

# Graph-based Analysis of Dynamic Systems

Benjamin Schiller

October 27, 2016

An ever increasing demand exists for the timely analysis of large dynamic systems in order to monitor, evaluate, and improve them. These systems are characterized by their constantly changing properties and the need for their subsequent analysis. The in-depth analysis of systems like biological, computer, and traffic networks promises great insights into their working principles, key properties, and the time-dependent changes thereof. It allows us to monitor the development of fundamental characteristics over time and thereby to understand how a dynamic system works. Using the results of an analysis, we can judge the impact of parameters on a system's properties and performance. These insights assist the development of new systems and the optimization of existing ones. The detailed analysis of a system allows us to identify components that are important to a specific task. The timely analysis of dynamic systems enables us to quickly react to anomalies and thereby guarantee the correct functionality of a system.

Examples of dynamic systems can be found in many areas such as computer networks, power grid management, traffic systems, social networks, and biological systems. The identification of vulnerable points in computer networks, or power grids, enables us to protect them against random failures and targeted attacks [19, 17, 3, 2, 10]. The surveillance of air as well as road traffic is a great asset in the prevention of delays and traffic jams [22, 28]. Analyzing similarities of users in constantly changing online social networks like Facebook, LinkedIn, and Twitter enables intricate advertisement schemes and is an important resource for the service providers [18]. The stability and heat resistance of a protein can be improved by identifying relevant interactions between amino acids and strengthening their corresponding bonds [14, 6, 24].

Analyzing a system means to determine its characteristics. Analyzing a dynamic system means to track the change of its characteristics over time. A general workflow for this process is illustrated in Figure 1. The state of a dynamic system, which changes over time, is observed and subsequently analyzed at certain points in time. The interpretation of the analysis results provides insights into the development of the corresponding system over time. The question which properties are relevant, therefore depends on the system itself and the analyst's objectives.

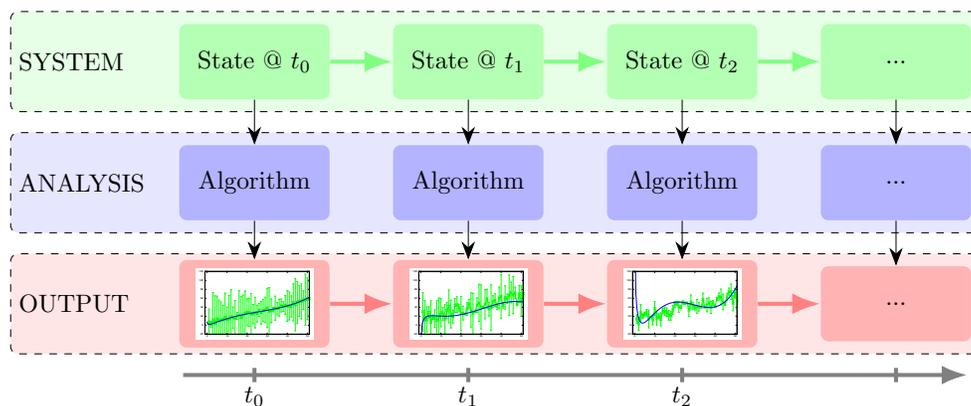


Figure 1: General process of the *Analysis of Dynamic Systems*

The analysis of a dynamic system is commonly performed with some frequency, e.g., every second, every minute, or every 24 hours. The system's state is analyzed for every corresponding point in time and a set of properties is output as a result. The duration between the analyzed points in time depends on the analysis-frequency. A high frequency results in the output of more data points in a time interval compared to an analysis at low frequency.

The separate computation of relevant properties for each point in time is the basic way to analyze a dynamic system. The analysis takes the complete state as input and computes the properties of interest for each point in time. Increasing the frequency of an analysis therefore leads to the frequent re-computation of all properties based on the complete state. The performance of this approach does not scale well with an increasing frequency and is therefore not a good fit for analyses that require the frequent output of analysis results.

The event-based maintenance of a system’s properties is another approach to analyze dynamic systems. Every event that occurs in the analyzed system is used to trigger an analysis that updates the current properties based on the implied change to the system. It takes advantage of the already computed properties and restricts the analysis to the single event. This approach is better suited in case an analysis at a high frequency is desired where only few changes occur between two points in time.

