

Algorithms for Large-Scale Network Analysis and the NetworKit Toolkit

Eugenio Angriman^{1(⊠)}, Alexander van der Grinten¹, Michael Hamann², Henning Meyerhenke¹, and Manuel Penschuck³

> ¹ Humboldt-Universität zu Berlin, Berlin, Germany {angrimae,avdgrinten,meyerhenke}@hu-berlin.de
> ² Karlsruhe Institute of Technology, Karlsruhe, Germany michael.hamann@kit.edu
> ³ Goethe University Frankfurt, Frankfurt am Main, Germany mpenschuck@ae.cs.uni-frankfurt.de

Abstract. The abundance of massive network data in a plethora of applications makes scalable analysis algorithms and software tools necessary to generate knowledge from such data in reasonable time. Addressing scalability as well as other requirements such as good usability and a rich feature set, the opensource software NETWORKIT has established itself as a popular tool for largescale network analysis. This chapter provides a brief overview of the contributions to NETWORKIT made by the SPP 1736. Algorithmic contributions in the areas of centrality computations, community detection, and sparsification are in the focus, but we also mention several other aspects – such as current software engineering principles of the project and ways to visualize network data within a NETWORKIT-based workflow.

Keywords: Network analysis · Algorithms · Software package

1 Introduction

Network phenomena surround us, be they social contact networks, organizational structures, or infrastructure networks such as the energy grid, roads or the (physical) internet. Purely virtual networks such as the world wide web, online social networks, or co-authorship networks can become particularly large and play an ever increasing role in our daily lives [8,62]. Traditional data analysis has been and is very successful in discovering knowledge from non-network (e.g., geometric or relational) data [50]. Yet, networks and their analysis are about "dependence, both between and within variables" [26]. To uncover implicit dependencies hidden in the data, it thus requires appropriate algorithmic techniques (some of which are also covered in Leskovec et al.'s textbook on mining massive datasets [50]).

Massive networks, often with billions of vertices and edges, pose challenges to many established analysis concepts and algorithms due to their prohibitive computational costs. This leads to the ongoing development of efficient and scalable algorithms. The open-source software package NETWORKIT¹ [75 SPP] aims to combine a broad

¹ https://networkit.github.io/.

range of such algorithms for the analysis of large networks and to make them accessible via consistent, easy to use, and well-documented frontends. For instance, it offers a feature-rich Python API which integrates into the large Python ecosystem for data analysis. Under the hood, the heavy lifting is carried out by performance-oriented algorithms that are implemented in C++ and often use multicore parallelism. The package is also well suited to develop and evaluate novel algorithmic approaches. As such, NET-WORKIT received numerous unique scalable algorithms and implementations in recent years, particularly designed to handle large inputs.

In this chapter, we present a high-level overview of NETWORKIT (Sect. 2) and portray algorithmic research results derived with and for NETWORKIT – mostly those obtained by projects of SPP 1736. We cover four main topics: centrality algorithms (Sect. 3), community detection (Sect. 4), graph sparsification (Sect. 5) as well as graph drawing and network visualization (Sect. 6). While these have been focus areas of NET-WORKIT development during the lifetime of SPP 1736, the package has been used in various other application contexts such as quantum chemistry [56 SPP] and digital humanities [47].

2 NETWORKIT—An Overview

NETWORKIT is in development since 2013. The architecture of the current codebase was released in 2014. At the time of writing, NETWORKIT has a regular release cycle with two new major releases per year. Staudt et al. [75 SPP] describe the package's state at the end of 2015. In this section, we consequently focus on the many additions of new functionality as well as improvements to the code quality that have been realized in the meantime. This concerns new performance-oriented graph algorithms, engineering to speed up existing algorithms, more software engineering guidelines and best practices, as well as the modernization and extension of NETWORKIT's integration with other tools within a rich ecosystem (as detailed in Sect. 2.2).

2.1 Design Considerations

NETWORKIT consists of several Python modules wrapping an independently usable core library that is written in C++. Both parts are connected using Cython and are tightly integrated to offer consistent interfaces for most features. The package is organized into multiple modules, each focusing on one (class of) network analytic problem(s). Important modules deal with network centrality (centrality), community detection (community and scd) as well as graph generation and perturbation (generators and randomization). Some novel algorithms in the centrality, community, and sparsification modules that were developed within SPP 1736 are described in more detail in Sects. 3 to 5. Other important modules that are not covered here include modules for graph algorithms in the language of linear algebra (algebraic, following the philosophy of GraphBLAS [45 SPP]), decomposition of graphs into components (components), distance computations (distance), reading and writing graphs (io), link prediction (linkprediction), graph coarsening (coarsening), and more.

As a graph data structure, NETWORKIT uses an adjacency array using dynamic arrays (std::vector) to store vertices and their neighborhoods. It also supports edge weights and edge IDs. This data structure was chosen over static ones such as CSR matrices since it allows for efficient dynamic updates. The design is complemented by several non-trivial algorithms that can efficiently update their results if the underlying graph changes (i.e., after adding and/or deleting edges).

Many of NETWORKIT's algorithms use OPENMP for shared-memory parallelism. In fact, several algorithms in NETWORKIT exhibit best-in-class parallel performance [36 SPP]. Based on an empirical comparison [46 SPP] between NETWORKIT and several distributed frameworks for data and network analysis, NETWORKIT's speed advantage usually remains true in comparison to distributed systems with eight-fold resource consumption. Ref. [46 SPP] finds that a shared-memory machine is sufficient to solve many network analytic problems on real-world instances and concludes that shared-memory parallelism should be preferred to distributed graph algorithms as long as the input graph fits into main memory.

2.2 Ecosystem

In recent years, NETWORKIT matured into an actively maintained open-source project with more than 140 000 lines of code and a steadily growing number of users and contributors. By now, the software package exceeds a critical size that warrants efforts beyond the development of new algorithmic features.

To ease contributions and uphold the code quality, NETWORKIT offers detailed guidelines and implements a thorough review process. We also make heavy use of unittests, static code analysis and automated code-formatting as part of our continuous integration pipeline, which targets the three major operating systems. As many new tests improve the coding standards, we continuously modernize the codebase. Still, backwards compatibility is a major concern and manifests itself, for instance, in long-term compiler support and in as few changes breaking the API as possible (preceded by a deprecation period of at least one major version release).

