

Cobwebs from the Past and Present: Extracting Large Social Networks using Internet Archive Data

Miroslav Shaltev¹, Jan-Hendrik Zab¹ and Philipp Kemkes¹, Stefan Siersdorfer, Sergej Zerr²

¹L3S Research Center, Hannover, Germany
{shaltev,zab,kemkes}@L3S.de
siersdorfer@outlook.de

²Electronics and Computer Science, University of Southampton, Southampton, UK
s.zerr@soton.ac.uk

ABSTRACT

Social graph construction from various sources has been of interest to researchers due to its application potential and the broad range of technical challenges involved. The World Wide Web provides a huge amount of continuously updated data and information on a wide range of topics created by a variety of content providers, and makes the study of extracted people networks and their temporal evolution valuable for social as well as computer scientists. In this paper we present SocGraph - an extraction and exploration system for social relations from the content of around 2 billion web pages collected by the Internet Archive over the 17 years time period between 1996 and 2013. We describe methods for constructing large social graphs from extracted relations and introduce an interface to study their temporal evolution.

1. INTRODUCTION

The advances of the computer science and technology in the last decades enabled the extraction and analysis of social networks from various types of structured and unstructured data sources. Some sources provide explicit and easy to extract information about user relations. This includes on-line platforms such as Facebook, Twitter, or LinkedIn that maintain user databases and offer software interfaces for accessing contacts, friends, or followers. However, in many cases information about social connections is hidden within unstructured data such as Web pages and archives. In the past, personal relationships have been extracted from textual and multimedia sources such as books, historical repositories [3, 1, 9] and web search engines [5, 2, 6]. Despite the work towards increasing the efficiency of finding entity relations on the web [7] only recently in [8] a method suitable for mining of large graphs has been proposed.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '16, July 17-21, 2016, Pisa, Italy

© 2016 ACM. ISBN 978-1-4503-4069-4/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2911451.2911467>

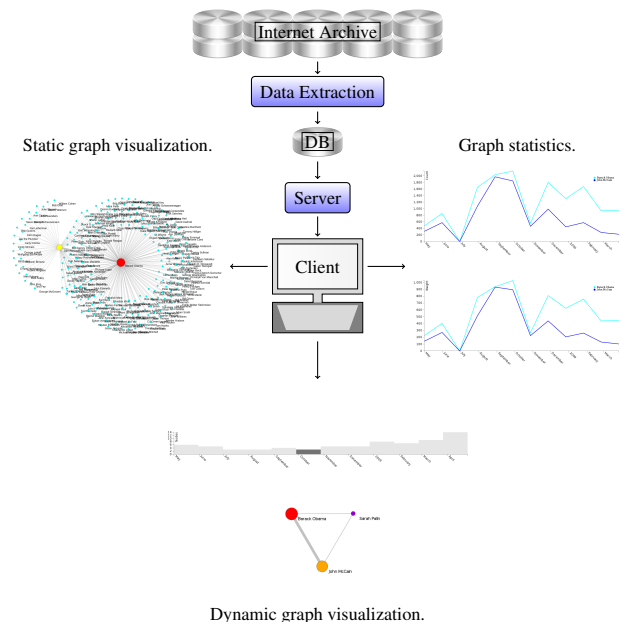


Figure 1: The overview over SocGraph system architecture and its main storage and visualization components.

In this work we introduce SocGraph¹ - a tool for the construction and analysis of social graphs extracted from the Internet Archive (IA) data and the exploration of the temporal evolution of communities in a variety of applications.

The objectives of the analysis include, but are not limited to shedding light on financial and business relations on the internet, identifying ad-hoc communities centered in social media on different events, such as solar eclipse and "Earth Hour" that are concerned with artificial light pollution of the night skies of our cities,² or describing and profiling of working groups in citizen science projects, where volunteers are supporting scientists by classifying astronomic and biological phenomena in raw data photographs. Inspired by [8] we move a step further in that direction and provide an interface for studying the *temporal evolution* of the social networks extracted directly from archived web page content. For social

¹<http://socgraph.l3s.uni-hannover.de>

²<http://www.stars4all.eu/>

and computer scientists our tool will provide a gateway to the information and knowledge about connections between people stored in the world wide web in the last few decades.