The frequency of an analysis has a high impact on its ability to describe the characteristics of a dynamic system accurately. Lowering the frequency of an analysis leads to the loss of information because fewer data points are output as results. This can be irrelevant in case a system’s properties change rarely and steadily, like the room temperature. In case a system’s properties change frequently and briefly, the expressiveness of analysis results decreases with the frequency. Interesting incidents may remain unobserved or be noticed too late, like the sudden increase of incoming traffic during a Denial-of-Service attack in a computer network. Therefore, it is commonly desirable to analyze a dynamic system at a high frequency.

The analysis results for a biological system at different frequencies are shown in Figure 2 as an example<sup>1</sup>. At the highest frequency, 1024 states of the system are analyzed while only 512, 256, 128, and 64 data points are contained in the results of the corresponding analyses with lower frequencies. Even for the second highest frequency, which outputs 512 data points, certain peaks are missed. The 64 data points, which result from the analysis at the lowest frequency, fail to depict the actual changes that occur in the system over time.

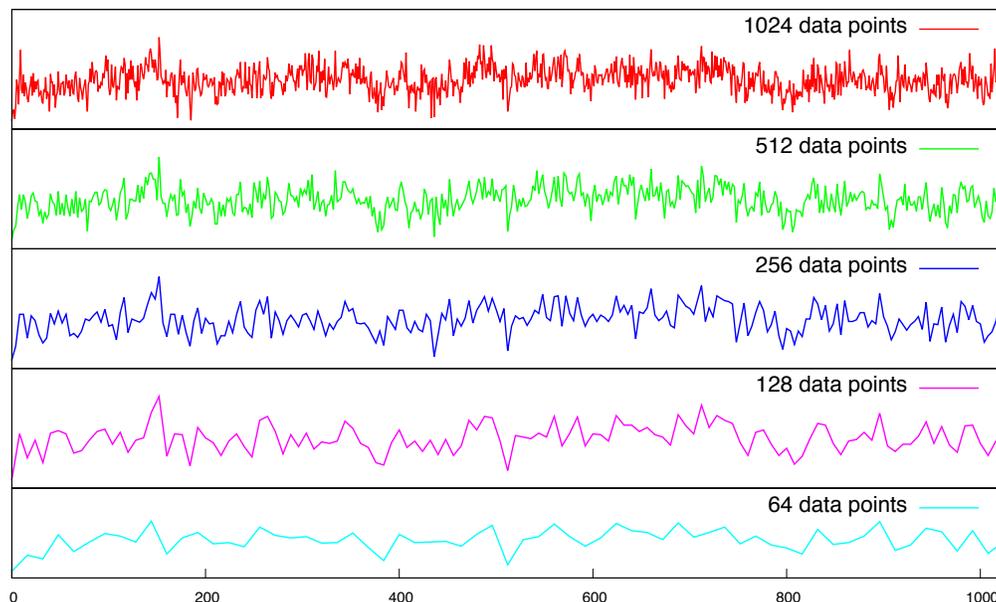


Figure 2: Results of an analysis performed at different frequencies

The characteristics of many systems are best described by the relationships between their components. Therefore, graphs are often used to model systems from various fields, including biology [1, 21], chemistry [20, 4], computer networks [13, 32], transportation networks [16], social sciences [8], and online social networks [27, 29]. The components of these systems are represented as vertices of a graph while their relations are modeled as edges that connect the corresponding vertices. These systems change over time as new components join and existing ones leave. Relations can be created and dissolved over time as well.

As examples, consider online social networks and air traffic systems. Users of an online social network are represented as vertices while their friendship relations are expressed as edges between them. Airports,

<sup>1</sup>Here, the frequencies of a single subgraph over time in the graph model of a molecular dynamics trajectory of the enzyme *para Nitro Butyrate Esterase-13* are shown. At the highest frequency of 2.5 ps, the system is analyzed at 1024 points in time. The lower analysis frequencies of 5 ps, 10 ps, 20 ps, or 40 ps result in the output of fewer data points.

the components of air traffic systems, are modeled as vertices, which are connected by edges that represent connecting flights. The creation of a new user profile in an online social network leads to the addition of a vertex to the graph while the shutdown of an airport results in the removal of the corresponding vertex. In case a new friendship relation is established between two users, the corresponding vertices are connected by an edge while the cancellation of the route between two airports leads to the removal of the edge connecting their respective vertices in the graph.