Users benefit from a welcoming community, ever-improving documentation, interactive examples showcasing most features, a regular release schedule, and growing support for package managers (currently brew, Conda, pip, and Spack). NETWORKIT naturally interacts with external projects such as GEPHI (see Sect. 6), SIMEXPAL [4 SPP], and NETWORKX as well as graph repositories and formats including KONECT, SNAP, and METIS; recent changes make it now even possible to develop standalone NET-WORKIT Python modules.

Graph data can not only be imported but also be synthesized. To this end, NETWORKIT offers versatile graph generators in the modules generators and randomization. Among others, they are designed to generate and supplement datasets for applications ranging from rapid prototyping to experimental campaigns. Here, we only mention the supported network *models* since Chap. 2 surveys novel generation *algorithms* obtained during SPP 1736. We include here citations to models or generators developed for/with NETWORKIT.

 Focus on community structure: Clustered-Random-Graph, LFR, PubWeb, R-MAT, Stochastic Block Model, Watts-Strogatz

- Prescribed degrees: Havel-Hakimi, Chung-Lu, Curveball and Global-Curveball [28 SPP], Edge-Switching
- Preferential attachment processes: Barabási-Albert, Dorogovtsev-Mendes
- Geometrically embedded: Hyperbolic Random Graph [52 SPP, 53 SPP, 54 SPP, 55 SPP, 19 SPP], Geometric Inhomogenous Random Graph [19 SPP], Mocnik [59, 60]
- Basic models: G(n, p), Lattice.

Several generators have dynamic variants simulating the evolution of graphs over time.

3 Centrality Algorithms

One of the most popular concepts used for the analysis of a graph G = (V, E) is *centrality*. Centrality measures assign a score to each vertex² (or group of vertices) based on its structural position or importance; these scores allow a corresponding vertex ranking [21]. As an example, the well-known PageRank [27] is a centrality measure originally devised for web page (and eventually search query) ranking. It is important to match the underlying research question with the appropriate centrality measure [77 SPP] and no single measure is universal. Thus, dozens of measures have been proposed in the literature [21].

As described in more detail below, the centrality research within NETWORKIT revolves not only around faster algorithms for computing individual scores and top-*k* rankings. Another emphasis is placed on two families of centrality-driven optimization problems (centrality improvement and group centrality) and how to scale approximation algorithms or heuristics for their solution to much larger input sizes. For a broader overview, also with a scalability focus, the reader is referred to Ref. [35 SPP].

It should also be noted that fast centrality algorithms can be useful in different (but related) contexts as well; e.g., scores of several centrality measures are used as shortcuts for more expensive influence maximization calculations [70 SPP]. Also, using score distributions for graph fingerprinting (putting graphs into classes where all members have similar distributions) is a conceivable use case with the need for numerous measures that can be computed quickly.

3.1 Individual Centrality Scores

We first discuss centrality measures for individual vertices, i.e., measures that assign a centrality score to each $v \in V$. During SPP 1736, our focus has been on two classes of centrality measures: centralities that make use of shortest path computations (i.e., (harmonic) closeness and betweenness) and algebraic centrality measures that consider more than just shortest paths (like Katz centrality and electrical closeness). Figure 1 depicts the distribution of these centralities for a single network, including the ED Walk centrality that we propose in Ref. [3 SPP].

² Edge centrality measures are ignored here in the interest of space.



Fig. 1. Histograms of the distribution of vertex centrality measures of the JAZZ network, which models the collaboration of Jazz musicians [34].

Betweenness. Betweenness centrality is based on the fraction of shortest paths a vertex participates in. NETWORKIT implements the well-known Brandes algorithm [23] for exact betweenness and several algorithms for betweenness approximation. For static graphs, it has an implementation of the KADABRA algorithm [22]; additionally, NET-WORKIT can approximate betweenness in dynamic graphs [15 SPP]. Both of these algorithms employ a sampling technique that was originally introduced by Riondato and Kornaropoulos [66]. More precisely, the algorithms sample pairs (*s*,*t*) of source and target vertices uniformly at random. For each (*s*,*t*), a single shortest path is sampled uniformly at random out of all shortest *s*-*t* paths. The algorithms count the number of occurrences of vertices on these paths; they differ in their stopping conditions. The multi-threaded implementation of the static KADABRA algorithm additionally exploits a fast data structure for asynchronous synchronization barriers [36 SPP]. To the best of our knowledge, NETWORKIT's implementation of KADABRA is the fastest betweenness approximation code that is available for multi-threaded machines.

In Ref. [39 SPP], this algorithm was extended to work with replicated graphs in distributed memory. The resulting algorithm obtains good parallel speedups and performs well even on multi-socket shared memory machines due to the fact that it can avoid NUMA bottlenecks. Since distributed memory algorithms are outside the scope of NETWORKIT, this implementation is available externally.

Closeness. Closeness centrality also uses the notion of shortest paths: it quantifies the importance of a vertex $v \in V$ depending on how close v is to all the other vertices of the graph [11]. It is defined as $c(v) := (n-1)/(\sum_{w \in V} d(v,w))$ and computing it for a single vertex requires to run a single-source shortest path (SSSP) algorithm. The textbook algorithm to identify the top-k vertices with highest closeness centrality computes c(v) for each vertex of the graph by running n SSSPs, which is impractical for large-scale networks. NETWORKIT improves on this by providing an algorithm which finds the top-k vertices with highest closeness centrality along with their exact value of $c(\cdot)$ [12 SPP]. Even though the worst-case running time of the algorithm is also

 $\Omega(|V||E|)$, experimental evaluation on real-world data shows that, for small values of k, the algorithm is in practice much more efficient than the textbook algorithm and other state-of-the-art strategies.

NETWORKIT additionally implements a batch-dynamic version of this algorithm [18 SPP,2 SPP], which also addresses harmonic centrality [21,67] – an alternative definition of closeness centrality introducing support for disconnected graphs. Experiments on both real-world and synthetic instances demonstrate that, for moderately large batches of edge updates, the dynamic algorithm is up to four orders of magnitude faster than a static recomputation from scratch.