2. SYSTEM ARCHITECTURE

The architecture of our system is schematically shown in Fig. 1. First, we analyze the web archive to detect co-mentions of entities in the web pages. In the next step, at the server, we extract the temporal statistics and construct the social graph by connecting extracted entity pairs using detected edges.

Finally, the user can access the application and create, visualize, modify and interact with the graphs by issuing new queries via a web browser based user interface. In the following we provide an overview of the system components and show how results are presented to the user in more detail.

2.1 Data

The Internet Archive (IA) is a non-profit organization crawling the World Wide Web since 1996. For our application we have access to web pages from about 1.8×10^9 distinct URLs, collected by the IA in the time period of 17 years between 1996 and 2013. Extracting data even from a small sample of all archived web pages is a computationally intensive and requires parallelization of processes.

We used Hadoop and Spark [10] technology on a dedicated 25 node computing cluster with 1.3TB main memory and 268 CPU cores to extract co-mentions of persons directly from the archived documents. We stored extracted names, patterns, URLs, date of the crawl and additional miscellaneous information in a relational database, which can be accessed from the server process in real time. The crawled documents of MIME type text have been encapsulated in 346,000 Web ARChive (WARC) files.

2.2 Extraction of Entity Pairs

The entity pairs are extracted from the body of the archived documents first by splitting the documents into sentences using the Stanford CoreNLP library [4]. Then extracted sentences are scanned, detecting the constructs matching the template

$$\langle person1 \rangle \langle pattern \rangle \langle person2 \rangle$$

such as $\langle Barack Obama \rangle \langle and his rival \rangle \langle John McCain \rangle$, using a person names dictionary and a sliding window with a pattern length of three words. To reduce noise in the data we exclude pairs with identical names and discard overly long sentences and patterns. We computed the weight of the node as $1/n$, where n is the number of persons co-occurring in the sentence. In Fig.2a we plot the distribution of the weight per entry in the data set. The distribution of the number of extracted pairs per year is shown in Fig.2b.

2.3 Graph construction

For building the social graph, we select extracted pairs matching the user query and user defined parameters such as the time period T , and merge them into a network considering the node weight and the edge weight between two nodes as follows:

$$w = 1/n, \quad (1)$$

where n is the number of co-mentioned persons in a sentence.




Figure 2: The extracted data: (a) distribution of the weight per entry in the data set, (b) number of extracted edges per year.

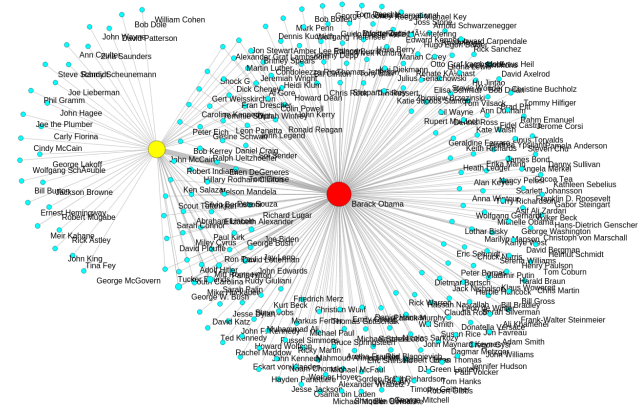


Figure 3: A static graph showing the extracted social network of Barack Obama and John McCain for the period May, 2008 to May, 2009.

We then define node weight \mathcal{W}_N over all considered entries i as

$$\mathcal{W}_N = \sum_i w_i \quad (2)$$

and the edge weight $\mathcal{W}_{E,kl}$ between two nodes k and l as

$$\mathcal{W}_{E,kl} = \sum_i w_{k,i} + w_{l,i}, \quad (3)$$

where $w_{k,i}$ and $w_{l,i}$ are the individual node weights. We refer to such graphs as static graphs. To study the temporal evolution of the social networks we also construct dynamic graphs, consisting of sequences of static graphs for disjoint intervals of length ΔT within the time period T .

