

Recent Advances in Fully Dynamic Graph Algorithms

Kathrin Hanauer  

Faculty of Computer Science, Universität Wien, Austria

Monika Henzinger  

Faculty of Computer Science, Universität Wien, Austria

Christian Schulz  

Faculty of Mathematics and Computer Science, Universität Heidelberg, Germany

Abstract

In recent years, significant advances have been made in the design and analysis of fully dynamic algorithms. However, these theoretical results have received very little attention from the practical perspective. Few of the algorithms are implemented and tested on real datasets, and their practical potential is far from understood. Here, we present a quick reference guide to recent engineering and theory results in the area of fully dynamic graph algorithms.

2012 ACM Subject Classification General and reference → Surveys and overviews; Networks → Network dynamics; Mathematics of computing → Graph algorithms

Keywords and phrases fully dynamic graph algorithms, survey

Digital Object Identifier 10.4230/LIPIcs.SAND.2022.1

Category Invited Talk

Funding This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 101019564, “The Design of Modern Fully Dynamic Data Structures (MoDynStruct)”), as well as from the Austrian Science Fund (FWF) and netIDEE SCIENCE project P 33775-N. Moreover, we have been partially supported by DFG grant SCHU 2567/1-2.



1 Introduction

A (fully) dynamic graph algorithm is a data structure that supports edge insertions, edge deletions, and answers certain queries that are specific to the problem under consideration. There has been a lot of research on dynamic algorithms for graph problems that are solvable in polynomial time by a static algorithm. The most studied dynamic problems are graph problems such as connectivity, reachability, shortest paths, or matching (see [115]). Typically, any dynamic algorithm that can handle edge insertions can be used as a static algorithm by starting with an empty graph and inserting all m edges of the static input graph step-by-step. A fundamental question that arises is which problems can be *fully dynamized*, which boils down to the question whether they admit a dynamic algorithm that supports updates in $\mathcal{O}(T(m)/m)$ time, where $T(m)$ is the static running time. Thus, for static problems that can be solved in near-linear time, the research community is interested in near-constant time updates. By now, such results have been achieved for a wide range of problems [115], which resulted in a rich algorithmic toolbox spanning a wide range of techniques. However, while there is a large body of theoretical work on efficient dynamic graph algorithms, until recently there has been very little on their empirical evaluation. For some classical dynamic algorithms, experimental studies have been performed, such as for fully dynamic graph clustering [76] and fully dynamic approximation of betweenness centrality [33]. However, for



© Kathrin Hanauer, Monika Henzinger, and Christian Schulz;
licensed under Creative Commons License CC-BY 4.0

1st Symposium on Algorithmic Foundations of Dynamic Networks (SAND 2022).

Editors: James Aspnes and Othon Michail; Article No. 1; pp. 1:1–1:47

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

other fundamental dynamic graph problems, the theoretical algorithmic ideas have received very little attention from the practical perspective. In particular, very little work has been devoted to engineering such algorithms and providing efficient implementations in practice. Previous surveys on the topic [249, 10] are more than twenty years old and do not capture the state-of-the-field anymore. In this work, we aim to survey recent progress in theory as well as in the empirical evaluation of fully dynamic graph algorithms and summarize methodologies used to evaluate such algorithms. Moreover, we point to theoretical results that we think have a good potential for practical implementations. Hence, this paper should help an unfamiliar reader by providing most recent references for various problems in fully dynamic graph algorithms. Lastly, there currently is a lack of fully dynamic real-world graphs available online – most of the instances that can be found to date are insertions-only. Hence, together with this survey we will also start a new open-access graph repository that provides fully dynamic graph instances^{1,2}.

We want to point out that there are also various dynamic graph models which we cannot discuss in any depth for space limitations. These are insertions-only algorithms, deletions-only algorithms, offline dynamic algorithms, algorithms with vertex insertions and deletions, kinetic algorithms, temporal algorithms, algorithms with a limit on the number of allowed queries, algorithms for the sliding-windows model, and algorithms for sensitivity problems (also called emergency planning or fault-tolerant algorithms). . We also exclude dynamic algorithms in other models of computation such as distributed algorithms and algorithms in the massively parallel computation (MPC) model. If the full graph is known at preprocessing time and vertices are “switched on and off”, this is called the *subgraph model*, whereas *algorithms under failures* deal with the case that vertices or edges are only “switched off”. We do not discuss these algorithms either.

Note that fully dynamic graph algorithms (according to our definition) are also sometimes called *algorithms for evolving graphs* or *for incremental graphs* or sometimes even *maintaining a graph online*.

2 Preliminaries

Let $G = (V, E)$ be a (un)directed graph with vertex set V and edge set E . Throughout this paper, let $n = |V|$ and $m = |E|$. The *density* of G is $d = \frac{m}{n}$. In the directed case, an edge $(u, v) \in E$ has *tail* u and *head* v and u and v are said to be *adjacent*. (u, v) is said to be an *outgoing* edge or *out-edge* of u and an *incoming* edge or *in-edge* of v . The *outdegree* $\deg^+(v)$ /*indegree* $\deg^-(v)$ /*degree* $\deg(v)$ of a vertex v is its number of (out-/in-) edges. The *out-neighborhood* (*in-neighborhood*) of a vertex u is the set of all vertices v such that $(u, v) \in E$ ($(v, u) \in E$). In the undirected case, $N(v) := \{u : \{v, u\} \in E\}$ denotes the *neighbors* of v . The degree of a vertex v is $\deg(v) := |N(v)|$ here. In the following, Δ denotes the maximum degree that can be found in any state of the dynamic graph. Our focus in this paper are *fully dynamic graphs*, where the number of vertices is fixed, but edges can be added and removed. We use $\tilde{O}(\cdot)$ to hide polylogarithmic factors.

¹ If you have access to fully dynamic instances, we are happy to provide them in our repository.

² <https://DynGraphLab.github.io>

2.1 Conditional Lower Bounds

There are lower bounds for fully dynamic graph algorithms based on various popular conjectures initiated by [183, 3, 117]. These lower bounds usually involve three parameters: the preprocessing time $p(m, n)$, the update time $u(m, n)$, and the query time $q(m, n)$. We will use the notation $(p(m, n), u(m, n), q(m, n))$ below to indicate that no algorithm with preprocessing time at most $p(m, n)$ exists that requires at most update time $u(m, n)$ and query time $q(m, n)$. Note that if the preprocessing time is larger than $p(m, n)$ or if the query time is larger than $q(m, n)$, then it might be possible to achieve an update time better than $u(m, n)$. In the same vein, if the preprocessing time is larger than $p(m, n)$ or if the update time is larger than $u(m, n)$, then it might be possible to achieve a query time better than $q(m, n)$. We will write *poly* to denote any running time that is *polynomial* in the size of the input.

Any conditional lower bound that is based on the OMv (Online Boolean Matrix-Vector Multiplication) conjecture [117] applies to both the (amortized or worst-case) running time of any fully dynamic algorithm *and also* to the worst-case running time of insertions-only and deletions-only algorithms. We will not mention this for each problem below and only state the lower bound, except in cases where as a result of the lower bound only algorithms for the insertions-only or deletions-only setting have been studied.

3 Fully Dynamic Graph Algorithms

In this section, we describe recent efforts in fully dynamic graph algorithms. We start by describing fundamental problems that we think belong to a basic toolbox of fully dynamic graph algorithms: strongly connected components, minimum spanning trees, cycle detection/topological ordering, matching, core decomposition, subgraph detection, diameter, as well as independent sets. Later on, we discuss problems that are closer to the application side. To this end we include fully dynamic algorithms for shortest paths, maximum flows, graph clustering, centrality measures, and graph partitioning.

3.1 (Strongly) Connected Components and BFS/DFS Trees

One of the most fundamental questions on graphs is whether two given vertices are connected by a path. In the undirected case, a path connecting two vertices u and w is a sequence of edges $\mathcal{P} = (\{u, v_0\}, \{v_0, v_1\}, \dots, \{v_k, w\})$. A *connected component* is a maximal set of vertices that are pairwise connected by a path. A graph is *connected* if there is exactly one connected component, which is V . In a directed graph, we say that a vertex u can *reach* a vertex w if there is a directed path from u to w , i.e., a sequence of directed edges $\mathcal{P} = ((u, v_0), (v_0, v_1), \dots, (v_k, w))$. A *strongly connected component* (SCC) is a maximal set of vertices that can reach each other pairwise. A directed graph is *strongly connected* if there is just one strongly connected component, which is V . The *transitive closure* of a graph G is a graph on the same vertex set with an edge $(u, w) \in V \times V$ if and only if u can reach w in G . Given an undirected graph, we can construct a directed graph from it by replacing each undirected edge $\{u, w\}$ by a pair of directed edges (u, w) and (w, u) and translate queries of connectedness into reachability queries on the directed graph. A breadth-first search (BFS) or depth-first search (DFS) traversal of a directed or undirected graph defines a rooted, spanning subtree that consists of the edges via which a new vertex was discovered. Apart from connectivity or reachability, BFS and DFS trees can be used to answer a variety of problems on graphs, such as testing bipartiteness, shortest paths in the unweighted setting, 2-edge connectivity, or biconnectivity.

Undirected Graphs (Connectivity)

Patrascu and Demaine [184] gave an (unconditional) lower bound of $\Omega(\log n)$ per operation for this problem, improving a bound of $\Omega(\log n / \log \log n)$ [123]. The first non-trivial dynamic algorithms for connectivity, and also for 2-edge connectivity, and 2-vertex connectivity [86, 121, 79, 80, 122] took time $\tilde{O}(\sqrt{n})$ per operation. Henzinger and King [125] were the first to give a fully dynamic algorithm with polylogarithmic time per operation for this problem. Their algorithm is, however, randomized. Holm et al. [127] gave the first deterministic fully dynamic algorithm with polylogarithmic time per operation. The currently fastest fully dynamic connectivity algorithm takes $\mathcal{O}(\log n (\log \log n)^2)$ amortized expected time per operation [132]. There also is a batch-dynamic parallel algorithm that answers k queries in $\mathcal{O}(k \log(1 + n/k))$ expected work and $\mathcal{O}(\log n)$ depth with $\mathcal{O}(\log n \log(1 + n/B))$ expected amortized work per update and $\mathcal{O}(\log^3 n)$ depth for an average batch size of B [6].