Electrical Closeness. Electrical resistance is a distance function on graphs that is constructed by interpreting the graph as a network of electrical resistors and by measuring the effective resistance between vertices in this network. If the usual distance function (based on shortest-path distances) in the definition of closeness is replaced by effective resistance, one obtains the definition of *electrical* closeness. This centrality measure has been gaining attention due to the fact that it considers paths of any length. NETWORKIT has an efficient approximation algorithm to compute electrical closeness [6 SPP]. This algorithm exploits a well-known connection between electrical networks and uniform spanning trees to approximate electrical closeness faster than previous numerical algorithms (including the numerical algorithm from Ref. [17 SPP]) and can handle graphs with hundreds of millions of edges.

As part of our work on electrical closeness, NETWORKIT gained support for various numerical algorithms. These are typically either used as subprocedures of our algorithms or for performance and/or quality comparisons; however, they can also be called as standalone numerical solvers. Experiments with an (in terms of theoretical analysis) fast Laplacian solver revealed severe limitations in practice [43 SPP] – which is why it was discarded. Instead, we included a fast implementation [17 SPP] of the lean algebraic multigrid algorithm (LAMG) [51], which is particularly well-suited to solve series of Laplacian linear systems with identical system matrices.

Katz Centrality. NETWORKIT also implements an approximation algorithm for Katz centrality that can handle graphs with billions of edges within a few minutes [38 SPP]. The algorithm utilizes lower and upper bounds on the centrality score of each vertex and improves these bounds until the Katz centrality ranking is computed with sufficient precision. In comparison to earlier combinatorial algorithms for Katz centrality, our algorithm is the first to obtain a provable approximation bound and/or the correctness of the ranking. It is also at least 50% faster than numerical methods. NETWORKIT provides a parallel implementation of this algorithm that can also handle dynamic graphs. In Ref. [38 SPP], we additionally provide a GPU-based implementation which is not part of NETWORKIT.

3.2 Improving One's Own Centrality

One possible way to improve one's ranking position in a web search is to attract links from influential web pages. For some time, this led to so-called link farming [49] for

search engine optimization. More generally, beyond web search, one wants to increase the centrality of a vertex by adding a specified number of new edges incident to it. Crescenzi et al. [30] addressed this problem for closeness centrality. As a follow-up to that work, Ref. [13 SPP] considered two betweenness centrality improvement problems: maximizing the betweenness *score* of a given vertex (MBI) and maximizing the *ranking position* of a given vertex (MRI). The paper proves that both problems are hard to approximate. Unless $\mathcal{P} = \mathcal{NP}$, MBI cannot be approximated within a factor greater than $1 - \frac{1}{2e}$ and for MRI there is no α -approximation algorithm for any constant $\alpha \leq 1$. The paper also proposes a simple greedy algorithm for MBI that performs well in practice and provides a (1 - 1/e)-approximation. This way, MBI can be approximated for (most) networks with up to 10^5 edges in a matter of seconds or a few minutes. The greedy algorithm's implementation builds, among others, upon a dynamic algorithm for betweenness centrality [16 SPP] that can update the betweenness scores of all vertices much faster after small graph changes (such as the insertion of one or few edges).

3.3 Group Centrality Optimization

Group centralities are network-analytic measures that quantify the importance of vertex groups [31]. In contrast to centrality measures that apply to individual vertices, the goal of these measures is to determine how well the entire group jointly "covers" the graph; i.e., the group centrality score is *not* determined by the scores of individual vertices.

NETWORKIT includes various group centrality algorithms to approximate sets of vertices that maximize the group centrality score. Most of the algorithms are based on submodular optimization. For example, NETWORKIT implements a greedy algorithm to approximate group degree and the group betweenness maximization algorithm by Mahmoody et al. [57]. New algorithms developed as part of SPP 1736 are the GED-Walk approximation algorithms from Ref. [3 SPP] and various group closeness algorithms; these algorithms are described below. A very recent addition to NETWORKIT is an approximation algorithm for group forest closeness centrality; for details we refer to Ref. [37 SPP].

Group Closeness. Group closeness measures the importance of a *group* of vertices $S \subset V$ as the reciprocal of the sum of the distances from *S* to the vertices in $V \setminus S$, where the distance from *S* to a vertex $v \in V$ is defined by the minimum $d(S, v) := \min_{u \in S} d(u, v)$. Finding the group S^* with highest group closeness is known to be an \mathcal{NP} -hard optimization problem [29, 1 SPP]. Thus, in practice, the problem is addressed on large-scale networks either with heuristics or approximation algorithms. NETWORKIT provides a greedy heuristic [14 SPP] that computes a set of vertices with high group centrality. On small enough instances where it is feasible to compute the optimum, it has been shown that the algorithm yields solutions with nearly optimal quality.

An alternative heuristic, which allows to trade quality for speed, is based on local search. NETWORKIT implements a family of local search heuristics for group closeness maximization that achieve different trade-offs between quality and running time [5 SPP]. In general, they are one to three orders of magnitude faster than the greedy algorithm. At the same time, our algorithms retain 80%—and, in numerous cases, even

more than 99%—of the greedy algorithm's solution quality. NETWORKIT also includes the first approximation algorithm for group closeness maximization [1 SPP] (for undirected graphs) which yields solutions with higher quality than the greedy algorithm at the cost of additional running time.

A major limitation of group closeness is that it can only handle (strongly) connected graphs – the distance between unreachable vertices is either undefined or infinite, and an infinite denominator results in group closeness score of zero. Another group centrality measure that also handles disconnected graphs is group harmonic centrality, which is defined as $GH(S) := \sum_{u \in V \setminus S} d(S, u)^{-1}$. Maximizing *GH* has been shown to be an \mathcal{NP} -hard problem [1 SPP] as well and two approximation algorithms for group harmonic maximization have been introduced in Ref. [1 SPP]; both of them are available in NETWORKIT.

