

The Open Graph Archive: A Community-Driven Effort^{*}

Christian Bachmaier¹, Franz J. Brandenburg¹, Philip Effinger²,
Carsten Gutwenger³, Jyrki Katajainen⁴, Karsten Klein³, Miro Spönemann⁵,
Matthias Stegmaier², and Michael Wybrow⁶

¹ University of Passau, Germany

² Eberhard-Karls-Universität Tübingen, Germany

³ Technische Universität Dortmund, Germany

⁴ University of Copenhagen, Denmark

⁵ Christian-Albrechts-Universität zu Kiel, Germany

⁶ Monash University, Australia

1 Introduction

A *graphbase*, a term coined by Knuth [7], is a database of graphs and computer programs that generate, analyze, manipulate, and visualize graphs. The terms *graph library* and *graph archive* are often used as synonyms for this term. Our vision is to provide an infrastructure and quality standards for a public graphbase, named the Open Graph Archive, that is accessible to researchers and other interested parties around the world via the worldwide web. This paper describes the current work undertaken towards this goal; the paper is also intended to be a call for participation since this will be a community-driven effort where most of the content will be provided by users of the system.

Our motives for building this universal graphbase are similar to Knuth’s motives for building the Stanford GraphBase [7]; we are just working on a larger scale. First, we want to provide standard sets of graphs to enable repeatability of experiments. We expect that the graphbase would be particularly interesting for researchers working in the areas of algorithm engineering and graph drawing. Second, we want to provide a single point of access for datasets relevant for people working with graphs. By annotating the datasets with their origin and other semantic information, we can help researchers to find publications relevant for their work. Third, a graphbase that is accessible worldwide can stimulate interesting theory development. As pointed out by Knuth [7], a graphbase can bridge the gap between theoreticians and practitioners. Fourth, the programs (and maybe also the datasets) available in a graphbase, if done well, can have a significant educational value.

Many existing collections, like the graphs available in the Stanford GraphBase [7] and the well-known Rome graphs [4], are static and only cover a small number of data sizes, types, and properties that may be relevant for the users.

^{*} This work was initiated at Schloss Dagstuhl in seminar 11191 on “Graph Drawing with Algorithm Engineering Methods”.

In order to allow collection and exchange of interesting graphs, it is important to make the graphbase extendable. The needs of the community will certainly change over time. Expandability has been recognized as an important goal by other researchers as well (see, e.g., [1,2]), but the available data collections seem to be relevant to a limited range of users only. Our goal is to support the use in a wide variety of application areas.

2 User Needs and Requirements

In order to investigate the relevance of and the requirements for a universal graphbase we conducted a survey among 30 participants of the Dagstuhl seminar 11191, coming from the graph drawing and algorithm engineering communities. The survey solicited a variety of open-ended textual responses. In this section we summarize the most interesting and commonly recurring feedback.

Describe two most important use cases for a graph archive. The most frequent use cases were to search for graphs with specific properties, and to benchmark and compare algorithms, both mentioned by 37% of the participants. Further answers were to share datasets (27%), and to replicate experiments and compare results (23%). Since these are fundamental aspects of experimental scientific processes, we can see that a graphbase would be an important tool for researchers of graph algorithms.

What services do you expect? We proposed nine services of which the survey participants could select those they considered important. As shown in Table 1, support for tags and arbitrary comments are the most crucial features. When asked for further important services, a handful of people wanted to know which publications refer to a specific graph or collection of graphs (17%).

Which category tags and analysis properties may be useful? Participants named 20 different application domains to categorize a graph or collection of graphs, e.g., biology, social networks, geography, software engineering. Furthermore, participants named 16 graph properties, most of which can be determined automatically. The most popular properties were connectivity (60%), including the num-

Add categorization tags	80%
Add comments, links, or further information	77%
Search for specific tags	77%
Automatic conversion of file formats	70%
Search for specific properties	60%
Add information on how graphs were created	60%
Add images (drawings of the graphs)	50%
Automatic analysis of graph properties	47%
Programmatic web service	23%

Table 1. Result of the multiple-choice question “Which services do you consider as critical for a graph archive?”

ber of k -connected components, and planarity (43%), including the best known crossing number for non-planar graphs.

Name two file formats you use most. The most frequently mentioned formats were GraphML (43%) and GML (33%). Since a total of 13 different formats were named, it is evident that a universal graphbase should not rely on one specific format, but offer support for several formats, preferably even converting automatically between formats.