The fully dynamic connectivity problem can be reduced to the maintenance of a spanning forest, using, e.g., dynamic trees [222, 7] or Euler tour trees [124, 232] (see also Section 3.2), for the components. If the graph is a forest, updates and queries can be processed in amortized $\mathcal{O}(\log n)$ time, whereas the theoretically fastest algorithms [141] to date for general graphs have polylogarithmic worst-case update time and $\mathcal{O}(\log n / \log \log n)$ worst-case query time, the latter matching the lower bound [123, 168]. The key challenge on general graphs is to determine whether the deletion of an edge of the spanning forest disconnects the component or whether a replacement edge can be found. There are also fully dynamic algorithms for more refined notions of connectivity: Two-edge connectivity [125, 126] and two-vertex connectivity [126] can also be maintained in polylogarithmic time per operation. See [134] for a survey on that topic.

Building on an earlier study by Alberts et al. [10], Iyer et al. [137] experimentally compared the Euler tour tree-based algorithms by Henzinger and King [124] and Holm et al. [126] to each other as well as several heuristics to achieve speedups in both candidates. The instances used in the evaluation were random graphs with random edge insertions and deletions, random graphs where a fixed set of edges appear and disappear dynamically, graphs consisting of cliques of equal size plus a set of inter-clique edges, where only the latter are inserted and deleted, as well as specially crafted worst-case instances for the algorithms. The authors showed that the running time of both algorithms can be improved distinctly via heuristics; in particular a sampling approach to replace deleted tree edges has proven to be successful. The experimental running time of both algorithms was comparable, but with the heuristics, the algorithm by Holm et al. [126] performed better.

Baswana et al. [25] gave the first algorithm for maintaining an undirected DFS tree with $o(m)$ update time and showed a conditional lower bound of $\Omega(n)$ on the update time in case of vertex updates and, if the tree is maintained explicitly, an unconditional lower bound of $\Omega(n)$ under edge updates. Their algorithm has a preprocessing time of $\mathcal{O}(m \log n)$, a worst-case update time of $\mathcal{O}(\sqrt{mn} \log^{2.5} n)$, and uses $\mathcal{O}(m \log^2 n)$ bits. Nakamura and Sadakane [172] improved the update time by $\text{polylog } n$ factors and the space required to $\mathcal{O}(m \log n)$. Recently, Baswana et al. [27] further reduced the update time down to $\mathcal{O}(\sqrt{mn} \log n)$. A parallel algorithm that uses m processors and $\mathcal{O}(\text{polylog } n)$ update time was given by Khan [145]. To the best of our knowledge, experimental evaluations have only been conducted to date with algorithms designed for the *incremental* setting, but not for fully-dynamic algorithms. No experimental studies on dynamically maintaining BFS trees are known to us.

Directed Graphs (Reachability, Strong Connectivity, Transitive Closure)

For directed graphs that are and remain acyclic, the same algorithms can be employed for reachability as for (undirected) connectivity in forests (see above). On general graphs, there is a conditional lower bound of $(\text{poly}, m^{1/2-\delta}, m^{1-\delta})$ for any small constant $\delta > 0$ based on the OMv conjecture. This bound even holds for the s - t reachability problem, where both s and t are fixed for all queries. The currently fastest algorithms for transitive closure are three Monte Carlo algorithms with one-sided error: Two by Sankowski [205] with $\mathcal{O}(1)$ or $\mathcal{O}(n^{0.58})$ worst-case query time and $\mathcal{O}(n^2)$ or $\mathcal{O}(n^{1.58})$ worst-case update time, respectively, and one by van den Brand, Nanongkai, and Saranurak [234] with $\mathcal{O}(n^{1.407})$ worst-case update and worst-case query time. There exists a conditional lower bound based on a variant of the OMv conjecture that shows that these running times are optimal [234]. Moreover, there are two deterministic, combinatorial algorithms: Roditty's algorithm with constant query time and $\mathcal{O}(n^2)$ amortized update time [198], as well as one by Roditty and Zwick [201] with an improved $\mathcal{O}(m + n \log n)$ amortized update time at the expense of $\mathcal{O}(n)$ worst-case query time.

Frigioni et al. [89] and later Krommidas and Zaroliagis [153] empirically studied the performance of an extensive number of algorithms for transitive closure, including those mentioned above. They also developed various extensions and variations and compared them not only to each other, but also to static, so-called “simple-minded” algorithms such as breadth-first and depth-first search. Their evaluation included random Erdős-Rényi graphs, specially constructed hard instances, as well as two instances based on real-world graphs. It showed that the “simple-minded” algorithms could outperform the dynamic ones distinctly and up to several factors, unless the query ratio was more than 65 % or the instances were dense random graphs.

In recent experimental studies by Hanauer et al. [110, 109], two relatively straightforward algorithms for single-source reachability could outperform the “simple-minded” algorithms of the earlier studies in a single-source setting by several orders of magnitude in practice both on random graphs as well as on real-world instances: SI maintains an arbitrary reachability tree which is re-constructed via a combined forward and backward breadth-first search traversal on edge deletions if necessary and is especially fast if insertions predominate, which can be handled in $\mathcal{O}(n + m)$ time. By contrast, it may take up to $\mathcal{O}(nm)$ time for a single edge removal. SES is an extension and simplification of Even-Shiloach trees [220], which originally only handle edge deletions. Its strength are hence instances with many deletions. As a plus, it is able to deliver not just any path as a witness for reachability, but even the shortest path (with respect to the number of edges). Furthermore, it internally maintains a BFS tree, which makes it viable also for numerous other applications, see above. Its worst-case update time is $\mathcal{O}(n + m)$, and, like SI, it answers queries in constant time. One key ingredient for the superior performance of both algorithms in practice are carefully chosen criteria for an abortion of the re-construction of their data structures and their re-building from scratch [110]. To query the transitive closure of a graph, a number of so-called “supportive vertices”, for which both in- and out-reachability trees are maintained explicitly, can be picked either once or periodically anew and then be used to answer both positive and negative reachability queries between a number of pairs of vertices decisively in constant time [109]. The fallback routine can be a simple static graph traversal and therefore be relatively expensive: With a random initial choice of supportive vertices and no periodic renewals, this approach has been shown to answer a great majority of reachability queries on both random and real-world instances in constant time already if the number of supportive vertices is very small, i.e., two or three.

These experimental studies clearly show the limitations of worst-case analysis: All implemented algorithms are fully dynamic with at least linear worst-case running time per operation and, thus, all perform “(very) poor” in the worst case. Still on all graphs used in the study the relatively simple new algorithms clearly outperformed the algorithms used in previous studies.

Yang et al. [246] were the first to give a fully dynamic algorithm for maintaining a DFS tree in a directed graph along with several optimizations to achieve speedups in practice. In an experimental evaluation on twelve real-world instances, they showed that the optimized version of their algorithm can handle edge insertions and deletions within few seconds on average for instances with millions of vertices. With regard to BFS trees, the already mentioned SES algorithm [110] is the only fully dynamic algorithm we are aware of that maintains a BFS tree on a directed graph.

3.2 Minimum Weight Spanning Trees

A minimum weight spanning tree (MST) of a connected graph is a subset of the edges such that all nodes are connected via the edges in the subset, the induced subgraph has no cycles and, lastly, has the minimum total weight among all possible subsets fulfilling the first two properties.

The lower bound of $\Omega(\log n)$ [184] on the time per operation for connectivity trivially extends to maintaining the weight of a minimum spanning tree. Holm et al. [127] gave the first fully dynamic algorithm with polylogarithmic time per operation for this problem. It was later slightly improved to $\mathcal{O}(\log^4 n / \log \log n)$ time per operation [128].

Amato et al. [133] presented the first experimental study of dynamic minimum spanning tree algorithms. In particular, the authors implemented different versions of Frederickson’s algorithm [85] which uses partitions and topology trees. The algorithms have been adapted with sparsification techniques to improve their performance. The update running times of these algorithms range from $\mathcal{O}(m^{2/3})$ to $\mathcal{O}(m^{1/2})$. The authors further presented a variant of Frederickson’s algorithm that is significantly faster than all other implementations of this algorithm. However, the authors also proposed a simple adaption of a partially dynamic data structure of Kruskal’s algorithm that was the fastest implementation on random inputs. Later, Cattaneo et al. [56, 57] presented an experimental study on several algorithms for the problem. The authors presented an efficient implementation of the algorithm of Holm et al. [127], proposed new simple algorithms for dynamic MST that are not as asymptotically efficient as the algorithm by Holm et al. but seem to be fast in practice, and lastly compared their algorithms with the results of Amato et al. [133]. The algorithm by Holm et al. uses a clever refinement of a technique by Henzinger and King [119] for developing fully dynamic algorithms starting from the deletions-only case. One outcome of their experiments is that simple algorithms outperform the theoretically more heavy algorithms on random and worst-case networks. On the other hand, on k -clique inputs, i.e. graphs that contain k cliques of size c plus $2k$ randomly chosen inter-clique edges, the implementation of the algorithm by Holm et al. outperformed the simpler algorithms.