2.4 Graphical user interface

The graphical user interface of SocGraph is accessible through a web browser, see Fig. 5. The ‘‘Person’’ input field is expecting a query such as a name of a person or a list of person names. The ‘‘Time span’’ selectors can be used to choose a time period of interest. Three filters are exposed to the user. It is possible (a) to limit the maximum number of considered database entries through providing the value in the ‘‘Maximal number of entries’’ input field, filter the nodes (b) by their weight through the adjustment of the ‘‘Weight per entry’’ slider and filter the edges (c) by their weight through the ‘‘Edge weight’’ slider.

The node weights are normalized in the range from zero to

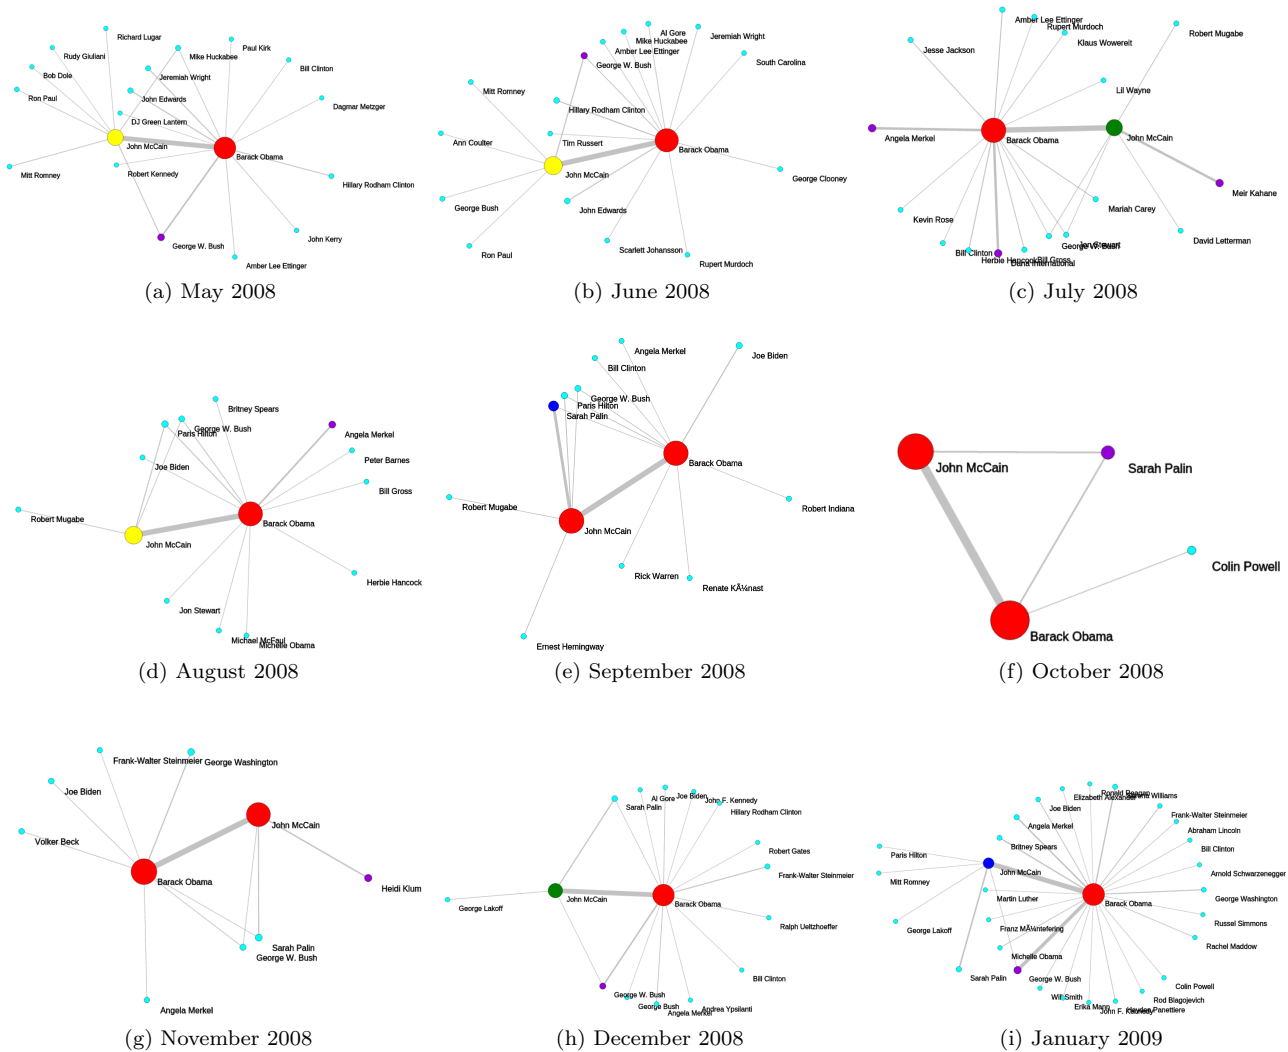


Figure 4: Temporal evolution of the social graph of Barack Obama and John McCain during the US election campaign 2008

one and a color scheme ranging from blue to red is utilized to indicate particular weight value from small to large, respectively. Overall, there are three modes of the SocGraph operation, namely:

- (i) A static graph can be computed and visualized over the complete dataset, providing an overview of all commentions for the requested persons available in the Internet Archive. This mode is triggered by the “Graph View” button in the “Update” field of the control panel on the right.
- (ii) For a specified sliding window a temporal statistical plot can be generated, showing the number of raw commentions as well as the weight of the persons of interest. This mode is triggered by the “Temporal statistic” button in the “Update” field of the control panel, the results are displayed in the “Temporal Statistic” tab of the demonstration interface.