A large number of graph measures has been developed to analyze graphs modeled from arbitrary systems [12, 27]. Some measures compute global properties for the whole graph like, e.g., global clustering coefficient [23], degree distribution [11], and rich-club coefficient [33]. They reveal properties of an entire graph and thereby relate to overall system properties. Other measures express properties for each vertex like, e.g., local clustering coefficient [30], betweenness centrality [5], and PageRank [26]. Their values can be directly related to the corresponding component of the modeled system and reflect their characteristics.

The representation of a dynamic system as a graph and its subsequent analysis is a great way to utilize the expressiveness of graphs. The meaning and interpretation of analysis results depends on the analyzed system and the model used to represent it as a graph. It is therefore crucial to first determine the system’s characteristics of interest. A system can then be modeled as a graph to reflect the desired properties. The graph measures that should be computed during analysis must then be selected to reflect the characteristics of interest.

Three system-specific tasks must therefore be solved to realize the graph-based analysis of a dynamic system: model a dynamic graph from the system ( $T1$ ), compute graph properties relevant to the system ( $T2$ ), and interpret the results to deduce characteristics of the underlying system ( $T3$ ). An overview of these tasks is given in Table 1.

<b>ID</b>	<b>Description</b>
$T1$	Model the dynamic system as a graph
$T2$	Compute meaningful properties of the dynamic graph
$T3$	Interpret the analysis results and deduce properties of the dynamic system

Table 1: System-specific tasks to solve for the graph-based analysis of a dynamic system

The identification of influential users in online social networks is of great interest for recommendation systems [7], targeted advertisements [31], and the early detection of hate groups [25]. Systems like Twitter are often modeled as follower graphs where users are represented as vertices and their follower relations are modeled by directed edges from follower to followee. The number of incoming edges of a user’s vertex, referred to as in-degree, is then considered as a measure of influence in the network. While this local measure denotes the influence of single users, the overall measure of the in-degree distribution represents the ratio between users of different influence levels.

The analysis of a dynamic system based on measures computed from its graph representation is a promising approach. A general workflow for the process of the *graph-based analysis of dynamic systems* is shown in Figure 3. A dynamic graph is modeled, which reflects the underlying system and all changes occurring to it over time. The dynamic graph’s properties are analyzed at specific points in time, as for the general process for the analysis of dynamic systems. The analysis results are interpreted to reveal the system’s characteristics.

Different approaches exist for the analysis of dynamic graphs. Snapshot-based approaches compute the graph properties of interest based on a single state of the graph [15], just like the separate computation introduced before. Stream-based approaches apply the idea of the event-based maintenance described earlier. They update the previously computed graph properties based on single updates to the graph instead of re-computing everything for each point in time [9].

The main issue that arises in this scenario is the performance of executing the analysis of a dynamic graph. Many algorithms have been developed for snapshot- as well as stream-based approaches to compute various graph measures. Unfortunately, there is no straight-forward way to predict which approach or algorithm performs best for the analysis of arbitrary dynamic graphs. The overall performance of dynamic graph analysis depends on two main factors: the approach and algorithm used for the computation of graph measures and the data structures used to represent the dynamic graph in memory. Other factors are the size, topology, and type of the analyzed graph, the type of changes that occur, and the analysis-frequency.

It poses a challenging problem to decide which approach or algorithm to use for the computation of certain graph properties for a specific graph. This creates the need for a general way to benchmark and compare