Tarjan and Werneck [227] performed experiments for several variants of dynamic trees data structure. The evaluated data structures have been used by Ribero and Toso [196], who focused on the case of changing weights, i.e. the edges of the graph are constant, but the edge weights can change dynamically. The authors also proposed and used a new data structure for dynamic tree representation called DRD-trees. In their algorithm the dynamic tree data structure is used to speed up connectivity queries that check whether two vertices belong to different subtrees. More generally, the authors compared different types of data structures

to do this task. In particular, the authors used the dynamic tree data structures that have been evaluated by Tarjan and Werneck [227]. The experimental evaluation demonstrated that the new structure reduces the computation time observed for the algorithm of Cattaneo et al. [56], and at the same time yielded the fastest algorithms in the experiments.

3.3 Cycle Detection and Topological Ordering

A cycle in a (directed) graph $G = (V, E)$ is a non-empty path $\mathcal{P} = (v_1, \dots, v_k = v_1)$ such that $(v_i, v_{i+1}) \in E$. A topological ordering of a directed graph is a linear ordering of its vertices from 1 to n such that for every directed edge (u, v) from vertex u to vertex v , u is ordered before v . In the static case, one can use a depth-first search (DFS) to compute a topological ordering of a directed acyclic graph or to check if a (un)directed graph contains a cycle.

Let $\delta > 0$ be any small constant. Based on the OMv conjecture [117] it is straightforward to construct a lower bound of $(\text{poly}, m^{1/2-\delta}, m^{1-\delta})$ for the (amortized or worst-case) running time of any fully dynamic algorithm that detects whether the graph contains any cycle. As any algorithm for topological ordering can be used to decide whether a graph contains a cycle, this lower bound also applies to any fully dynamic topological ordering algorithm. Via dynamic matrix inverse one can maintain fully dynamic directed cycle detection in $\mathcal{O}(n^{1.407})$ [234], which is conditionally optimal based on a variant of the OMv conjecture.

Pearce and Kelly [187, 188] were the first to evaluate algorithms for topological ordering in the presence of edge insertions and deletions. In their work, the authors compared three algorithms that can deal with the online topological ordering problem. More precisely, the authors implemented the algorithms by Marchetti-Spaccamela et al. [164] and Alpern et al. [12] as well as a newly developed algorithm. Their new algorithm is the one that performed best in their experiments. The algorithm maintains a node-to-index map, called $n2i$, that maps each vertex to a unique integer in $\{1 \dots n\}$ and ensures that for any edge (u, v) in G , it holds $n2i[u] < n2i[v]$. When an insertion (u, v) invalidates the topological ordering, affected nodes are updated. The set of affected nodes are identified using a forward DFS from v and backward DFS from u . The two sets are then separately sorted into increasing topological order and afterwards a remapping to the available indices is performed. The algorithm by Marchetti-Spaccamela et al. [164] is quite similar to the algorithm by Pearce and Kelly. However, it only maintains the forward set of affected nodes and obtains a correct solution by shifting the affected nodes up in the ordering (putting them after u). Alpern et al. [12] used a data structure to create new priorities between existing ones in constant worst-case time. The result by Pearce and Kelly has later been applied to online cycle detection and difference propagation in pointer analysis by Pearce et al. [189]. Furthermore, Pearce and Kelly [186] later extended their algorithm to be able to provide more efficient batch updates.

3.4 (Weighted) Matching

The matching problem is one of the most prominently studied combinatorial graph problems having a variety of practical applications. A matching \mathcal{M} of a graph $G = (V, E)$ is a subset of edges such that no two elements of \mathcal{M} have a common end point. Many applications require matchings with certain properties, like being maximal (no edge can be added to \mathcal{M} without violating the matching property) or having maximum cardinality.

In the dynamic setting, there is a conditional lower bound of $(\text{poly}, m^{1/2-\delta}, m^{1-\delta})$ (for any small constant $\delta > 0$) for the size of the maximum cardinality matching based on the OMv conjecture [117]. Of course, maintaining an actual maximum matching is only harder than

maintaining the size of a maximum matching. Thus upper bounds have mostly focused on approximately maximum matching. However, also here we have to distinguish (a) algorithms that maintain the *size* of an approximately maximum matching and (b) algorithms that maintain an *approximately maximum matching*.

- (a) Improving Sankowski's $O(n^{1.495})$ update time bound [207], van den Brand et al. [234] maintain the exact size of a maximum matching in $O(n^{1.407})$ update time. To maintain the approximate size of the maximum matching, dynamic algorithms use the duality of maximum matching and vertex cover and maintain instead a $(2 + \epsilon)$ -approximate vertex cover. This line of work lead to a sequence of papers [135, 40, 42, 39], resulting finally in a deterministic $(2 + \epsilon)$ -approximation algorithm that maintains a hierarchical graph decomposition with $O(1/\epsilon^2)$ amortized update time [47]. The algorithm can be turned into an algorithm with worst-case $O(\log^3 n)$ time per update [43].
- (b) One can trivially maintain a maximal matching in $O(n)$ update time by resolving all trivial augmenting paths, i.e. cycle-free paths that start and end on a unmatched vertex and where edges from \mathcal{M} alternate with edges from $E \setminus \mathcal{M}$, of length one. As any maximal matching is a 2-approximation of a maximum matching, this leads to a 2-approximation algorithm. Onak and Rubinfeld [181] presented a randomized algorithm for maintaining an $O(1)$ -approximate matching with $O(\log^2 n)$ expected amortized time per edge update. Baswana, Gupta, and Sen [26] gave an elegant algorithm that maintains a *maximal* matching with amortized update time $O(\log n)$. It is based on a hierarchical graph decomposition and was subsequently improved by Solomon to amortized constant expected update time [223]. For worst-case bounds, the best results are a $(1 + \epsilon)$ -approximation in $O(\sqrt{m}/\epsilon)$ update time by Gupta and Peng [104] (see [178] for a 3/2-approximation in the same time), a $(3/2 + \epsilon)$ -approximation in $O(m^{1/4}/\epsilon^{2.5})$ time by Bernstein and Stein [37], and a $(2 + \epsilon)$ -approximation in $O(\text{polylog } n)$ time by Charikar and Solomon [59] and Arar et al. [17]. Recently, Grandoni et al. [100] gave an incremental matching algorithm that achieves a $(1 + \epsilon)$ -approximate matching in constant deterministic amortized time. Finally, Bernstein et al. [36] improved the maximal matching algorithm of Baswana et al. [26] to $O(\log^5 n)$ worst-case time with high probability.

Despite this variety of different algorithms, to the best of our knowledge, there have been only limited efforts so far to engineer and evaluate these algorithms on real-world instances. Henzinger et al. [116] initiated the empirical evaluation of algorithms for this problem in practice. To this end, the authors evaluated several dynamic maximal matching algorithms as well as an algorithm that is able to maintain the maximum matching. They implemented the algorithm by Baswana, Gupta and Sen [26], which performs edge updates in $O(\sqrt{n})$ time and maintains a 2-approximate maximum matching, the algorithm of Neiman and Solomon [178], which takes $O(\sqrt{m})$ time to maintain a 3/2-approximate maximum matching, as well as two novel dynamic algorithms, namely a random walk-based algorithm as well as a dynamic algorithm that searches for augmenting paths using a (depth-bounded) blossom algorithm. Their experiments indicate that an optimum matching can be maintained dynamically more than an order of magnitude faster than the naive algorithm that recomputes maximum matchings from scratch. Second, all non-optimum dynamic algorithms that have been considered in this work were able to maintain near-optimum matchings in practice while being multiple orders of magnitudes faster than the naive exact dynamic algorithm. The study concludes that in practice an extended random walk-based algorithms is the method of choice.

For the *weighted* dynamic matching problem, Anand et al. [14] proposed an algorithm that can maintain an 4.911-approximate dynamic maximum weight matching that runs in amortized $\mathcal{O}(\log n \log C)$ time where C is the ratio of the weight of the highest weight edge to the weight of the smallest weight edge. Furthermore, a sequence [41, 1, 39, 46, 44] of work on fully dynamic set cover resulted in $(1 + \epsilon)$ -approximate weighted dynamic matching algorithms, with $\mathcal{O}(1/\epsilon^3 + (1/\epsilon^2) \log C)$ amortized and $\mathcal{O}((1/\epsilon^3) \log^2(Cn))$ worst-case time per operation based on various hierarchical hypergraph decompositions. Gupta and Peng [105] maintain a $(1 + \epsilon)$ -approximation under edge insertions/deletions that runs in time $\mathcal{O}(\sqrt{m} \epsilon^{-2 - \mathcal{O}(1/\epsilon)} \log N)$ time per update, if edge weights are in between 1 and N . Their result is based on rerunning a static algorithm from time to time, a trimming routine that trims the graph to a smaller equivalent graph whenever possible and in the weighted case, a partition of the weights of the edges into intervals of geometrically increasing size. Stubbs and Williams [225] presented metatheorems for dynamic weighted matching. Here, the authors reduced the dynamic maximum weight matching problem to the dynamic maximum cardinality matching problem in which the graph is unweighted. The authors proved that using this reduction, if there is an α -approximation for maximum cardinality matching with update time T in an unweighted graph, then there is also a $(2 + \epsilon)\alpha$ -approximation for maximum weight matching with update time $\mathcal{O}(\frac{T}{\epsilon^2} \log^2 N)$. Their basic idea is an extension of the algorithm of Crouch and Stubbs [64] who tackled the problem in the streaming model. Here, the reduction is to take matchings from weight-threshold based subgraphs of the dynamic graph, i.e. the algorithm maintains maximal matchings in $\log C$ subgraphs, where subgraph i contains all edges having weight at least $(1 + \epsilon)^i$. The resulting matchings are then greedily merged together by considering the matched edges in descending order of i (heaviest edges first). Recently, the approach by Stubbs and Williams has been evaluated experimentally and has been compared against a new random walk-based approach [16] which gives a $(1 + \epsilon)$ approximation w.h.p.. When inserting or deleting an edge, the random walk-based approach finds random simple paths (using random walks) and solves those paths using dynamic programming to improve the maintained matching. In practice, the random walk-based approach outperforms the approach by Stubbs and Williams significantly.