GED-Walk. GED-Walk (GED = group exponentially decaying) is an algebraic group centrality measure that was introduced in Ref. [3 SPP]. Similarly to Katz centrality (which only applies to individual vertices), GED-Walk counts the number of *walks* (and not paths) in the graph. Unlike Katz centrality, it counts walks that *cross* the group of vertices (instead of counting walks that *start* (or end) at certain vertices). Computing GED scores can essentially be done via sparse matrix-vector multiplication; hence, the measure can be computed faster than centrality measures that involve the computation of shortest paths. In Ref. [3 SPP], we propose a greedy algorithm that computes a group with approximately maximal GED-Walk centrality. The algorithmic approach is based on techniques derived from our Katz algorithm [38 SPP] and iteratively refines bounds on the group centrality score. In experiments, GED-Walk maximization turns out to be at least one order of magnitude faster than the corresponding greedy algorithms for group betweenness and group closeness. When applied within semi-supervised vertex classification, GED-Walk improves the accuracy compared to various existing measures.

4 Community Detection

Community detection aims to detect subgraphs that are internally densely and externally sparsely connected. From this fuzzy idea, many formalizations and algorithms have been developed [32]. A division of the graph into disjoint communities is the most frequently studied setting. The most popular quality measure for this setting is modularity [63]. As it is \mathcal{NP} -hard to find the (clustering with) optimal modularity score [24], heuristics are used in practice. A very popular one is the Louvain algorithm [20]. While it is already quite fast, it is purely sequential in its original formulation and thus does not exploit the many cores available in modern processors. Already the earliest work in NETWORKIT includes the development of a parallel variant of the Louvain algorithm named PLM [72]. This first work also includes a fast parallel label propagation algorithm named PLP and an ensemble algorithm that combines several runs of PLP with a final step where PLM is used. Later improvements to PLM, including the parallelization of additional steps, made PLM so fast that it outperformed the ensemble approach both in terms of speed and quality [74 SPP]. Further, a refinement round similar to Ref. [68]

11

has been introduced that further increases the quality at the expense of a slightly longer running time. PLM was later used in a case study on correspondences between clusterings [33 SPP]. With such correspondences one can reveal how one clustering differs from another one, e.g., when computed with different algorithms or after minor graph changes.

If only a community around a specific vertex or a set of vertices (so-called seed vertices) is desired, we do not need to detect communities that cover the whole graph. Many such algorithms greedily add new vertices until a local minimum of a certain quality function is reached. A first study on such local community detection algorithms [71 SPP] based on NETWORKIT has shown that they are quite slow and imprecise in comparison to PLM. A more recent study [41 SPP] shows that many local community algorithms detect a community in which the seed is not strongly connected. Only algorithms that employ further guidance, e.g., using edge scores based on triangles, are able to correctly identify a community detection algorithms can be improved by starting with the largest clique in the subgraph induced by the neighbors of the seed vertex. For this, the possibility to combine two local community detection algorithms has been added to NETWORKIT – a first one that detects the clique and then a second one that expands this clique into a community [41 SPP]. This allows changing both the seeding strategy and the latter expansion step.

For the experimental evaluation of community detection algorithms, suitable input instances are required [7]. Ideally, instances from applications of community detection with known ground truth communities should be used for this. However, they are frequently either quite small, unavailable due to privacy concerns or commercial interests, or the available ground truth data cannot be recovered from the graph's structure [32,65]. For this reason, synthetically generated benchmark graphs with generated ground truth communities are frequently used. The most popular one is the LFR benchmark graph generator [48], of which NETWORKIT also provides an implementation for the case of unweighted, undirected graphs with disjoint communities [73 SPP] (see also Chap. 2). Due to a partial parallelization and more efficient data structures, experiments show a speedup compared to the original implementation of 18 to 70 using 16 cores [73 SPP]. When the similarity between a detected and a (possible) ground truth community is low, it is often not clear if such a similarity could also be achieved by chance. Therefore, Hamann et al. [41 SPP] also introduced a simple baseline algorithm using a BFS that stops when the same number of vertices as contained in the ground truth community have been visited and returns them as community. Together with additional methods for the evaluation of the found communities, NETWORKIT thus provides a comprehensive framework for the development, evaluation, and application of local community detection algorithms.

Nastos and Gao [61] suggest quasi-threshold graphs, i.e., graphs that do not contain a path or cycle of four vertices as vertex-induced subgraph, as a model for communities in social networks. As a given graph is usually not a quasi-threshold graph, they suggest to insert and delete as few edges as possible to transform a graph into a quasi-threshold graph. The connected components are then considered as communities. The first scalable heuristic for this problem [25 SPP] has been implemented in NETWORKIT, for details we refer to Chap. 7.

5 Graph Sparsification

Centrality measures suggest that certain vertices or edges are more important than others. In graph *sparsification*, the idea is to exploit this fact to obtain a subset of the vertices and/or edges that preserve key properties of the graph, i.e., to select vertices and edges that are important for these properties. Properties of the graph can be preserved either directly or in a scaled version. For example, the degree distribution cannot be exactly preserved when we remove edges, but we can preserve the general shape of the degree distribution. Graph sparsification can provide insights into the structure of a graph, as it provides insights on how much redundancy there is and which edges are important for certain properties. An application of these insights is speeding up other network analysis tasks or making them possible in the first place by reducing the graph's size such that the running time and memory requirements are reduced [69]. Further, some of these sparsification techniques can also remove noise from the graph such that, e.g., more informative drawings can be generated [64 SPP]. In NETWORKIT, we provide a set of edge sparsification algorithms [40 SPP]. Given a graph G = (V, E), they identify subsets $E' \subset E$ of the edges such that G' = (V, E') preserves certain properties of G. We currently do not consider vertex sparsification, i.e., filtering vertices while maintaining properties of the graph - since in many network analysis tasks (like vertex centralities or community detection), we are interested in a result for every vertex. If some vertices were no longer part of the graph, we would need to extrapolate their results, requiring an additional post-processing step for every network analysis task.