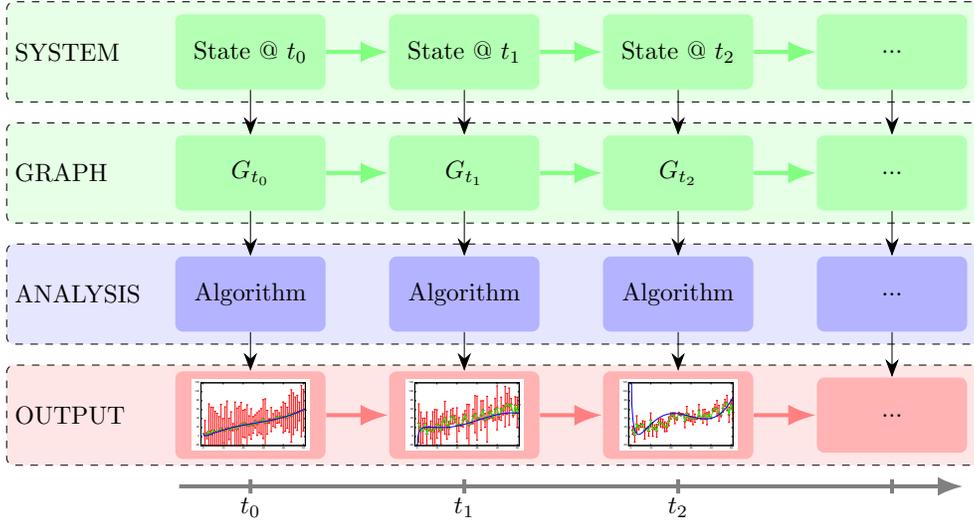


Figure 3: General process of the *Graph-based Analysis of Dynamic Systems*

different approaches and algorithms. Such a benchmarking approach should allow for the comparison of approaches and algorithms for specific classes of dynamic graphs and analysis frequencies.

Many stream-based algorithms have been developed to compute certain graph measures. Still, there exist many graph measures for which no stream-based algorithms have yet been developed. It is therefore promising to investigate the development of new stream-based algorithms to speed up the analysis of dynamic graphs.

The performance of stream-based algorithms scales well when increasing the analysis-frequency. It does not scale well with the growth of the analyzed graph, especially for graph measures that investigate large fractions of the whole graph for each processed update. We should therefore investigate possibilities to distribute the analysis of dynamic graphs among multiple processing units.

The performance of dynamic graph analysis highly depends on the data structures used to represent and maintain the dynamic graph in memory. Which representation performs best depends on the patterns of read operations during analysis and write operations during graph maintenance as well as the number of elements stored. It is not always easy to foresee which data structure performs best for a given scenario. It highly depends on the structure and type of the dynamic graph, the algorithms used for its analysis, and the analysis-frequency. It is therefore necessary to provide a way to benchmark and compare different data structures for the representation of dynamic graphs.

The benchmarking and comparison of data structures provides insights into their corresponding benefits and drawbacks. Even with this knowledge, it cannot easily be decided when to use which data structure. We should therefore investigate the automatic selection of data structures for the analysis of dynamic graphs.

In the next Section, we state five research questions that are based on these general problems for the efficient analysis of dynamic graphs.

## 1 Research Questions

In this Section, we pose five research questions that should be investigated and answered in order to provide the efficient graph-based analysis of dynamic systems as mentioned before. An overview of these research questions is given in Table 2.

**Q1** Different approaches exist to analyze dynamic graphs in general. For each approach and graph measure, a multitude of algorithms has been developed. Which algorithm performs best depends on many factors, like the size and topology of the graph, the number of changes between two analyzed states, and the type of changes that occur. It is not straight-forward to predict the best algorithm for a given scenario. We therefore require a way to benchmark and compare different algorithms in various scenarios to help analysts decide which algorithm they should use to achieve the best performance.

ID	Description
<i>Q1</i>	How can algorithms for dynamic graph analysis be benchmarked and compared?
<i>Q2</i>	How can graph measures be computed efficiently for dynamic graphs?
<i>Q3</i>	How can we speed up the analysis of dynamic graphs using distributed processing?
<i>Q4</i>	How can different graph representations for dynamic graphs be benchmarked and compared?
<i>Q5</i>	How can we determine the most efficient representation of a dynamic graph?

Table 2: Research questions investigated in this thesis

***Q2*** A large number of algorithms has been developed so far to efficiently compute various graph measures using snapshot-based approaches. For many measures, no stream-based algorithms have yet been developed, even though, they promise great performance gains over snapshot-based algorithms for the analysis of dynamic graphs at high frequency. This provides opportunities for the development of faster algorithms for the stream-based analysis of dynamic graphs. We should therefore investigate the efficient computation of graph measures in dynamic graphs.