3.5 k -Core Decomposition

A k -core of a graph is a maximal connected subgraph in which all vertices have degree at least k . The k -core decomposition problem is to compute the core number of every node in the graph. It is well-known that a k -core decomposition can be computed in linear time for a static graph. The problem of maintaining the k -core decomposition in a fully dynamic graph has not received much attention by the theoretical computer science community: Sun et al. [226] showed that the insertion and deletion of a single edge can change the core value of all vertices. They also gave a $(4 + \epsilon)$ -approximate fully dynamic algorithm with polylogarithmic running time. The algorithm can be implemented in time $\mathcal{O}(\log^2 n)$ in graphs using the algorithm of [45]. It dynamically maintains $\mathcal{O}(\log_{(1+\epsilon)} n)$ many (α, β) -decompositions of the graph, one for each β -value that is a power of $(1 + \epsilon)$ between 1 and $(1 + \epsilon)n$. An (α, β) -decomposition of a graph $G = (V, E)$ is a decomposition Z_1, \dots, Z_L of V into $L := 1 + \lceil (1 + \epsilon) \log n \rceil$ levels such that $Z_{i+1} \subseteq Z_i$ for all $1 \leq i < L$, $Z_1 = V$, and the following invariants are maintained: (1) All vertices v on level Z_i with $\deg_{Z_i}(v) > \alpha\beta$ belong to Z_{i+1} and (2) all vertices v on level Z_i with $\deg_{Z_i}(v) < \beta$ do not belong to Z_{i+1} . There are no further lower bounds, neither conditional nor unconditional, and no faster algorithms known for maintaining an approximate k -core decomposition.

Miorandi and De Pellegrini [169] proposed two methods to rank nodes according to their k -core number in fully dynamic networks. The focus of their work is to identify the most influential spreaders in complex dynamic networks. Li et al. [157] used a filtering method to only update nodes whose core number is affected by the network update. More precisely, the authors showed that nodes that need to be updated must be connected via a path to the endpoints of the inserted/removed edge and the core number must be equal to the smaller core number of the endpoints. Moreover, the authors presented efficient algorithms to identify such nodes as well as additional techniques to reduce the size of the nodes that need updates. Similarly, Sariyüce et al. [208] proposed the k -core algorithm TRAVERSAL and gave additional rules to prune the size of the subgraphs that are guaranteed to contain the vertices whose k -core number can have changed. Note that these algorithm can have a high variation in running time for the update operations depending on the size of the affected subgraphs. Zhang et al. [250] noted that due to this reason it can be impractical to process updates one by one and introduced the k -order concept which can reduce the cost of the update operations. A k -order is defined as follows: a node u is ordered before v in the k -order if u has a smaller core number than v or when the vertices have the same core number, if the linear time algorithm to compute the core decomposition would remove u before v . A recent result by Sun et al. [226] also contains experimental results. However, their main focus is on hypergraphs and there are no comparisons against the algorithms mentioned above.

Aridhi et al. [18] gave a distributed k -core decomposition algorithm in large dynamic graphs. The authors used a graph partitioning approach to distribute the workload and pruning techniques to find nodes that are affected by the changes. Wang et al. [242] gave a parallel algorithm that appears to significantly outperform the TRAVERSAL algorithm. Jin et al. [138] presented a parallel approach based on matching to update core numbers in fully dynamic networks. Specifically, the authors showed that if a batch of inserted/deleted edges forms a matching, then the core number update step can be performed in parallel. However, the type of the edges has to be the same (i.e. only insertions, or only deletions) in each update. Hua et al. [130] noted that previous algorithms become inefficient for high superior degree vertices, i.e., vertices that have many neighbors that have a core number that is larger than its own core number. For example, the matching-based approach of Jin et al. [138] can only process one edge associated to a vertex in each iteration. Their new algorithm can handle multiple insertions/deletions per iteration.

It would be interesting to evaluate the algorithm of Sun et al. [226] which maintains a $(4 + \epsilon)$ -approximate core number, on graphs to see how far from the exact core numbers these estimates are and how its running time compares to the above approaches. Note that an (α, β) -decomposition actually gives a $(2\alpha + \epsilon)$ approximation and α has to be chosen to be slightly larger than 2 only to guarantee polylogarithmic updates. Thus, it would be interesting to also experiment with smaller values of α .

3.6 Motif Search and Motif Counting

Two graphs are isomorphic if there is a bijection between the vertex sets of the graphs that preserves adjacency. Given a graph pattern H (or multiple H_i), motif counting counts the subgraphs of G that are isomorphic to H (H_i respectively). In the work that is currently available there is a subset of work that focuses on the special case of counting triangles or wedges, i.e., paths of length two, in dynamic networks.

There is a conditional lower bound of $(\text{poly}, m^{1/2-\delta}, m^{1-\delta})$ even for the most fundamental problem of detecting whether a graph contains a triangle [117]. The same lower bound also extends to various four-vertex subgraphs [108], whereas there is a lower bound of

(poly, $m^{1-\delta}$, $m^{2-\delta}$) for counting 4-cliques as well as *induced* connected four-vertex subgraphs. A fully dynamic algorithm with $\mathcal{O}(\sqrt{m})$ update time was recently given independently by Kara et al. [142, 143] for counting triangles. Subsequently, Lu and Tao [161] studied the trade-off between update time and approximation quality and presented a new data structure for exact triangle counting whose complexity depends on the arboricity of the graph. The result by Kara et al. was also extended to general k -clique counting by Dhulipala et al. [74]. Motivated by the fact that real-world graphs in certain applications often have small h -index h (i.e., there are at most h vertices of degree at least h), Eppstein and Spiro [82] showed that the undirected triangle count can be maintained in $\mathcal{O}(h)$ time. Eppstein et al. [81] later extended this result to maintaining the counts of directed triangles in amortized $\mathcal{O}(h)$ time and of undirected four-vertex subgraphs in amortized $\mathcal{O}(h^2)$. Note that h can be as large as $\mathcal{O}(\sqrt{m})$, resulting in an amortized time complexity of $\mathcal{O}(m)$ per update for four-vertex subgraphs in general. Only very recently, Hanauer et al. [108] showed how to reduce this to amortized $\mathcal{O}(m^{2/3})$ time per update for all four-vertex subgraphs except the 4-clique. This is currently an active area of research.

In our description of the empirical work for this problem we start with recent work that mainly focuses on triangle counting. Pavan et al. [185] introduced neighborhood sampling to count and sample triangles in a one-pass streaming algorithm. In neighborhood sampling, first a random edge in the stream is sampled and in subsequent steps, edges that share an endpoint with the already sampled edges are sampled. The algorithm outperformed their implementations of the previous best algorithms for the problem, namely the algorithms by Jowhari and Ghodsi [140] and by Buriol et al. [54]. Note that the method does not appear to be able to handle edge deletions. Bulteau et al. [53] estimated the number of triangles in fully dynamic streamed graphs. Their method adapts 2-path sampling to work for dynamic graphs. The main idea of 2-path sampling is to sample a certain number of 2-paths and compute the ratio of 2-paths in the sample that are complete triangles. The total number of 2-paths in the graph is then multiplied with the ratio to obtain the total number of triangles in the graph. This approach fails, however, if one allows deletions. Thus, the contribution of the paper is a novel technique for sampling 2-paths. More precisely, the algorithm first streams the graph and sparsifies it. Afterwards, the sampling technique is applied on the sparsified graph. The core contribution of the authors is to show that the estimate obtained in the sparsified graph is similar to the number of triangles in the original graph. For graphs with constant transitivity coefficient, the authors achieve constant processing time per edge. Makkar et al. [163] presented an exact and parallel approach using an inclusion-exclusion formulation for triangle counting in dynamic graphs. The algorithm is implemented in **cuSTINGER** [84] and runs on GPUs. The algorithm computes updates for batches of edge updates and also updates the number of triangles each vertex belongs to. The **TRIEST** algorithm [224] estimates local and global triangles. An input parameter of the algorithm is the amount of available memory. The algorithm maintains a sample of the edges using reservoir sampling and random pairing to exploit the available memory as much as possible. The algorithm reduces the average estimation error by up to 90 % w.r.t. to the previous state-of-the-art. Han and Sethu [107] proposed a new sampling approach, called **edge-sample-and-discard**, which generates an unbiased estimate of the total number of triangles in a fully dynamic graph. The algorithm significantly reduces the estimation error compared to **TRIEST**. The **MASCOT** algorithm [159, 158] focuses on local triangle counting, i.e. counting the triangles adjacent to every node. In their work, the authors provide an unbiased estimation of the number of local triangles.

We now report algorithms that can count more complex patterns. The neighborhood sampling method of Pavan et al. [185] can also be used for more complex patterns, for example Pavan et al. also presented experiments for 4-cliques. Shiller et al. [212] presented the

stream-based (insertions and deletions) algorithm **StreaM** for counting undirected 4-vertex motifs in dynamic graphs. Ahmed et al. [8] presented a general purpose sampling framework for graph streams. The authors proposed a martingale formulation for subgraph count estimation and showed how to compute unbiased estimate of subgraph counts from a sample at any point during the stream. The estimates for triangle and wedge counting obtained are less than 1 % away from the true number of triangles/wedges. The algorithm outperformed their own implementation of **TRIÈST** and **MASCOT**. Mukherjee et al. [171] gave an exact counting algorithm for a given set of motifs in dynamic networks. Their focus is on biological networks. The algorithm computes an initial embedding of each motif in the initial network. Then for each motif its embeddings are stored in a list. This list is then dynamically updated while the graph evolves. Liu et al. [160] estimated motifs in dynamic networks. The algorithm uses exact counting algorithms as a subroutine, and hence can speed up any exact algorithm at the expense of accuracy. The main idea of their algorithm is to partition the stream into time intervals and find exact motif counts in subsets of these intervals. Recently, Wang et al. [241] improved on the result of Liu et al.. The improvement stems from a generic edge sampling algorithm to estimate the number of instances of any k -vertex ℓ -edge motif in a dynamic network. The main idea of the algorithm is to first uniformly at random draw random edges from the dynamic network, then exactly count the number of local motifs and lastly estimate the global count from the local counts. The experimental evaluation showed that their algorithm is up to 48.5 times faster than the previous state-of-the-art while having lower estimation errors.

