientific collaborations of an author will be able to use keywords in order to search for this person, e.g., “Christos Faloutsos”. Then, using the “Focus on node”, the user can quickly explore all Faloutsos’ collaborations by following the “Christos Faloutsos · *has-author* · article · *has-author*” paths.

Beyond simple navigation, the attendees will be able to perform a multi-level graph exploration. In particular, they will be able to modify the abstraction level as well as the abstraction criteria (i.e., Node degree, PageRank, HITS). For example, by selecting either PageRank or HITS as the abstraction criterion in the Notre Dame web graph, the users will be able to view different layers of the graph that contain only the “important” nodes (e.g., sites whose PageRank score is above a threshold).

A video presenting the basic functionality of our prototype is available at: vimeo.com/117547871.

Acknowledgement. This work was partially supported by the EU project “SlideWiki” (688095).

REFERENCES

- [1] N. Bikakis *et al.*, “Towards Scalable Visual Exploration of Very Large RDF Graphs,” in *ESWC*, 2015.
- [2] M. Bastian, S. Heymann, and M. Jacomy, “Gephi: An Open Source Software for Exploring and Manipulating Networks,” in *ICWSM*, 2009.
- [3] T. Hastrup, R. Cyganiak, and U. Bojars, “Browsing Linked Data with Fenfire,” in *WWW*, 2008.
- [4] J. Abello, F. van Ham, and N. Krishnan, “ASK-GraphView: A Large Scale Graph Visualization System,” *TVCG*, vol. 12, no. 5, 2006.
- [5] D. Archambault, T. Munzner, and D. Auber, “GrouseFlocks: Steerable Exploration of Graph Hierarchy Space,” *TVCG*, vol. 14, no. 4, 2008.
- [6] D. Auber, “Tulip - A Huge Graph Visualization Framework,” in *Graph Drawing Software*, 2004.
- [7] J. F. Rodrigues *et al.*, “GMine: A System for Scalable, Interactive Graph Visualization and Mining,” in *VLDB*, 2006.
- [8] Z. Lin *et al.*, “Demonstrating Interactive Multi-resolution Large Graph Exploration,” in *ICDM*, 2013.
- [9] C. Tominski, J. Abello, and H. Schumann, “CGV - An interactive graph visualization system,” *Computers & Graphics*, vol. 33, no. 6, 2009.
- [10] M. Zinsmaier *et al.*, “Interactive Level-of-Detail Rendering of Large Graphs,” *IEEE Trans. Vis. Comput. Graph.*, vol. 18, no. 12, 2012.
- [11] S. Sundara *et al.*, “Visualizing large-scale RDF data using Subsets, Summaries, and Sampling in Oracle,” in *ICDE*, 2010.
- [12] N. Bikakis and T. Sellis, “Exploration and Visualization in the Web of Big Linked Data: A Survey of the State of the Art,” in *LWDM*, 2016.
- [13] G. Karypis and V. Kumar, “Multilevel Graph Partitioning Schemes,” in *ICPP*, 1995.