***Q3*** The scalability of dynamic graph analysis at high frequencies can be achieved using stream-based algorithms. For large graphs and complex graph measures, the analysis using stream-based algorithms does not suffice to provide an efficient analysis of the corresponding graph. We should therefore investigate how to speedup the dynamic graph analysis using distributed processing.

***Q4*** The data structures used to represent a dynamic graph in memory have a big impact on the overall performance of dynamic graph analysis. It is therefore crucial to provide means for their proper benchmarking and comparison in order to judge their respective benefits and drawbacks for specific operations. We therefore require a way to benchmark and compare different data structures for the representation of dynamic graphs for different graph types and sizes as well as arbitrary combinations of algorithms for the computation of graph measures.

***Q5*** There exist many efficient representations of dynamic graphs but it is unclear which one performs best in a given scenario. Their performance is influenced by many factors which make it hard to foresee the best representation beforehand. Therefore, we should investigate how to determine the most efficient graph representation of a dynamic graph to use it during an analysis.

In this thesis, we investigate all these questions. In the next Section, we give an overview of our contributions to answer them.

## 2 Contributions

In this thesis, we investigate the five research questions posed before and present four related contributions:

1. A benchmarking framework for dynamic graph analysis,
2. three novel algorithms that enable the efficient analysis of dynamic graphs,
3. an approach for the parallelization of dynamic graph analysis, and
4. a novel paradigm to select and adapt the data structures for dynamic graph analysis.

In addition, we present three use cases for the graph-based analysis of dynamic systems. They serve as examples for demonstrating the three system-specific tasks *T1*, *T2*, and *T3*. In the remainder of this Section, we give an overview of our four contributions and the three investigated use cases. Then, we detail collaborations and related publications that contribute to the content of this thesis.

**Benchmarking Framework** It is crucial to compare the performance of different algorithms for the analysis of dynamic graphs and their representation using various data structures. We therefore present a benchmarking framework for the analysis of dynamic graphs, called *Dynamic Network Analyzer* (DNA). It allows the generation of dynamic graphs of various types and arbitrary sizes to support the benchmarking and comparison of algorithms and data structures. The framework also supports the development of new algorithms of different types with verifications for their correctness and the automatic determination of the precision of their results. Its extensive visualization components for dynamic graphs and analysis results enables us to monitor them during their analysis. Overall, it provides means to benchmark and compare algorithms (*Q1*) and graph data structures (*Q4*). It supports the development of new algorithms and their subsequent performance analysis. DNA thereby enables us to investigate new approaches for the efficient analysis of dynamic graphs (*Q2*).

**Algorithms** The performance of dynamic graph analysis highly depends on the algorithms used to compute the graph measures of interest for a dynamic graph. We present and evaluate new stream-based algorithms to speed up the analysis of dynamic graphs and thereby provide means to analyze them at high frequencies (*Q2*). We introduce novel stream-based algorithms for the computation of degree distribution, rich-club coefficient, and  $k$ -vertex motif frequencies. We evaluate and compare their performance to existing snapshot-based algorithms using synthetic as well as real-world dynamic graphs. Our results show that these novel algorithms are able to speed up the analysis of dynamic graphs compared to their snapshot-based counterparts. The stream-based algorithm for the computation of  $k$ -vertex motif frequencies, called *StreaM<sub>k</sub>*, achieves great speedups compared to existing approaches and thereby enables the analysis of highly dynamic systems originating from biology.

**Parallelization of Dynamic Graph Analysis** The performance of dynamic graph analysis highly depends on the graph’s size and the complexity of considered graph measures. The analysis does not scale well with an increase in graph size, especially for complex graph measures. We therefore present a novel approach for the distributed processing of dynamic graph analysis, called *parallel Dynamic Graph Analysis* (pDNA). The computational workload is distributed among workers based on a partition of the vertex set. Corresponding subgraphs are assigned to each worker, which computes the respective measure on its local graph view. The changes to the main graph are propagated accordingly to the workers. The results from all workers are then aggregated into the measures for the whole graph for each point in time in a collation step. We evaluate pDNA for the analysis of five graph measures on real-world dynamic graphs. Our results reveal great speedups for the distributed analysis of dynamic graphs using pDNA, especially for complex graph measures.