Dhulipala et al. [74] recently gave parallel batch-dynamic algorithms for k -clique counting. Their first algorithm is a batch-dynamic parallel algorithm for triangle counting that has amortized work $\mathcal{O}(\Delta\sqrt{\Delta+m})$ and $\mathcal{O}(\log^*(\Delta+m))$ depth with high probability. The algorithm is based on degree thresholding which divides the vertices into vertices with low- and high-degree. Given the classification of the vertex, different updates routines are used. A multicore implementation of the triangle counting algorithm is given. Experiments indicate that the algorithms achieve 36.54 to 74.73-times parallel speedups on a machine with 72 cores. Lastly, the authors developed a simple batch-dynamic algorithm for k -clique counting that has expected $\mathcal{O}(\Delta(m+\Delta)\alpha^{k-4})$ work and $\mathcal{O}(\log^{k-2}n)$ depth with high probability, for graphs with arboricity α .

To summarize for this problem the empirical work is far ahead of the theoretical work and it would be interesting to better understand the theoretical complexity of motif search and motif counting.

3.7 Diameter

The *eccentricity* of a vertex is the greatest distance between the vertex and any other vertex in the graph. Based on this definition, the *diameter* of a graph is defined as the maximum eccentricity over all vertices in the graph. The *radius* is the minimum eccentricity of all vertices. Through recomputation from scratch it is straightforward to compute a 2-approximation for diameter and radius and a $(2+\epsilon)$ -approximation for radius in linear time.

Anacona et al. [15] recently showed that under the strong exponential time hypothesis (SETH) there can be no $(2-\epsilon)$ -approximate fully dynamic approximation algorithm for any of these problems with $\mathcal{O}(m^{1-\delta})$ update or query time for any $\delta > 0$. There also exist non-trivial (and sub- n^2 time) fully dynamic algorithms for $(1.5+\epsilon)$ approximate diameter (and also for radius and eccentricities) [234]. In this paper, the authors also construct a non-trivial algorithm for exact diameter. We are not aware of any experimental study for fully dynamic diameter.

3.8 Independent Set and Vertex Cover

Given a graph $G = (V, E)$, an *independent set* is a set $S \subseteq V$ such that no vertices in S are adjacent to one another. The *maximum independent set problem* is to compute an independent set of maximum cardinality, called a *maximum independent set* (MIS). The *minimum vertex cover* problem is equivalent to the maximum independent set problem: S is a minimum vertex cover C in G iff $V \setminus S$ is a maximum independent set $V \setminus C$ in G . Thus, an algorithm that solves one of these problems can be used to solve the other. Note, however, that this does not hold for approximation algorithms: If C' is an α -approximation of a minimum vertex cover, then $V \setminus C'$ is not necessarily an α -approximation of a maximum independent set. Another related problem is the *maximal independent set* problem. A set S is a maximal independent set if it is an independent set such that for any vertex $v \in V \setminus S$, $S \cup \{v\}$ is not independent.

As computing the size of an MIS is NP-hard, all dynamic algorithms of independent set study the maximal independent set problem. Note, however, that unlike for matching a maximal independent set does not give an approximate solution for the MIS problem, as shown by a star graph. In a sequence of papers [19, 103, 20, 60, 31] the running time for the maximal independent set problem was reduced to $\mathcal{O}(\log^4 n)$ expected worst-case update time.

While quite a large amount of engineering work has been devoted to the computation of independent sets/vertex covers in static graphs, the amount of engineering work for the dynamic independent set problem is very limited. Zheng et al. [252] presented a heuristic fully dynamic algorithm and proposed a lazy search algorithm to improve the size of the maintained independent set. A year later, Zheng et al. [251] improved the result such that the algorithm is less sensitive to the quality of the initial solution used for the evolving MIS. In their algorithm, the authors used two well known data reduction rules, degree one and degree two vertex reduction, that are frequently used in the static case. Moreover, the authors can handle batch updates. Bhore et al. [48] focused on the special case of MIS for independent rectangles which is frequently used in map labelling applications. The authors presented a deterministic algorithm for maintaining a MIS of a dynamic set of uniform rectangles with amortized sub-logarithmic update time. Moreover, the authors evaluated their approach using extensive experiments.

3.9 Shortest Paths

One of the most studied problems on weighted dynamic networks is the maintenance of shortest path information between pairs of vertices. In the most general setting, given an undirected, dynamic graph with dynamically changing edge weights representing distances, we are interested in the shortest path between two arbitrary vertices s and t (*all-pairs shortest path problem*). For the *single-source shortest path problem*, the source vertex s is fixed beforehand and the dynamic graph algorithm is only required to answer distance queries between s and an (arbitrary) vertex t which is specified by the query operation. In the *s-t shortest path problem* both s and t are fixed beforehand and the data structure is only required to return the distance between s and t as answer to a query. In all cases, the analogous problem can also be cast on a directed graph, asking for a shortest path from s to t instead.

Let $\delta > 0$ be a small constant. There is a conditional lower bound of $(\text{poly}, m^{1/2-\delta}, m^{1-\delta})$ for any small constant $\delta > 0$ based on the OMv conjecture, even for s - t shortest paths [117]. This lower bound applies also to any algorithm that gives a better than $5/3$ -approximation. For planar graphs the product of query and update time is $\Omega(n^{1-\delta})$ based on the APSP

conjecture [2]. As even the partially dynamic versions have shown to be at least as hard as the static all-pairs shortest paths problem [199, 2], one cannot hope for a *combinatorial* fully dynamic all-pairs shortest paths algorithm with $\mathcal{O}(n^{3-\delta})$ preprocessing time, $\mathcal{O}(n^{2-\delta})$ amortized update time, and constant query time. The state-of-the-art algorithms come close to this: For directed, weighted graphs, Demetrescu and Italiano [72] achieved an *amortized* update time of $\tilde{\mathcal{O}}(n^2)$, which was later improved by a polylogarithmic factor by Thorup [228]. Both of these algorithms actually allow vertex insertions and deletion, not just edge updates. There is also a fully dynamic $2^{O(k^2)}$ -approximation algorithm that takes time $\tilde{\mathcal{O}}(\sqrt{mn}^{1/k})$ per update and $\mathcal{O}(k^2)$ per update for any positive integer k [5].

With respect to *worst-case* update times, the currently fastest algorithms are randomized with $\tilde{\mathcal{O}}(n^{2+2/3})$ update time [4, 106]. Moreover, Probst Gutenberg and Wulff-Nilsen [106] presented a deterministic algorithm with $\tilde{\mathcal{O}}(n^{2+5/7})$ update time, thereby improving a 15 years old result by Thorup [229]. Van den Brand and Nanongkai [233] showed that Monte Carlo-randomized $(1+\epsilon)$ -approximation algorithms exist with $\tilde{\mathcal{O}}(n^{1.823}/\epsilon^2)$ worst-case update time for the fully dynamic single-source shortest path problem and $\tilde{\mathcal{O}}(n^{2.045}/\epsilon^2)$ for all-pairs shortest paths, in each case with positive real edge weights and constant query time. Slightly faster exact and approximative algorithms exist in part for the “special cases” of *unweighted* graphs [206, 199, 4, 106, 234, 233] (all edges have unit weight) and/or *undirected* graphs [200, 233] (every edge has a reverse edge of the same weight). More details on shortest paths algorithms including fully dynamic algorithms are given in the survey of Madkour et al. [162].

The first experimental study for fully dynamic single-source shortest paths on directed graphs with positive real edge weights was conducted by Frigioni et al. [87], who evaluated Dijkstra’s seminal static algorithm [75] against a fully dynamic algorithm by Ramalingam and Reps [195] (RR) as well as one by Frigioni et al. [88] (FMN). RR is based on Dijkstra’s static algorithm and maintains a spanning subgraph consisting of edges that belong to at least one shortest s - t path for some vertex t . After an edge insertion, the spanning subgraph is updated starting from the edge’s head until all affected vertices have been processed. In case of an edge deletion, the affected vertices are identified as a first step, followed by an update of their distances. The resulting worst-case update time is $\mathcal{O}(x_\delta + n_\delta \log n_\delta) \subseteq \mathcal{O}(m + n \log n)$, where n_δ corresponds to the number of vertices affected by the update, i.e., whose distance from s changes and x_δ equals n_δ plus the number of edges incident to an affected vertex. Similarly, Frigioni et al. [88] analyzed the update complexity of their algorithm FMN with respect to the change in the solution and showed a worst-case running time of $\mathcal{O}(|U_\delta| \sqrt{m} \log n)$, where U_δ is the set of vertices where either the distance from s must be updated or their parent in the shortest paths tree. The algorithm assigns each edge (u, v) a forward (backward) level, which corresponds to the difference between the (sum of) v ’s (u ’s) distance from s and the edge weight, as well as an owner, which is either u or v , and used to bound the running time. Incident outgoing and incoming edges of a vertex that it does not own are kept in a priority queue each, with the priority corresponding to the edge’s level. In case of a distance update at a vertex, only those edges are scanned that are either owned by the vertex or have a priority that indicates a shorter path. Edge insertion and deletion routines are based on Dijkstra’s algorithm and handled similar as in RR, but using level and ownership information. The experiments were run on three types of input instances: randomly generated ones, instances crafted specifically for the tested algorithms, and random updates on autonomous systems networks. The static Dijkstra algorithm is made dynamic in that it is re-run from scratch each time its shortest paths tree is affected by an update. The evaluation showed that the dynamic algorithms can speed up the update time by 95 %