- (iii) Finally, SocGraph provides a functionality to construct and visualize the individual graphs for particular time periods as well as to create an animation of graph evolution over time. The length of particular periods (measured in days) can be entered in the “Sliding window” input of the “Animation” field of the control panel. The user obtains the dynamic graph by clicking on the “Refresh” button (left to the “Play” button) of the player control strip. Once the data is loaded a time line with the number of edges in each sliding window period appears in the “Graph View” tab. To start the animation, the user should click the “Play” button. The animation speed can be controlled by the adjustment of the value (in seconds) in the “Speed” input. The playback can be paused with the “Pause button. Note, that the player also allows step by step forward and backward rendering of the graphs, triggered by the corresponding player control buttons.

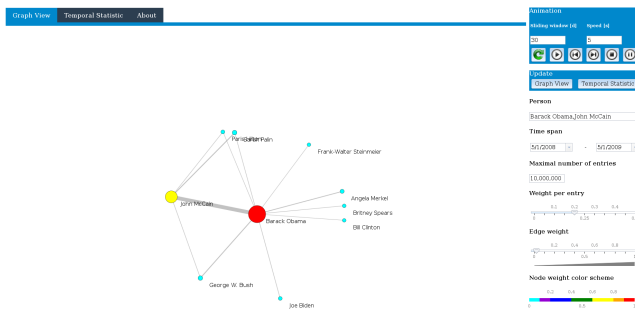


Figure 5: The graphical user interface of SocGraph as shown in a web browser.

3. DEMONSTRATION OVERVIEW

In the demonstration we will primarily show how the SocGraph time travel graph system works and how the social networks are constructed from the content of IA web pages. We will demonstrate the graphical interface usage for static and dynamic graph visualization. Additionally, we can elaborate in more detail on the person pair extraction process and explain the underlying parallelization algorithms.