**Data Structure Selection** The data structures used to represent a dynamic graph in memory have a high impact on the performance of algorithm execution and the maintenance of the dynamic graph over time. The performance for the execution of single operations on data structures is well understood and can be investigated using benchmarks. The sizes of data structures and the frequencies of operations executed on them are not easy to foresee in the context of dynamic graph analysis. This makes it hard to predict which data structures perform best for the analysis of certain graph measures on a specific dynamic graph. We therefore present a novel paradigm for the selection and adaptation of the data structures for dynamic graph analysis. Our approach determines efficient data structures to represent a dynamic graph in memory during analysis (*Q5*). It consists of two parts: a compile-time and a run-time selection of efficient data structures. The compile-time approach is applicable to dynamic graphs whose workload does not change significantly over time. The run-time approach enables the exchange of data structures during run-time and thereby supports scenarios in which the workload changes significantly during the analysis. We evaluate our approaches using synthetic workloads as well as workloads that originate from the analysis of real-world dynamic graphs. In both cases, our approaches achieve great speedups compared to baseline data structures.

**Use Cases** We show the general applicability of the graph-based analysis for dynamic systems that originate from three different fields: social networks, computer networks, and biological networks. We showcase the development of new graph models to grow Web-of-Trust graphs over time. We illustrate the expressiveness of dynamic graph analysis as feature generators for the detection of intrusions in a computer network. We present the analysis of dynamic graphs modeled from molecular dynamics trajectories and show that the graph-based analysis is more expressive than existing analysis approaches from that field.

**Collaborations** While I am the sole author of this thesis, its content is the result of extensive discussions with coworkers, collaborators, and my supervisor Prof. Dr. Thorsten Strufe. The DNA framework has been used and extended by many students that I worked with. Various metrics and algorithms have been implemented by Tim Grube, Benedict Jahn, Bastian Laur, Christoph Schott, Maurice Wendt, Jan Wiese, and Marcel Wunderlich. Generators for dynamic graphs have been implemented by Benedict Jahn, Christoph Schott, Maurice Wendt, Jan Wiese, René Wilmes, and Marcel Wunderlich. The first version of the graph data structures component was implemented by Nico Haase and extended with interfaces to graph databases by Matthias Jordan. Large parts of the aggregation and visualization components have been implemented by René Wilmes. The algorithm  $\text{StreaM}_k$  has been evaluated using trajectories from molecular dynamics simulations provided by Sven Jager. These datasets are also used in the evaluation of pDNA and the selection of efficient data structures. Jeronimo Castrillon and Clemens Deusser helped designing the evaluation of the compile- and run-time approaches for the selection of efficient data structures. Dirk Kohlweyer and Jan Seedorf co-developed the Web-of-Trust model. Matthias Fischer, Jan Reubold, and René Wilmes collaborated in the development and evaluation of gIDS. Kay Hamacher, Michael Vogel, and Sven Jager interpreted the results of the analyzed molecular dynamics simulations. I use these collaborative results in agreement with my collaborators. To indicate that part of the results presented in this thesis are the outcome of collaborations, I use the pronoun ‘we’ instead of ‘I’ throughout this thesis.

**Related Publications** An initial version of the DNA framework was published and presented at SummerSim 2013 [s6]. The algorithm  $\text{StreaM}$ , was published and presented at WABI 2016 [s1].  $\text{StreaM}_k$ , the algorithm was included in an article entitled “Motif-based Analysis of Molecular Dynamics Simulations” and submitted to the SIAM Interdisciplinary Journal of Multiscale Modeling and Simulation. The algorithm  $\text{StreaM}$ , a predecessor to  $\text{StreaM}_k$ , was published and presented at AICoB 2015 [s5]. The compile-time approach for the selection of efficient data structures was published and presented at ComplexNetworks 2015 [s3]. An extension of this approach as well as the run-time approach was published in the Journal of Applied Network Science in September 2016 [s4]. The models for generating instances of the Web-Of-Trust have been published and presented at LCN 2015 [s8]. The model has also been used in a demonstration that was published and presented at ICN 2014 [s9]. The framework GTNA, used for the development and analysis of the model was published and presented at SpringSim 2010 [s2] and an extension at SummerSim 2013 [s7]. The analyses results of molecular dynamics trajectories were published and presented at AICoB 2015 [s5] and are included in the article submitted to the Journal of Multiscale Modeling and Simulation.