over the static algorithm. Furthermore, RR turned out to be faster in practice than FMN except on autonomous systems instances, where the spanning subgraph was large due to many alternative shortest paths. In a follow-up work, Demetrescu et al. [71, 70] extended this study to dynamic graphs with arbitrary edge weight, allowing in particular also for negative weights. In addition to the above mentioned algorithm by Ramalingam and Reps [195] in a slightly lighter version (RRL) and the one by Frigioni et al. [88] (FMN), their study also includes a simplified variant of the latter which waives edge ownership (DFMN), as well as a rather straightforward dynamic algorithm (DF) that in case of a weight increase on an edge (u, v) first marks all vertices in the shortest paths subtree rooted at v and then finds for each of these vertices an alternative path from s using only unmarked vertices. The new weight of these vertices can be at most this distance or the old distance plus the amount of weight increase on (u, v) . Therefore, the minimum is taken as a distance estimate for the second step, where the procedure is as in Dijkstra's algorithm. In case of a weight decrease on an edge (u, v) the first step is omitted. As Dijkstra's algorithm is employed as a subroutine, the worst-case running time of DF for a weight change is $\mathcal{O}(m + n \log n)$. For updates, all algorithms use a technique introduced by Edmonds and Karp [78] to transform the weight $w(u, v)$ of each edge (u, v) to a non-negative one by replacing it with the reduced weight $w(u, v) - (d(v) - d(u))$, where $d(\cdot)$ denotes the distance from s . This preserves shortest paths and allows Dijkstra's algorithm to be used during the update process. The authors compared these dynamic algorithms to re-running the static algorithm by Bellman and Ford on each update from scratch on various randomly generated dynamic instances with mixed incremental and decremental updates on the edge weights, always avoiding negative-length cycles. Their study showed that DF is the fastest in practice on most instances, however, in certain circumstances RR and DFMN are faster, whereas FMN turned out to be too slow in practice due to its complicated data structures. The authors observed a runtime dependency on the interval size of the edge weights; RR was the fastest if this interval was small, except for very sparse graphs. DFMN on the other hand was shown to perform better than DF in presence of zero-length cycles, whereas RR is incapable of handling such instances. It is interesting to note here that the differences in running time are only due to the updates that increase distances, as all three candidates used the same routine for operations that decrease distances. The static algorithm was slower than the dynamic algorithms by several orders of magnitude.

The first fully dynamic algorithm for all-pairs shortest paths in graphs with positive integer weights less than a constant C was presented by King [147], with an amortized update time of $\mathcal{O}(n^{2.5} \sqrt{C \log n})$. For each vertex v , it maintains two shortest paths trees up to a distance d : one outbound with v as source and one inbound with v as target. A so-called stitching algorithm is used to stitch together longer paths from shortest paths of distance at most d . To achieve the above mentioned running time, d is set to $\sqrt{nC \log n}$. The space requirement is $\mathcal{O}(n^3)$ originally, but can be reduced to $\tilde{\mathcal{O}}(n^2 \sqrt{nC})$ [148]. For non-negative, real-valued edge weights, Demetrescu and Italiano [72] proposed an algorithm with an amortized update time of $\mathcal{O}(n^2 \log^3 n)$, which was later improved to $\mathcal{O}(n^2 (\log n + \log^2((n+m)/n)))$ by Thorup [228]. The algorithm uses the concept of *locally shortest paths*, which are paths such that each proper subpath is a shortest path, but not necessarily the entire path, and *historical shortest paths*, which are paths that have once been shortest paths and whose edges have not received any weight updates since then. The combination of both yields so-called locally historical paths, which are maintained by the algorithm. To keep their number small, the original sequence of updates is transformed into an equivalent, but slightly longer *smoothed sequence*. In case of a weight update, the algorithm discards all maintained paths containing the updated

edge and then computes new locally historical paths using a routine similar to Dijkstra's algorithm. Both algorithms have constant query time and were evaluated experimentally in a study by Demetrescu and Italiano [73] against RRL [71] on random instances, graphs with a single bottleneck edge, which serves as a bridge between two equally-sized complete bipartite graphs and only its weight is subject to updates, as well as real-world instances obtained from the US road networks and autonomous systems networks. Apart from RRL, the study also comprises the Dijkstra's static algorithm. Both these algorithms are designed for single-source shortest paths and were hence run once per vertex. All algorithms were implemented with small deviations from their respective theoretical description to speed them up in practice. The study showed that RRL and the algorithm based on locally historical paths (LHP) can outperform the static algorithm by a factor of up to 10 000, whereas the algorithm by King only achieves a speedup factor of around 10. RRL turned out to be especially fast if the solution changes only slightly, but by contrast exhibited the worst performance on the bottleneck instances unless the graphs were sparse. In comparison, LHP was slightly slower on sparse instances, but could beat RRL as the density increased. The authors also point out differences in performance that depend mainly on the memory architecture of the machines used for benchmarking, where RRL could better cope with small caches or memory bandwidth due to its reduced space requirements and better locality in the memory access pattern, whereas LHP benefited from larger caches and more bandwidth.

Buriol et al. [55] presented a technique that reduces the number of elements that need to be processed in a heap after an update for various dynamic shortest paths algorithms by excluding vertices whose distance changes by exactly the same amount as the weight change and handling them separately. They showed how this improvement can be incorporated into RR [195], a variant similar to RRL [195], the algorithm by King and Thorup [148] (KT), and DF [70] and achieves speedups of up to 1.79 for random weight changes and up to 5.11 for unit weight changes. Narváez et al. [173] proposed a framework to dynamize static shortest path algorithms such as Dijkstra's or Bellman-Ford [32]. In a follow-up work [174], they developed a new algorithm that fits in this framework and is based on the linear programming formulation of shortest paths and its dual, which yields the problem in a so-called *ball-and-string model*. The authors experimentally showed that their algorithm needs fewer comparisons per vertex when processing an update than the algorithms from their earlier work, as it can reuse intact substructures of the old shortest path tree.

To speed up shortest paths computations experimentally, Wagner et al. [239] introduced a concept for pruning the search space by *geometric containers*. Here, each edge (u, v) is associated with a set of vertices called container, which is a superset of all vertices w whose shortest u - w path starts with (u, v) . The authors assume that each vertex is mapped to a point in two-dimensional Euclidean space and based on this, suggest different types of geometric objects as containers, such as disks, ellipses, sectors or boxes. All types of container only require constant additional space per edge. The experimental evaluation on static instances obtained from road and railway networks showed that using the bounding box as container reduces the query time the most in comparison to running the Dijkstra algorithm without pruning, as the search space could be reduced to 5 % to 10 %. This could be preserved for dynamic instances obtained from railway networks if containers were grown and shrunk in response to an update, with a speedup factor of 2 to 3 over a recomputation of the containers from scratch. For bidirectional search, reverse containers need to be maintained additionally, which about doubled the absolute update time. Delling and Wagner [69] adapted the static ALT algorithm [90] to the dynamic setting. ALT is a variant of bidirectional A^* search that uses a small subset of vertices called *landmarks*, for which distances from and to all other

vertices are precomputed, and the triangle inequality to direct the search for a shortest path towards the target more efficiently. The authors distinguish between an eager and a lazy dynamic version of ALT, where the eager one updates all shortest path trees of the landmarks immediately after an update. The lazy variant instead keeps the preprocessed information as long as it still guarantees correctness, which holds as long as the weight of an edge is at least its initial weight, however at the expense of a potentially larger search space. The choice of landmarks remains fixed. The experimental study on large road networks showed that queries in the lazy version are almost as fast as in the eager version for short distances or if no edges representing motorways are affected, but slower by several factors for longer distances, larger changes to the weight of motorway edges, or after many updates. Schultes and Sanders [216] combined and generalized different techniques that have been successfully used in the static setting, such as separators, highway hierarchies, and transit node routing in a multi-level approach termed *highway-node routing*: For the set of vertices V_i on each level i , $V_i \subseteq V_{i-1}$, and the overlay graph G_i is defined on V_i with an edge $(s, t) \in V_i \times V_i$ iff there is a shortest s - t path in G_{i-1} that contains no vertices in V_i except for s and t . Queries are carried out by a modified Dijkstra search on this graph hierarchy. The authors extended this approach also to the dynamic setting and consider two scenarios: a server scenario, where in case of edge weight changes the sets of highway nodes V_i are kept and the graphs G_i are updated, and a mobile scenario, where only those vertices that are potentially affected are determined and the query routine needs to be aware of possibly outdated information during a search. In an experimental evaluation on a very large road network with dynamically changing travel times as weights it is shown that the dynamic highway-node routing outperformed recomputation from scratch as well as dynamic ALT search with 16 landmarks clearly with respect to preprocessing, update, and query time as well as space overhead.