With its diverse set of network analysis algorithms, NETWORKIT provides the ideal testbed for sparsification algorithms. A study [40 SPP] compares a set of six existing and one novel sparsification algorithm as well as five novel variants of the existing algorithms using NETWORKIT. The study shows that these sparsification algorithms can be classified into three groups: those that primarily preserve edges within densely connected areas, those that primarily preserve connectivity between different areas, and those that are almost or completely random. The algorithms in the first group strengthen the formation of communities and either keep or increase the average local clustering coefficient as already suggested by previous work [69,64 SPP]. The novel local degree technique, on the other hand, keeps distances in the graph and thus the diameter small, see Fig. 2 for an example. As the results show, it is also good at preserving vertex centralities. Completely random filtering also works surprisingly well at preserving a wide range of network properties. The study shows that all methods perform better for most measures if, instead of directly filtering edges globally, a vertex of degree d keeps its top d^e neighbors for some exponent e < 1. This local filtering step has been proposed before [69] for a single sparsification algorithm and the study suggests to apply it to all considered algorithms. In particular, this preserves connectivity of the graph quite well and in general leads to a more even distribution of the preserved edges.

All of these sparsification algorithms can be decomposed into two steps: A first step that assigns each edge a score and a second step that only keeps a certain fraction of the highest-rated edges. Even the local filtering step can be implemented as a



Fig. 2. Drawing using GEPHI [9] of the JAZZ network [34] (left) and a sparsified version containing 15% of the edges (right) using the novel local degree algorithm. Vertex size and color is proportional to degree. (Color figure online)

transformation of edge scores. This makes it possible to easily combine existing and new algorithms. Further, the resulting scores can be considered as edge centrality measures that permit a ranking of the edges. With the help of visualization software like GEPHI [9] (Sect. 6), the scores can also be visualized or used for interactive filtering of edges.

6 Graph Drawing and Network Data Visualization

In exploratory network analysis, one needs to evaluate several properties of the network, which requires writing code to run algorithms and plot their results. To speed up this process, NETWORKIT provides a dedicated profiling module that allows non-expert users to run several network analysis algorithms as a single program and visualize their results in a graphical report that can be rendered in a Jupyter Notebook or exported as an HTML or a LATEX document. As thoroughly explained in Ref. [75 SPP], first the report lists global properties of the networks such as the size and the density. Then it provides an overview of the distribution of several centrality networks as histograms (as shown in Fig. 1, Sect. 3), followed by a more detailed statistical analysis. Finally, the report includes a matrix with the Spearman correlation coefficients between the rankings of the vertices according to the considered centrality measures; an example for the JAZZ network is shown in Fig. 3.

When dealing with large graphs, statistical overviews as the ones mentioned are indispensable, since the well-known vertex-edge diagrams do not even scale to graphs of medium size (without further adjustments). For small graphs, however, visualizations such as those diagrams can be very valuable. In general, the goal of graph visualization [10] is to represent graphs in a form that is meaningful to the human eye. Popular

+1.000	ED Walk					
+0.955	+1.000	PageRank				
+1.000	+0.956	+1.000	Katz Centrality			
+0.645	+0.736	+0.645	+1.000 Betweenness			
+0.967	+0.927	+0.967	+0.677	+1.000	Harmonic Closeness	
+0.982	+0.952	+0.982	+0.659	+0.954	+1.000	Electrical Closeness

Fig. 3. Spearman's correlation coefficients between vertex rankings obtained with different centrality measures for the JAZZ network. Darker [lighter] block shades indicate higher [smaller] correlation values.



Fig. 4. Visualization example with GEPHI of the KARATE graph. Red vertices have the highest harmonic centrality, blue vertices the lowest. (Color figure online)

application areas for graph visualization are biology (e.g., genetic maps), chemistry (e.g., protein functions) [42], social network analysis [47], and many more. GEPHI [9] is a popular Java-based GUI application to explore and visualize graphs. NETWORKIT's gephi module [40 SPP] allows to use GEPHI to visualize graphs along with additional vertex- or edge attributes with minimal effort. Figure 4 shows the visualization in GEPHI of the popular KARATE graph obtained by the ForceAtlas2 graph drawing algorithm [44] and by coloring the vertices according to their harmonic centrality score.

Graph drawing actually precedes visualization in most cases. It is the process of computing meaningful coordinates for the graph vertices where such information is not supplied with the graph. NETWORKIT's approach for the most part is to use the graph drawing capability in GEPHI. It has, however, also an implementation of an algorithm for the maxent-stress objective function, following Ref. [58 SPP]. Here, the main intention is to solve an optimization problem that computes the three-dimensional structure of biomolecules, given distance information between some atom pairs. To this end, the original algorithm received several application-specific adaptations [76 SPP], e.g., to be able to handle noisy data appropriately. As a result, the new algorithm by far outperforms its competitors in terms of speed and flexibility, and often even produces a superior solution quality.

7 Conclusions

The main design goals of NETWORKIT (speed, rich feature set, usability, and integration into an ecosystem) prove to be very useful for users, but they can also be challenging for the developers. One lesson learned to keep an academic open-source project of this size manageable and alive, is to combine best practices in both software engineering and algorithm engineering [4 SPP]. For example, a proper modularization allows easier reuse and combination of components, leading to a better extensibility and maintainability. These keywords are well-known in software engineering, but they also have their effect in algorithm design and implementation – in particular a simplified exploration of the design space in experimental algorithmics. NETWORKIT has already proved to be very useful in this respect for developers.

We have seen that approximation and parallelism can bring us a long way regarding scalability. They are the obvious, but certainly not the only choices for acceleration: exploiting the structure of the data, e.g., small vs. large diameter [12 SPP], can yield significant speedups on real-world data—even in the context of exact computations and potentially on top of parallelism.

NETWORKIT is constantly improved and extended – according to the resources available to the project. There are numerous ideas for larger updates from various angles – of which we mention only two representative ones: inherent support for attributes within (some of) the algorithms and further/improved integration with other tools. The latter is particularly geared towards a closer connection with machine learning, both on an algorithmic and a software tool level. Given the current interest in machine learning for data analysis, complete workflows within one seamless toolchain including NET-WORKIT and tools such as SCIKIT-LEARN can be expected to be very attractive for users from many domains.