Existing archives and collections. Responses for existing archives showed that GraphArchive [6] from the University of Tübingen and the datasets from the DIMACS implementation challenges [5] were both known by a handful of people (20% and 13%, respectively). These numbers are quite low and might also be biased towards the archives used by the researchers that participated in the seminar. They also suggest that there is currently no commonly used and accepted graph archive service. Regarding graph collections, participants mostly worked with randomly generated graphs, as well as with the popular Rome [4] and AT&T graphs [3].

Community contributions. Several participants of the survey declared that they would be willing to provide human resources (students, testing and development time), a hardware platform, or even money. This reaffirms that there is definite interest and enthusiasm for such a system, and also that the project should take advantage of this through involvement of the community.

Technical and service requirements. The survey results and subsequent community discussions indicate that potential users agree on a core set of important features, as well as a larger list of desirable functionality. However, several questions regarding the interface, architecture, and content remain open. Below we list the most relevant issues that need to be discussed or dealt with.

Storage. Graphs must be stored persistently under a unique ID for identification and access. Should graphs be stored in their original submission format, or converted by the system or the user into a unique storage format? In file conversions it is important that as much information as possible is preserved.

Metadata. There is a variety of metadata that can be stored with a graph, e.g., creator, description of the underlying data or the generator, additional keywords, and links to corresponding experiments or publications. Some of this data should be defined as mandatory properties, whereas other parts may be added as generic text properties. Useful keywords/tags for categorization need to be defined. Some tags could be attributes for graphs or collections of graphs, and some could list their structural and semantic properties.

Searching. Based on the survey results and our own experience, we assume that a graphbase should allow the user to search using both graph properties (number of nodes, etc.) and annotations (categories, origin, etc.).

Data analysis. Automatic analysis of basic graph properties must be possible. However, we are not sure if there should be a restriction on the computational complexity of the analysis or on the size of the analyzed graphs, or if users should be allowed to upload that information, e.g., the crossing number of a graph.

Programs. In addition to datasets, it must be possible to store programs like graph generators, analyzers, or visualizers. If the graphbase contains randomly generated collections of graphs with certain attributes, it would be useful to provide access to the programs used for their generation.

Ownership and copyright. The ownership of uploaded graphs must be clear from the outset. The content should be as freely usable as possible with fair attribution to the original authors. Contributors will need to take responsibility for their submitted graphs and collections of graphs.

Existing collections. Existing popular collections should be made identifiable and accessible via the system.

Possible extensions. Further useful extensions may include the following:

- Automatic file conversion could be provided as an additional service and the programs providing these conversions could also be made available.
- A series of drawings (layouts) for submitted graphs could be provided, or even automatic layout on demand, and the programs used for drawing the graphs could be made publicly available.
- Special support for browsing collections of graphs could be provided. For this purpose a hierarchical classification system can be useful.
- Structure-based searching could be supported, e.g., find graphs containing a clique of a specific size.
- Versioning of individual graphs as well as the possibility to store a series of dynamic graphs could be supported.
- A web-service API could be provided to allow interrogation of the graphbase by computer programs, rather than via a web browser.

3 A Working Prototype: GraphArchive

In this section we give an overview of *GraphArchive*, a platform for exchanging and archiving graphs meant as a prototype for the Open Graph Archive. It is developed at the University of Tübingen and was designed as a successor to *GraphDB*, a now discontinued first attempt at creating a web-based graphbase. GraphArchive is an interactive online system built with modern web technologies. Below we list the main features of the existing prototype, followed by a short description of its software architecture. For more details, we refer to [6]. The working system can be accessed online at

<http://graphdrawing.org/grapharchive/>.

Main features. The features of GraphArchive, as listed below, have been chosen to support the goal of providing an open and easily accessible system.

Web-based user interface. The user interface is provided via a browser. A web portal offers all functionality that is needed to handle graphs, including uploading datasets, inspection and management of existing graphs, searching for specific graphs, and downloading datasets. Registration is performed online using a registration form, which is processed automatically. Standard techniques are used to prevent registration by spam bots.

Minimal permission management. There are no groups of users that define rights for small circles of users. Licenses for graphs limiting their usage are not encouraged in our open approach. However, if necessary, a license can be attached to a selected graph. After confirming registration by going through the opt-in e-mail process, a user has access to all graphs and can initiate queries without restrictions.

Categorization of graphs. For search queries, graphs can be assigned to the field(s) of application that they originate from. This enables researchers from different fields to use GraphArchive as a common platform.

Automatic graph analysis. After upload, graphs (with $< 100,000$ vertices) are automatically analyzed in order to provide consistent data. Consistency is very important for queries on graph properties.