Misra and Oommen [170] presented algorithms for single-source shortest paths that are based on learning automata and designed to find “statistical” shortest paths in a stochastic graph with stochastically changing edge weights. The algorithms are extensions of RR [195] and FMN [88] and shown to be superior to the original versions of RR and FMN by several orders of magnitude once they have converged. Chan and Yang [58] studied the problem of dynamically updating a single-source shortest path tree under multiple concurrent edge weight updates. They amended the algorithm by Narváez et al. [174] (MBS), for which they showed that it may misbehave in certain circumstances and suggested two further algorithms: MFP is an optimized version of an algorithm by Ramalingam and Reps [194] (DynamicSWSF-FP), which can handle multiple updates at once. The second algorithm is a generalization of the dynamic Dijkstra algorithm proposed by Narváez et al. [173]. In a detailed evaluation, they showed that an algorithm obtained by combining the incremental phase of MBS and the decremental phase of their dynamization of Dijkstra’s algorithm performed best on road networks, whereas the dynamized Dijkstra’s algorithm was best on random networks. An extensive experimental study on single-source shortest path algorithms was conducted by Bauer and Wagner [29]. They suggested several tuned variants of DynamicSWSF-FP [194] and evaluated them against FMN [88], different algorithms from the framework by Narváez et al. [173], as well as RR [195] on a diverse set of instances. The algorithms from the Narváez framework showed similar performance in case of single-edge updates and were the fastest on road networks and generated grid-like graphs. By contrast, the tuned variants of DynamicSWSF-FP behaved less consistent. RR was superior on Internet networks, whereas FMN was the slowest, especially on sparse instances. Interestingly, the authors showed that for batch updates with a set of randomly chosen edges, the algorithms behave similar as for single-edge updates, as there was almost no interference. The picture changed slightly for

simulated node failures and strongly for simulated traffic jams. RR and a tuned variant of DynamicSWSF-FP showed the best performance for simulated node failures, and two tuned variants of DynamicSWSF-FP dominated in case of simulated traffic jams. Notably, the algorithms from the Narváez framework were faster here if instead of in batches, the updates were processed one-by-one. In follow-up works, D’Andrea et al. [65] evaluated several batch-dynamic algorithms for single-source shortest paths, where the batches are homogeneous, i.e., all updates are either incremental or decremental. Their study contains RR [195], a tuned variant of DynamicSWSF-FP [194] described by Bauer and Wagner [29] (TSWSF), as well as a new algorithm DDFLP, which is designed specifically to handle homogeneous batches and uses similar techniques as FMN [88]. The instance set comprised road and Internet networks as well as randomly generated graphs according to the Erdős-Rényi model (uniform degree distribution) and the Barabási-Albert model (power-law degree distribution). Batch updates were obtained from simulated node failure and recovery, simulated traffic jam and recovery, as well as randomly selected edges for which the weights were either increased or decreased randomly. The evaluation confirmed the results by Bauer and Wagner [29] and showed that DDFLP and TSWSF are best in case of update scenarios like node failures or traffic jams and otherwise TSWSF and RR, where RR is preferable to TSWSF if the interference among the updates is low and vice versa. DDFLP generally benefited from dense instances. Singh and Khare [221] presented the first batch-dynamic parallel algorithm for single-source shortest paths for GPUs and showed in experiments that it outperforms the (sequential) tuned DynamicSWSF-FP algorithm [29] by a factor of up to 20 if the distances of up to 10 % of the nodes are affected.

For real-time shortest path computations on networks with fixed topology, but varying metric, Delling et al. [68] suggested a three-stage approach: In the first, preprocessing step, a metric-independent, moderate amount of auxiliary data is obtained from the network’s topology. It is followed by a customization step, which is run for each metric and produces few additional data. Whereas the first phase is run only once and can therefore use more computation time, the second phase must complete within seconds in real-life scenarios. Shortest path queries form the third phase and must be fast enough for actual applications. For the first, metric-independent stage, the authors describe an approach based on graph partitioning, where the number of *boundary edges*, i.e., edges between different partitions, is to be minimized. For the second stage, they compute an overlay network consisting of shortest paths between all pairs of *boundary nodes*, i.e. nodes that are incident to at least one boundary edge. An s - t query is then answered by running a bidirectional Dijkstra algorithm on the graph obtained by combining the overlay graph with the subgraphs induced by the partitions containing s and t , respectively. The authors also considered various options for speedups, such as a sparsification of the overlay network, incorporating goal-directed search techniques, and multiple levels of overlays. An experimental evaluation on road networks with travel distances and travel times as metrics showed that their approach allows for real-time queries and needs only few seconds for the metric-dependent customization phase.

Arc flags belong in the category of goal-directed techniques to speed up shortest path computations and have been successfully used in the static setting [28]. To this end, the set of vertices is partitioned into a number of regions. Each edge receives a label consisting of a flag for each region, which tells whether there is a shortest path starting with this edge and ending in the region. The technique is related to geometric containers and uses the arc flags to prune a (bidirectional) Dijkstra search. Berettini et al. [38] were the first to consider arc flags in a dynamic setting, however only for the case of weight increases. Their main idea is to maintain a threshold for each edge and region that gives the increase in weight

required for the edge to lie on a shortest path. On a weight increase, the thresholds are updated and used to determine when to change an arc flag. Although this potentially reduces the quality of the arc flags with each update, the experimental evaluation showed that the increase in query time is very small as long as the update sequence is short. With respect to the update time, a significant speedup could be achieved over recomputing arc flags from scratch. To refresh arc flags more exactly and in a fully dynamical setting, D’Angelo et al. [66] introduced a data structure called *road signs*. Road signs complement arc flags and store for each edge e and region R the set of boundary nodes contained in any shortest path starting with e and ending in R . In case of a weight increase, the algorithm first identifies all affected nodes whose shortest path to a boundary node changed and then updates all road signs for all outgoing edges of an affected node. In case of a weight decrease on edge (u, v) , the authors observed that all shortest paths containing (u, v) remain unchanged. However, shortest paths starting with other outgoing or incoming edges of u might require updates, as well as other paths containing an incoming edge of u . In an experimental study on road networks, the authors compared their algorithm against one that recomputes arc flags from scratch as well as the algorithm by Berettini et al. [38] (BDD). To mimic traffic jams and similar occurrences, the weight of a randomly chosen edge increases and then decreases by the same amount, however not necessarily in subsequent updates. The evaluation showed that updating both road signs and arc flags is by several factors faster than recomputing arc flags from scratch. On instances with weight increases only, the authors showed that their new algorithm outperforms BDD distinctly both for updates and queries.

A further speedup technique for shortest path queries are *2-hop cover labelings*, where the label $L(v)$ of each node v is a carefully chosen set of nodes U_v along with the distance between v and u for each $u \in U$. For each pair of vertices s and t , the shortest s - t path can be obtained by intersecting U_s and U_t and taking the minimum over all combinations of s - x and x - t paths for all nodes $x \in U_s \cap U_t$. In the static setting, a 2-hop cover labeling can be computed based on a breadth-first search that is run once for every vertex (“naive landmark labeling”). Akiba et al. [9] introduced *pruned landmark labeling* (PLL), which constitutes a more refined approach and uses pruned breadth-first searches instead. The authors developed an incremental algorithm for PLL, which was complemented by D’Angelo et al. [67] to a fully dynamic algorithm. The experimental evaluation showed that the algorithm achieves speedups of several orders of magnitude over a recomputation from scratch, while at the same time preserving the quality of the labeling, which makes this speedup technique suitable for practical use in dynamic scenarios.

Hayashi et al. [111] proposed a method to support shortest paths queries on unweighted networks with billions of edges by combining a bidirectional breadth-first search, which is optimized for the structure of small-world networks, with landmarks. To this end, the authors choose high-degree vertices and store shortest path trees as well as those of a subset of their neighbors in a so-called “bit-parallel” form. This increases the number of landmarks, which in turn generally speeds up the search and in particular for high-degree vertices, and at the same time keeps the memory requirements comparatively small. After an edge insertion or deletion, the bit-parallel shortest paths trees are updated accordingly. The experimental evaluation on twelve real-world instances having between 1.5 million and 3.7 billion edges showed that the new algorithm was able to process queries on average in less than 8 ms and even considerably less on many instances. The average edge insertion and deletion times were less than 1.3 ms and 8.1 s, respectively, after an initialization time of less than 1 h. The incremental algorithm by Akiba et al. [9], which was included in the study, was distinctly faster on queries, but on some instances several factors slower on insertions. However, it failed to complete the preprocessing step within 10 h or required more than 128 GB of memory on half of all instances.

3.10 Maximum Flows and Minimum Cuts

An instance of the maximum flow/minimum cut problem consists of an edge-weighted directed graph $G = (V, E, c)$ along with two distinguished vertices s and t . The edge weights c are positive and commonly referred to as *capacities*. An $(s-t)$ flow f is a non-negative weight function on the edges such $f(e) \leq c(e)$ for all $e \in E$ (capacity constraints) and except for s and t , the total flow on the incoming edges of each vertex must equal the total flow on the outgoing edges (conservation constraints). The *excess* of a vertex v is the total flow on its incoming edges minus that on its outgoing edges, which must be zero for all vertices except s and t . The value of a flow f then is the excess of t . The task is to find a flow of maximum value. An $(s-t)$ cut is a subset of edges $C \subseteq E$ whose removal makes t unreachable from s , and its value is the sum of the capacities of all edges in the cut. The well-known max-flow min-cut theorem states that the maximum value of a flow equals the minimum value of a cut. The fastest static algorithm whose running time does not depend on the size of the largest edge weight computes an optimal solution in $\mathcal{O}(nm)$ time [182].

In the dynamic setting, there is a conditional lower bound of $(\text{poly}, m^{1/2-\delta}, m^{1-\delta})$ (for any small constant $\delta > 0$) for the size of the maximum $s-t$ flow even in unweighted, undirected graphs based on the OMv conjecture [117]. Recently Chen et al. [62] gave an $\mathcal{O}(\log n \log \log n)$ -approximate fully dynamic maximum flow algorithm in time $\tilde{\mathcal{O}}(n^{2/3+o(1)})$ per update and Goranci et al. [95] gave a $n^{o(1)}$ -approximate fully dynamic algorithm in time $n^{o(1)}$ worst-case update time and $\mathcal{O}(\log^{1/6} n)$ query time. In the unweighted setting Jin and Sun [139] gave a datastructure that can be constructed for any fixed positive integer $c = (\log n)^{o(1)}$ and that answers for any pair (s, t) of vertices that are parameters of the query in time $n^{o(1)}$ where s and t are c -edge connected.

For *global minimum cuts* in the unweighted setting Thorup and Karger [231] presented a $\sqrt{2} + o(1)$ -approximation algorithm that takes polylogarithmic time per update and query and Thorup [230] designed a $(1 + \epsilon)$ -approximate algorithm in $\tilde{\mathcal{O}}(\sqrt{n})$ time per update and query.