References

- 1 SPP. Angriman, E., Becker, R., D'Angelo, G., Gilbert, H., van der Grinten, A., Meyerhenke, H.: Group-harmonic and group-closeness maximization - approximation and engineering. In: ALENEX. SIAM (2021)
- 2 SPP. Angriman, E., Bisenius, P., Bergamini, E., Meyerhenke, H.: Computing top-*k* closeness centrality in fully-dynamic graphs. Taylor & Francis (2021). Currently in review
- 3 SPP. Angriman, E., van der Grinten, A., Bojchevski, A., Zügner, D., Günnemann, S., Meyerhenke, H.: Group centrality maximization for large-scale graphs. In: ALENEX, pp. 56–69. SIAM (2020). https://doi.org/10.1137/1.9781611976007.5
- 4 SPP. Angriman, E., et al.: Guidelines for experimental algorithmics: a case study in network analysis. Algorithms 12(7), 127 (2019). https://doi.org/10.3390/a12070127
- 5 SPP. Angriman, E., van der Grinten, A., Meyerhenke, H.: Local search for group closeness maximization on big graphs. In: BigData, pp. 711–720. IEEE (2019). https://doi.org/10. 1109/BigData47090.2019.9006206
- 6 SPP. Angriman, E., Predari, M., van der Grinten, A., Meyerhenke, H.: Approximation of the diagonal of a Laplacian's pseudoinverse for complex network analysis. In: ESA, pp. 6:1–6:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020). https://doi.org/10. 4230/LIPIcs.ESA.2020.6

- Bader, D.A., Meyerhenke, H., Sanders, P., Schulz, C., Kappes, A., Wagner, D.: Benchmarking for graph clustering and partitioning. In: Alhajj, R., Rokne, J. (eds.) Encyclopedia of Social Network Analysis and Mining, pp. 73–82. Springer, New York (2014). https://doi.org/10.1007/978-1-4614-6170-8_23
- Barabási, A.L., Pósfai, M.: Network Science. Cambridge University Press, Cambridge (2016)
- 9. Bastian, M., Heymann, S., Jacomy, M.: Gephi: an open source software for exploring and manipulating networks. In: ICWSM. The AAAI Press (2009)
- 10. Battista, G.D., Eades, P., Tamassia, R., Tollis, I.G.: Graph Drawing: Algorithms for the Visualization of Graphs. Prentice-Hall, Hoboken (1999)
- 11. Bavelas, A.: A mathematical model for group structures. Hum. Organ. 7(3), 16–30 (1948)
- 12 SPP. Bergamini, E., Borassi, M., Crescenzi, P., Marino, A., Meyerhenke, H.: Computing topk closeness centrality faster in unweighted graphs. ACM Trans. Knowl. Discov. Data 13(5), 53:1–53:40 (2019). https://doi.org/10.1145/3344719
- 13 SPP. Bergamini, E., Crescenzi, P., D'Angelo, G., Meyerhenke, H., Severini, L., Velaj, Y.: Improving the betweenness centrality of a node by adding links. ACM J. Exp. Algorithmics 23, 1–32 (2018). https://doi.org/10.1145/3166071
- 14 SPP. Bergamini, E., Gonser, T., Meyerhenke, H.: Scaling up group closeness maximization. In: ALENEX, pp. 209–222. SIAM (2018). https://doi.org/10.1137/1.9781611975055.18
- 15 SPP. Bergamini, E., Meyerhenke, H.: Approximating betweenness centrality in fully dynamic networks. Internet Math. 12(5), 281–314 (2016). https://doi.org/10.1080/15427951. 2016.1177802
- 16 SPP. Bergamini, E., Meyerhenke, H., Ortmann, M., Slobbe, A.: Faster betweenness centrality updates in evolving networks. In: SEA, pp. 23:1–23:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017). https://doi.org/10.4230/LIPIcs.SEA.2017.23
- 17 SPP. Bergamini, E., Wegner, M., Lukarski, D., Meyerhenke, H.: Estimating current-flow closeness centrality with a multigrid Laplacian solver. In: CSC, pp. 1–12. SIAM (2016). https://doi.org/10.1137/1.9781611974690.ch1
- 18 SPP. Bisenius, P., Bergamini, E., Angriman, E., Meyerhenke, H.: Computing top-k closeness centrality in fully-dynamic graphs. In: ALENEX, pp. 21–35. SIAM (2018). https://doi. org/10.1137/1.9781611975055.3
- 19 SPP. Bläsius, T., Friedrich, T., Katzmann, M., Meyer, U., Penschuck, M., Weyand, C.: Efficiently generating geometric inhomogeneous and hyperbolic random graphs. In: ESA, pp. 21:1–21:14. Schloss Dagstuhl Leibniz-Zentrum für Informatik (2019). https://doi.org/10.4230/LIPIcs.ESA.2019.21
 - Blondel, V.D., Guillaume, J., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. J. Stat. Mech: Theory Exp. 2008(10), P10008 (2008). https://doi. org/10.1088/1742-5468/2008/10/p10008
 - 21. Boldi, P., Vigna, S.: Axioms for centrality. Internet Math. **10**(3–4), 222–262 (2014). https://doi.org/10.1080/15427951.2013.865686
 - Borassi, M., Natale, E.: KADABRA is an adaptive algorithm for betweenness via random approximation. ACM J. Exp. Algorithmics 24(1), 1.2:1–1.2:35 (2019). https://doi. org/10.1145/3284359
 - Brandes, U.: A faster algorithm for betweenness centrality. J. Math. Sociol. 25(2), 163– 177 (2001). https://doi.org/10.1080/0022250X.2001.9990249
 - 24. Brandes, U., et al.: On modularity clustering. IEEE Trans. Knowl. Data Eng. **20**(2), 172–188 (2008). https://doi.org/10.1109/TKDE.2007.190689
- 25 SPP. Brandes, U., Hamann, M., Strasser, B., Wagner, D.: Fast quasi-threshold editing. In: Bansal, N., Finocchi, I. (eds.) ESA 2015. LNCS, vol. 9294, pp. 251–262. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48350-3_22