We will explore the social networks of Barack Obama (node N_{BO}) and John McCain (node N_{MC}) for the one year period from May, 2008 to May, 2009, roughly corresponding to the US presidential election, as an example. The weight per entry has been set to 0.2, which means that we allow data records with up to four additional persons. The edge weight filter has been set to 0.025, in order to not overload the graphs and focus on the interesting entities. The static graph is plotted in Fig.3. A dynamic graph has been constructed with a sliding window of 30 days. The series of graphs are shown in Fig.4. For instance we observe that Hillary Clinton (N_{HC}) node was connected to (N_{BO}) from the beginning and disappeared in June 2008, corresponding to the time point, where Hillary Clinton endorsed Barack Obama and withdrew her candidacy. In following the sizes of both nodes, N_{BO} and N_{MC} , remain similar until November, where the actual election took place and N_{MC} drastically reduced the weight already in following month after Barack Obama became a president of the United States. Shortly before in September, Sarah Palins' (N_{SP}) talk on the side of John McCain had positive impact on the votes for this candidate. This fact is also reflected in our graph where N_{SP} appeared in September connected to N_{MC} .

4. DISCUSSION

In this paper we introduced a demonstration of SocGraph - a social graph extracting system for large networks from Internet Archive data. In contrast to other research concerned with graph construction from web related data, we are focused on the temporal evolution of social networks implicitly contained in the stored web pages. In our future work we plan to include pattern filtering techniques, integrate data from search engines and evaluate event identification, as well as sentiment analysis of the personal relationships and their evolution. Our system will be offered as a service within the EU Project Alexandria as part of a Web Observatory accessible to social and computer scientists as well as to general public for social network visualization and evolution analysis. The demonstration is available on our web

page: <http://socgraph.l3s.uni-hannover.de> and can be used with any web browser. Additionally, the web page contains a summary of the demo applications as well as a short video tutorial.

5. ACKNOWLEDGMENTS

This work is partly funded by the European Research Council under ALEXANDRIA (ERC 339233), by the European Commission under grant agreements 619525 (QualiMaster) and 688135 (STARS4ALL).

6. REFERENCES

- [1] C. Bird, A. Gourley, P. Devanbu, M. Gertz, and A. Swaminathan. Mining email social networks. In *MSR Workshop 2006*.
- [2] X. Canaleta, P. Ros, A. Vallejo, D. Vernet, and A. Zaballos. A system to extract social networks based on the processing of information obtained from internet. In *ICC Association for Artificial Intelligence 2008*.
- [3] D. K. Elson, N. Dames, and K. R. McKeown. Extracting social networks from literary fiction. In *Association for Computational Linguistics (ACL) System Demonstrations, 2010*.
- [4] C. D. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. J. Bethard, and D. McClosky. The Stanford CoreNLP natural language processing toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60, 2014.
- [5] Y. Matsuo, J. Mori, M. Hamasaki, K. Ishida, T. Nishimura, H. Takeda, K. Hasida, and M. Ishizuka. Polyphonet: An advanced social network extraction system from the web. In *Proceedings of World Wide Web Conference 2006, Semantic Web Track*.
- [6] M. K. M. Nasution and S. A. Noah. Superficial method for extracting social network for academics using web snippets. In *RSKT '10*.
- [7] R. Nuray-Turan, Z. Chen, D. V. Kalashnikov, and S. Mehrotra. Exploiting web querying for web people search in weps2. In *Web People Search Evaluation Workshop (WePS 2009)*.
- [8] S. Siersdorfer, P. Kemkes, H. Ackermann, and S. Zerr. Who with whom and how?: Extracting large social networks using search engines. In *CIKM '15*, pages 1491–1500. ACM.
- [9] L. Wieneke, M. Düring, G. Sillaume, C. Lallemand, V. Croce, M. Lazzaro, F. S. Nucci, C. Pasini, P. Fraternali, M. Tagliasacchi, M. Melenhorst, J. Novak, I. Micheel, E. Harloff, and J. G. Moron. Building the social graph of the history of european integration - A pipeline for humanist-machine interaction in the digital humanities. In *SocInfo 2013 International Workshops*.
- [10] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: Cluster computing with working sets. In *Proceedings of the 2Nd USENIX Conference on Hot Topics in Cloud Computing, HotCloud'10*, pages 10–10, Berkeley, CA, USA, 2010. USENIX Association.