## References

- [1] Tero Aittokallio and Benno Schwikowski. Graph-based methods for analysing networks in cell biology. *Briefings in bioinformatics*, 7(3):243–255, 2006.
- [2] Réka Albert, István Albert, and Gary L Nakarado. Structural vulnerability of the north american power grid. *Physical review E*, 69(2):025103, 2004.
- [3] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. *nature*, 406(6794):378–382, 2000.
- [4] Alexandru T Balaban. Applications of graph theory in chemistry. *Journal of Chemical Information and Computer Sciences*, 25(3):334–343, 1985.
- [5] Marc Barthelemy. Betweenness centrality in large complex networks. *The European Physical Journal B-Condensed Matter and Complex Systems*, 38(2):163–168, 2004.
- [6] Wayne J Becketl and John A Schellman. Protein stability curves. *Biopolymers*, 26(11):1859–1877, 1987.
- [7] Betim Berjani and Thorsten Strufe. A recommendation system for spots in location-based online social networks. In *Proceedings of the 4th Workshop on Social Network Systems*, page 4. ACM, 2011.
- [8] Stephen P Borgatti, Ajay Mehra, Daniel J Brass, and Giuseppe Labianca. Network analysis in the social sciences. *science*, 323(5916):892–895, 2009.
- [9] Buriol et al. Counting triangles in data streams. In *ACM SIGMOD-SIGACT*, 2006.

- [10] Paolo Crucitti, Vito Latora, and Massimo Marchiori. A topological analysis of the italian electric power grid. *Physica A: Statistical mechanics and its applications*, 338(1):92–97, 2004.
- [11] SN Dorogovtsev, JFF Mendes, and AN Samukhin. Size-dependent degree distribution of a scale-free growing network. *Physical Review E*, 63(6):062101, 2001.
- [12] Shimon Even. *Graph algorithms*. Cambridge University Press, 2011.
- [13] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. In *ACM SIGCOMM computer communication review*, volume 29, pages 251–262. ACM, 1999.
- [14] Alan R Fersht, Andreas Matouschek, and Luis Serrano. The folding of an enzyme: I. theory of protein engineering analysis of stability and pathway of protein folding. *Journal of molecular biology*, 224(3):771–782, 1992.
- [15] Andrew V Goldberg. *Efficient graph algorithms for sequential and parallel computers*. PhD thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 1987.
- [16] Roger Guimera, Stefano Mossa, Adrian Turttschi, and LA Nunes Amaral. The worldwide air transportation network: Anomalous centrality, community structure, and cities’ global roles. *Proceedings of the National Academy of Sciences*, 102(22):7794–7799, 2005.
- [17] Petter Holme, Beom Jun Kim, Chang No Yoon, and Seung Kee Han. Attack vulnerability of complex networks. *Physical Review E*, 65(5):056109, 2002.
- [18] Matthew O Jackson et al. *Social and economic networks*, volume 3. Princeton university press Princeton, 2008.
- [19] Sushil Jajodia, Steven Noel, and Brian OBerry. Topological analysis of network attack vulnerability. In *Managing Cyber Threats*, pages 247–266. Springer, 2005.
- [20] Dusanka Janezic, Ante Milicevic, Sonja Nikolic, and Nenad Trinajstic. *Graph-theoretical matrices in chemistry*. CRC Press, 2015.
- [21] Oliver Mason and Mark Verwoerd. Graph theory and networks in biology. *IET systems biology*, 1(2):89–119, 2007.
- [22] David Myr. Real time vehicle guidance and forecasting system under traffic jam conditions, November 12 2002. US Patent 6,480,783.
- [23] Mark EJ Newman, Duncan J Watts, and Steven H Strogatz. Random graph models of social networks. *Proceedings of the National Academy of Sciences*, 99(suppl 1):2566–2572, 2002.
- [24] Frank H Niesen, Helena Berglund, and Masoud Vedadi. The use of differential scanning fluorimetry to detect ligand interactions that promote protein stability. *Nature protocols*, 2(9):2212–2221, 2007.
- [25] Andre Oboler. The rise and fall of a facebook hate group. *First Monday*, 13(11), 2008.
- [26] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: bringing order to the web. 1999.
- [27] John Scott. *Social network analysis*. Sage, 2012.
- [28] Claire Tomlin, George J Pappas, and Shankar Sastry. Conflict resolution for air traffic management: A study in multiagent hybrid systems. *IEEE Transactions on automatic control*, 43(4):509–521, 1998.
- [29] Stanley Wasserman and Katherine Faust. *Social network analysis: Methods and applications*, volume 8. Cambridge university press, 1994.
- [30] Duncan J Watts and Steven H Strogatz. Collective dynamics of small-worldnetworks. *nature*, 393(6684):440–442, 1998.
- [31] Wan-Shiou Yang, Jia-Ben Dia, Hung-Chi Cheng, and Hsing-Tzu Lin. Mining social networks for targeted advertising. In *Proceedings of the 39th Annual Hawaii International Conference on System Sciences (HICSS’06)*, volume 6, pages 137a–137a. IEEE, 2006.