- 26. Brandes, U., Robins, G., McCranie, A., Wasserman, S.: What is network science? Netw. Sci. 1(1), 1–15 (2013). https://doi.org/10.1017/nws.2013.2
- Brin, S., Page, L.: Reprint of: the anatomy of a large-scale hypertextual web search engine. Comput. Netw. 56(18), 3825–3833 (2012). https://doi.org/10.1016/j.comnet. 2012.10.007
- 28 SPP. Carstens, C.J., Hamann, M., Meyer, U., Penschuck, M., Tran, H., Wagner, D.: Parallel and I/O-efficient randomisation of massive networks using global curveball trades. In: ESA, pp. 11:1–11:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2018). https://doi.org/10.4230/LIPIcs.ESA.2018.11
 - Chen, C., Wang, W., Wang, X.: Efficient maximum closeness centrality group identification. In: Cheema, M.A., Zhang, W., Chang, L. (eds.) ADC 2016. LNCS, vol. 9877, pp. 43–55. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46922-5_4
 - Crescenzi, P., D'Angelo, G., Severini, L., Velaj, Y.: Greedily improving our own closeness centrality in a network. ACM Trans. Knowl. Discov. Data 11(1), 9:1–9:32 (2016). https://doi.org/10.1145/2953882
 - Everett, M.G., Borgatti, S.P.: The centrality of groups and classes. J. Math. Sociol. 23(3), 181–201 (1999). https://doi.org/10.1080/0022250X.1999.9990219
 - 32. Fortunato, S., Hric, D.: Community detection in networks: a user guide. Phys. Rep. 659, 1–44 (2016). https://doi.org/10.1016/j.physrep.2016.09.002
- 33 SPP. Glantz, R., Meyerhenke, H.: Many-to-many correspondences between partitions: introducing a cut-based approach. In: SDM, pp. 1–9. SIAM (2018). https://doi.org/10.1137/ 1.9781611975321.1
 - Gleiser, P.M., Danon, L.: Community structure in jazz. Adv. Complex Syst. 6(4), 565– 574 (2003). https://doi.org/10.1142/S0219525903001067
- 35 SPP. van der Grinten, A., Angriman, E., Meyerhenke, H.: Scaling up network centrality computations - a brief overview. IT - Inf. Technol. 62(3–4), 189–204 (2020). https://doi.org/ 10.1515/itit-2019-0032
- 36 SPP. Grinten, A., Angriman, E., Meyerhenke, H.: Parallel adaptive sampling with almost no synchronization. In: Yahyapour, R. (ed.) Euro-Par 2019. LNCS, vol. 11725, pp. 434– 447. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29400-7_31
- 37 SPP. van der Grinten, A., Angriman, E., Predari, M., Meyerhenke, H.: New approximation algorithms for forest closeness centrality - for individual vertices and vertex groups. In: SDM, pp. 136–144. SIAM (2021)
- 38 SPP. van der Grinten, A., Bergamini, E., Green, O., Bader, D.A., Meyerhenke, H.: Scalable Katz ranking computation in large static and dynamic graphs. In: ESA, pp. 42:1– 42:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2018). https://doi.org/10. 4230/LIPIcs.ESA.2018.42
- 39 SPP. van der Grinten, A., Meyerhenke, H.: Scaling betweenness approximation to billions of edges by MPI-based adaptive sampling. In: IPDPS, pp. 527–535. IEEE (2020). https:// doi.org/10.1109/IPDPS47924.2020.00061
- 40 SPP. Hamann, M., Lindner, G., Meyerhenke, H., Staudt, C.L., Wagner, D.: Structurepreserving sparsification methods for social networks. Soc. Netw. Anal. Min. **6**(1), 22:1– 22:22 (2016). https://doi.org/10.1007/s13278-016-0332-2
- 41 SPP. Hamann, M., Röhrs, E., Wagner, D.: Local community detection based on small cliques. Algorithms **10**(3), 90 (2017). https://doi.org/10.3390/a10030090
 - Herman, I., Melançon, G., Marshall, M.S.: Graph visualization and navigation in information visualization: a survey. IEEE Trans. Vis. Comput. Graph. 6(1), 24–43 (2000). https://doi.org/10.1109/2945.841119
- 43 SPP. Hoske, D., Lukarski, D., Meyerhenke, H., Wegner, M.: Engineering a combinatorial Laplacian solver: lessons learned. Algorithms 9(4), 72 (2016). https://doi.org/10.3390/ a9040072