Multi-criterion search. Queries can be performed on multiple parameters, specifying graph properties, categories, author, name, and upload date. Also, parameters can be added later to further narrow down the result set.

Graph visualization. An image of a graph is valuable if a user wants to visually inspect the properties of a graph. Layouts are computed automatically in the background and can also be changed after upload.

Unique links to graphs. A URI associated with each graph allows for a permanent reference to be used in publications. By giving the URI, the user can quickly jump to a particular dataset. Reference annotations can be assigned to a graph in order to highlight publications and/or websites that refer to or make use of the graph.

Visual comparison of multiple graphs. For a quick comparison of graphs, we support simultaneous presentation of multiple graphs. Properties are displayed for all graphs. Boolean properties, e.g., directed/undirected, are presented visually on a scale (properties can be shared by (a) no graph, (b) a subset of the displayed graphs, or (c) all graphs).

Several file formats. When supporting many application domains it is impossible to dictate the file format used. Therefore, we aim at supporting as many formats as possible. The system is extendable and allows for addition of further formats in the future. For downloading graphs, a user can choose the format that fits best to his or her work environment. We provide cross conversion between different formats (the users can select any supported format and the system performs the conversion automatically).

Import/export of multiple graphs. We allow upload/download of several graphs simultaneously in zip-compressed form. In an upload process, each file in a compressed archive can be optionally processed individually (for property analysis and layout computation).

Guest access for non-registered users. If a user wants to check a specific graph, he or she can access a detailed view of the graph using its URI. All properties and attributes of that graph are made visible via a guest account.

Software architecture. GraphArchive is built with common web technologies. The application is written in PHP5⁷ and uses Apache2⁸ for online presentation. For graph analysis and layout computation, we use the Java graph library yFiles;⁹ these computations are handled in the background via PHP/Java bridge.¹⁰ Data storage is managed by PostgreSQL database management system.¹¹

More details and a descriptive walk-through showing a typical use case of the system can be found in [6]. For more news and information on the system and its current development status, please consult the system website.

4 Outlook

Our hope is to stimulate discussion on the initial system proposal and trigger community growth around the Open Graph Archive. The success of this project requires a passionate and enthusiastic community. We urge you to step up and participate by critiquing the existing system, helping the development effort, or contributing material to the graphbase.

References

1. R. F. Boisvert, R. Pozo, K. Remington, R. Barrett, and J. J. Dongarra. The matrix market: A web resource for test matrix collections. In *Quality of Numerical Software: Assessment and Enhancement*, volume 76 of *IFIP Conference Series*, pages 125–137. Chapman & Hall, 1997. Graphs available at <http://math.nist.gov/MatrixMarket>.
2. T. A. Davis and Y. Hu. The University of Florida sparse matrix collection. *ACM Trans. Math. Software*, 38(1), 2011. Graphs available at <http://www.cise.ufl.edu/research/sparse/matrices/>.
3. G. Di Battista, A. Garg, G. Liotta, A. Parise, , R. Tamassia, E. Tassinari, F. Vargiu, and L. Vismara. Drawing directed acyclic graphs: An experimental study. *Internat. J. Comput. Geom. Appl.*, 10(6):623–648, 2000. Graphs available at <http://www.graphdrawing.org/>.
4. G. Di Battista, A. Garg, R. Tamassia, E. Tassinari, and F. Vargiu. An experimental comparison of four graph drawing algorithms. *Comput. Geom.*, 7(5–6):303–325, 1997. Graphs available at <http://www.graphdrawing.org/>.
5. 10th DIMACS implementation challenge: Graph partitioning and graph clustering. <http://www.cc.gatech.edu/dimacs10/downloads.shtml>, accessed August 2011.
6. P. Effinger, M. Kaufmann, S. Meinert, and M. Stegmaier. GraphArchive: An online graph data store. Technical Report WSI-2011-03, Wilhelm-Schickard-Institut, Eberhard-Karls-Universität Tübingen, 2011.
7. D. Knuth. *The Stanford GraphBase: A Platform for Combinatorial Computing*. ACM Press, 1994.

⁷ See project homepage: <http://www.php.net>, accessed July 2011.

⁸ See project homepage: <http://www.apache.org>, accessed July 2011.

⁹ Developed by yWorks GmbH: <http://www.yworks.com>, accessed July 2011.

¹⁰ See project homepage: <http://php-java-bridge.sourceforge.net/pjb/index.php>, accessed July 2011.

¹¹ See project homepage: <http://www.postgresql.org/>, accessed July 2011.