- [32] Ellen W Zegura, Kenneth L Calvert, and Michael J Donahoo. A quantitative comparison of graph-based models for internet topology. *IEEE/ACM Transactions on Networking (TON)*, 5(6):770–783, 1997.
- [33] Shi Zhou and Raúl J Mondragón. The rich-club phenomenon in the internet topology. *IEEE Communications Letters*, 8(3):180–182, 2004.

## Own Publications

- [s1] Sven Jager, Benjamin Schiller, Thorsten Strufe, and Kay Hamacher. Stream-t\_g: Algorithms for analyzing coarse grained rna dynamics based on markov models of connectivity-graphs. In *International Workshop on Algorithms in Bioinformatics*, pages 197–209. Springer, 2016.
- [s2] Benjamin Schiller, Dirk Bradler, Immanuel Schweizer, Max Mühlhäuser, and Thorsten Strufe. Gtna: a framework for the graph-theoretic network analysis. In *Proceedings of the 2010 Spring Simulation Multiconference*, page 111. Society for Computer Simulation International, 2010.
- [s3] Benjamin Schiller, Jeronimo Castrillon, and Thorsten Strufe. Efficient data structures for dynamic graph analysis. In *2015 11th International Conference on Signal-Image Technology & Internet-Based Systems (SITIS)*, pages 497–504. IEEE, 2015.
- [s4] Benjamin Schiller, Clemens Deusser, Jeronimo Castrillon, and Thorsten Strufe. Compile- and run-time approaches for the selection of efficient data structures for dynamic graph analysis. *Journal of Applied Network Science*, 1(9), 2016.
- [s5] Benjamin Schiller, Sven Jager, Kay Hamacher, and Thorsten Strufe. Stream-a stream-based algorithm for counting motifs in dynamic graphs. In *International Conference on Algorithms for Computational Biology*, pages 53–67. Springer, 2015.
- [s6] Benjamin Schiller and Thorsten Strufe. Dynamic network analyzer building a framework for the graph-theoretic analysis of dynamic networks. In *Proceedings of the 2013 Summer Computer Simulation Conference*, page 49. Society for Modeling & Simulation International, 2013.
- [s7] Benjamin Schiller and Thorsten Strufe. Gtna 2.0-a framework for rapid prototyping and evaluation of routing algorithms. In *Proceedings of the 2013 Summer Computer Simulation Conference*, page 23. Society for Modeling & Simulation International, 2013.
- [s8] Benjamin Schiller, Thorsten Strufe, Dirk Kohlweyer, and Jan Seedorf. Growing a web of trust. In *Local Computer Networks (LCN), 2015 IEEE 40th Conference on*, pages 100–108. IEEE, 2015.
- [s9] Jan Seedorf, Bilal Gill, Dirk Kutscher, Benjamin Schiller, and Dirk Kohlweyer. Demo overview: fully decentralised authentication scheme for icn in disaster scenarios. In *Proceedings of the 1st international conference on Information-centric networking*, pages 191–192. ACM, 2014.