- 44. Jacomy, M., Venturini, T., Heymann, S., Bastian, M.: ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software. PLoS ONE **9**(6), e98679 (2014)
- 45 SPP. Kepner, J., et al.: Mathematical foundations of the GraphBLAS. In: HPEC, pp. 1–9. IEEE (2016). https://doi.org/10.1109/HPEC.2016.7761646
- 46 SPP. Koch, J., Staudt, C.L., Vogel, M., Meyerhenke, H.: An empirical comparison of big graph frameworks in the context of network analysis. Soc. Netw. Anal. Min. **6**(1), 84:1–84:20 (2016). https://doi.org/10.1007/s13278-016-0394-1
 - Kreutel, J.: Augmenting network analysis with linked data for humanities research. In: Kremers, H. (ed.) Digital Cultural Heritage, pp. 1–14. Springer, Cham (2020). https:// doi.org/10.1007/978-3-030-15200-0_1
 - Lancichinetti, A., Fortunato, S., Radicchi, F.: Benchmark graphs for testing community detection algorithms. Phys. Rev. E 78, 046110 (2008). https://doi.org/10.1103/ PhysRevE.78.046110
 - 49. Langville, A.N., Meyer, C.D.: Google's PageRank and Beyond The Science of Search Engine Rankings. Princeton University Press, Princeton (2006)
 - 50. Leskovec, J., Rajaraman, A., Ullman, J.D.: Mining of Massive Datasets, 2nd edn. Cambridge University Press, Cambridge (2014)
 - Livne, O.E., Brandt, A.: Lean algebraic multigrid (LAMG): fast graph Laplacian linear solver. SIAM J. Sci. Comput. 34(4), B499–B522 (2012). https://doi.org/10.1137/ 110843563
- 52 SPP. von Looz, M., Meyerhenke, H.: Querying probabilistic neighborhoods in spatial data sets efficiently. In: Mäkinen, V., Puglisi, S.J., Salmela, L. (eds.) IWOCA 2016. LNCS, vol. 9843, pp. 449–460. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-44543-4_35
- 53 SPP. von Looz, M., Meyerhenke, H.: Updating dynamic random hyperbolic graphs in sublinear time. ACM J. Exp. Algorithmics 23, 1–30 (2018). https://doi.org/10.1145/3195635
- 54 SPP. von Looz, M., Meyerhenke, H., Prutkin, R.: Generating random hyperbolic graphs in subquadratic time. In: Elbassioni, K., Makino, K. (eds.) ISAAC 2015. LNCS, vol. 9472, pp. 467–478. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48971-0_40
- 55 SPP. von Looz, M., Özdayi, M.S., Laue, S., Meyerhenke, H.: Generating massive complex networks with hyperbolic geometry faster in practice. In: HPEC, pp. 1–6. IEEE (2016). https://doi.org/10.1109/HPEC.2016.7761644
- 56 SPP. von Looz, M., Wolter, M., Jacob, C.R., Meyerhenke, H.: Better partitions of protein graphs for subsystem quantum chemistry. In: Goldberg, A.V., Kulikov, A.S. (eds.) SEA 2016. LNCS, vol. 9685, pp. 353–368. Springer, Cham (2016). https://doi.org/10.1007/ 978-3-319-38851-9_24
 - Mahmoody, A., Tsourakakis, C.E., Upfal, E.: Scalable betweenness centrality maximization via sampling. In: KDD, pp. 1765–1773. ACM (2016). https://doi.org/10.1145/ 2939672.2939869
- 58 SPP. Meyerhenke, H., Nöllenburg, M., Schulz, C.: Drawing large graphs by multilevel maxent-stress optimization. IEEE Trans. Vis. Comput. Graph. 24(5), 1814–1827 (2018). https://doi.org/10.1109/TVCG.2017.2689016
 - Mocnik, F.B.: The polynomial volume law of complex networks in the context of local and global optimization. Sci. Rep. 8(1), 1–10 (2018). https://doi.org/10.1038/s41598-018-29131-0
 - Mocnik, F.-B., Frank, A.U.: Modelling spatial structures. In: Fabrikant, S.I., Raubal, M., Bertolotto, M., Davies, C., Freundschuh, S., Bell, S. (eds.) COSIT 2015. LNCS, vol. 9368, pp. 44–64. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23374-1_3

- 61. Nastos, J., Gao, Y.: Familial groups in social networks. Soc. Netw. **35**(3), 439–450 (2013). https://doi.org/10.1016/j.socnet.2013.05.001
- 62. Newman, M.: Networks, 2nd edn. Oxford University Press, Oxford (2018)
- 63. Newman, M.E.J., Girvan, M.: Finding and evaluating community structure in networks. Phys. Rev. E **69**, 026113 (2004). https://doi.org/10.1103/PhysRevE.69.026113
- 64 SPP. Nocaj, A., Ortmann, M., Brandes, U.: Untangling the hairballs of multi-centered, small-world online social media networks. J. Graph Algorithms Appl. 19(2), 595–618 (2015). https://doi.org/10.7155/jgaa.00370
 - Peel, L., Larremore, D.B., Clauset, A.: The ground truth about metadata and community detection in networks. Sci. Adv. 3(5), e1602548 (2017). https://doi.org/10.1126/sciadv. 1602548
 - Riondato, M., Kornaropoulos, E.M.: Fast approximation of betweenness centrality through sampling. In: WSDM, pp. 413–422. ACM (2014). https://doi.org/10.1145/ 2556195.2556224
 - 67. Rochat, Y.: Closeness centrality extended to unconnected graphs: the harmonic centrality index. In: ASNA, Applications of Social Network Analysis (2009)
 - Rotta, R., Noack, A.: Multilevel local search algorithms for modularity clustering. ACM J. Exp. Algorithmics 16, 27 (2011). https://doi.org/10.1145/1963190.1970376
 - Satuluri, V., Parthasarathy, S., Ruan, Y.: Local graph sparsification for scalable clustering. In: SIGMOD Conference, pp. 721–732. ACM (2011). https://doi.org/10.1145/ 1989323.1989399
- 70 SPP. Şimşek, M., Meyerhenke, H.: Combined centrality measures for an improved characterization of influence spread in social networks. J. Complex Netw. 8(1), cnz048 (2020). https://doi.org/10.1093/comnet/cnz048
- 71 SPP. Staudt, C., Marrakchi, Y., Meyerhenke, H.: Detecting communities around seed nodes in complex networks. In: BigData, pp. 62–69. IEEE Computer Society (2014). https:// doi.org/10.1109/BigData.2014.7004373
 - Staudt, C., Meyerhenke, H.: Engineering high-performance community detection heuristics for massive graphs. In: ICPP, pp. 180–189. IEEE Computer Society (2013). https://doi.org/10.1109/ICPP.2013.27
- 73 SPP. Staudt, C.L., Hamann, M., Gutfraind, A., Safro, I., Meyerhenke, H.: Generating realistic scaled complex networks. Appl. Netw. Sci. 2, 36 (2017). https://doi.org/10.1007/ s41109-017-0054-z
- 74 SPP. Staudt, C.L., Meyerhenke, H.: Engineering parallel algorithms for community detection in massive networks. IEEE Trans. Parallel Distrib. Syst. 27(1), 171–184 (2016). https:// doi.org/10.1109/TPDS.2015.2390633
- 75 SPP. Staudt, C.L., Sazonovs, A., Meyerhenke, H.: NetworKit: a tool suite for large-scale complex network analysis. Netw. Sci. 4(4), 508–530 (2016). https://doi.org/10.1017/nws. 2016.20
- 76 SPP. Wegner, M., Taubert, O., Schug, A., Meyerhenke, H.: Maxent-stress optimization of 3D biomolecular models. In: ESA, pp. 70:1–70:15. Schloss Dagstuhl Leibniz-Zentrum für Informatik (2017). https://doi.org/10.4230/LIPIcs.ESA.2017.70
- 77 SPP. Zweig, K.A.: Network Analysis Literacy A Practical Approach to the Analysis of Networks. Lecture Notes in Social Networks. Springer, Vienna (2016). https://doi.org/ 10.1007/978-3-7091-0741-6

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