Kumar and Gupta [154] extended the preflow-push approach [93] to solve maximum flow in static graphs to the dynamic setting. A preflow is a flow under a relaxed conservation constraint in that the excess of all vertices except s must be non-negative. Vertices with positive excess are called *active*. Preflow-push algorithms, also called push-relabel algorithms, use this relaxed variant of a flow during the construction of a maximum flow along with distance labels on the vertices. Generally speaking, they push flow out of active vertices towards vertices with smaller distance (to t) and terminate with a valid flow (i.e., observing conservation constraints). In case of an edge insertion or deletion, Kumar and Gupta first identify affected vertices via forward and backward breadth-first search while observing and updating distance labels and then follow the scheme of a basic preflow-push algorithm, however restricted to the set of affected vertices. The authors evaluated their algorithm only for the incremental setting on a set of randomly generated instances against the static preflow-push algorithm in [93] and found that their algorithm is able to reduce the number of push and relabel operations significantly as long as the instances are sparse and the number of affected vertices remains small.

Many important fields of application for the maximum flow/minimum cut problem stem from computer vision. In this area, the static algorithm of Boykov and Kolmogorov [50] (BK) is widely used due to its good performance in practice on computer vision instances and despite its pseudo-polynomial worst-case running time of $\mathcal{O}(nm \cdot \text{OPT})$, with OPT being the value of a maximum flow/minimum cut. Interestingly, however, a study by Verma and Batra [238] shows that its practical superiority only holds for sparse instances. BK follows

the Ford-Fulkerson method of augmenting flow along s - t paths, but uses two search trees grown from s and t , respectively, to find such paths. Kohli and Torr [149, 150] extended BK to the fully dynamic setting by updating capacities and flow upon changes and discuss an optimization that tries to recycle the search trees. They experimentally compared their algorithm to repeated executions of the static algorithm on dynamic instances obtained from video sequences and achieve a substantial speedup. They also observed that reusing the search trees leads to longer s - t paths, which affects the update time negatively as the instances undergo many changes.

Goldberg et al. [91] developed EIBFS, a generalization of their earlier algorithm IBFS, that by contrast also extends to the dynamic setting in a straightforward manner. IBFS in turn is a modification of BK that ensures that the two trees grown from s and t are height-minimal (i.e., BFS trees) and is closely related to the concept of blocking flows. The running time of EIBFS in the static setting and thus the initialization in the dynamic setting, is $\mathcal{O}(mn \log(n^2/m))$ with dynamic trees or $\mathcal{O}(mn^2)$ without. The algorithm works with a so-called pseudoflow, which observes capacity constraints, but may violate conservation constraints. It maintains two vertex-disjoint forests S and T , where the roots are exactly those vertices with a surplus of incoming flow and those with a surplus of outgoing flow, respectively, and originally only contain s and t . The steps of the algorithm consist in growth steps, where S or T are grown level-wise, augmentation steps, which occur if a link between the forests has been established and flow is pushed to a vertex in the other forest and further on to the root, and adoption steps, where vertices in T with surplus incoming flow or vertices in S with surplus outgoing flow are either adopted by a new parent in the same forest or become a root in the other forest. In case of an update in the dynamic setting, the invariants of the forests are restored and flow is pushed where possible, followed by alternating augmentation and adoption steps if necessary. The authors also mention that resetting the forests every $\mathcal{O}(m)$ work such that they contain only vertices with a surplus outgoing or incoming flow seemed to be beneficial in practice. In their experimental evaluation of EIBFS against the algorithm by Kohli and Torr as well as an altered version thereof and a more naive dynamization of IBFS, they showed for different dynamic real-world instances from the field of computer vision that EIBFS is the fastest on eight out of fourteen instances and relatively robust: In contrast to its competitors, it always takes at most roughly twice the time of the fastest algorithm on an instance. Notably, no algorithm is dominated by another one across all instances.

Zhu et al. [253] described a dynamic update strategy based on augmenting and de-augmenting paths as well as cancelling cyclic flows. The latter serves as a preparatory step and only reroutes flow in the network without increasing or decreasing the total s - t flow and is only necessary in a decremental update operation. They experimentally evaluated the effectiveness of their algorithm for online semi-supervised learning, where real-world big data is classified via minimum cuts, and showed that their algorithm outperforms state-of-the-art stream classification algorithms. A very similar algorithm was proposed by Greco et al. [101]. The authors compared it experimentally against EIBFS and the dynamic BK algorithm by Kohli and Torr as well as a number of the currently fastest static algorithms. Their experiments were conducted on a set of instances from computer vision where equally many edges are randomly chosen to be inserted and deleted, respectively. They showed that their algorithm is with one exception always the fastest on average in performing edge insertions if compared to the average update time of the competitors, and on half of all instances also in case of edge deletions. On the remaining instances, the average update time of EIBFS dominated.

For the global minimum cut problem, Henzinger et al. [118] implemented an algorithm for large dynamic graphs under both edge insertions and deletions. For edge insertions, the algorithm uses the approach of Henzinger [120] and Goranci et al. [94], which maintain a compact data structure of all minimum cuts in a graph and invalidate only the minimum cuts that are affected by an edge insertion. For edge deletions, the algorithms use the push-relabel algorithm of Goldberg and Tarjan [92] to certify whether the previous minimum cut is still a minimum cut. The algorithm outperformed static approaches by up to five orders of magnitude on large graphs.

3.11 Graph Clustering

Graph clustering is the problem of detecting tightly connected regions of a graph. More precisely, a clustering \mathcal{C} is a partition of the set of vertices, i.e. a set of disjoint *clusters* of vertices V_1, \dots, V_k such that $V_1 \cup \dots \cup V_k = V$. However, k is usually not given in advance and some objective function that models intra-cluster density versus inter-cluster sparsity, is optimized. It is common knowledge that there is neither a single best strategy nor objective function for graph clustering, which justifies a plethora of existing approaches. Moreover, most quality indices for graph clusterings have turned out to be NP-hard to optimize and are rather resilient to effective approximations, see, e.g. [21, 51], allowing only heuristic approaches for optimization. There has been a wide-range of algorithms for static graph clustering, the majority are based on the paradigm of intra-cluster density versus inter-cluster sparsity. For dynamic graphs, there has been a recent survey on the topic of community detection [202]. The survey covers features and challenges of dynamic community detection and classifies published approaches. Here we focus on engineering results and extend their survey in that regard with additional references as well as results that appeared in the meantime. Most algorithms in the area optimize for modularity. Modularity has recently been proposed [179]. The core idea for modularity is to take coverage, i.e. the fraction of edges covered by clusters, minus the expected value of the same quantity in a network with the same community divisions, but random connections between the vertices. The commonly used formula is as follows: $\text{mod}(\mathcal{C}) := \frac{m(\mathcal{C})}{m} - \frac{1}{4m^2} \sum_{C \in \mathcal{C}} (\sum_{v \in C} \deg(v))^2$.

Miller and Eliassi-Rad [167] adapted a dynamic extension of Latent Dirichlet Allocation for dynamic graph clustering. Latent Dirichlet Allocation has been originally proposed for modeling text documents, i.e. the algorithm assumes that a given set of documents can be classified into k topics. This approach has been transferred to graphs [113] and was adapted by the authors for dynamic networks. Aynaud and Guillaume [23] tracked communities between successive snapshots of the input network. They first noted that using standard community detection algorithms results in stability issues, i.e. little modifications of the network can result in wildly different clusterings. Hence, the authors propose a modification of the Louvain method to obtain stable clusterings. This is done by modifying the initialization routine of the Louvain method. By default, the Louvain method starts with each node being in its own clustering. In the modified version of Aynaud and Guillaume, the algorithm keeps the clustering of the previous time step and uses this as a starting point for the Louvain method which results in much more stable clusterings. Bansal et al. [24] also reused the communities from previous time steps. However, their approach is based on greedy agglomeration where two communities are merged at each step to optimize the modularity objective function. The authors improved the efficiency of dynamic graph clustering algorithms by limiting recomputation to regions of the network and merging processes that have been affected by insertion and deletion operations. Görke et al. [96] showed that the structure of minimum s - t -cuts in a graph allows for efficient updates of

clusterings. The algorithm builds on partially updating a specific part of a minimum-cut tree and is able to maintain a clustering fulfilling a provable quality guarantee, i.e. the clusterings computed by the algorithm are guaranteed to yield a certain expansion. To the best of our knowledge, this is the only dynamic graph clustering algorithm that provides such a guarantee. Later, Görke et al. [99, 98] formally introduced the concept of smoothness to compare consecutive clusterings and provided a portfolio of different update strategies for different types of local and global algorithms. Moreover, their fastest algorithm is guaranteed to run in time $\Theta(\log n)$ per update. Their experimental evaluation indicates that dynamically maintaining a clustering of a dynamic random network saves time and at the same time also yields higher modularity than recomputation from scratch. Alvari et al. [13] proposed a dynamic game theory method to tackle the community detection problem in dynamic social networks. Roughly speaking, the authors model the process of community detection as an iterative game performed in a dynamic multiagent environment where each node is an agent who wants to maximize its total utility. In each iteration, an agent can decide to join, switch, leave, or stay in a community. The new utility is then computed by the best outcome of these operations. The authors use neighborhood similarity to measure structural similarity and optimize for modularity. The experimental evaluation is limited to two graphs. Zakrzweska and Bader [248] presented two algorithms that update communities. Their first algorithm is similar to the dynamic greedy agglomeration algorithm by Görke et al. [98]. The second algorithm is a modification of the first one that runs faster. This first is achieved by more stringent backtracking of merges than Görke et al. [98], i.e. merges are only undone if the merge has significantly changed the modularity score of the clustering. Moreover, the authors used a faster agglomeration scheme during update operations that uses information about previous merges to speed up contractions. Recently, Zhuang et al. [254] proposed the DynaMo algorithm which also is a dynamic algorithm for modularity maximization, however the algorithm processes network changes in batches.

3.12 Centralities

We will describe three popular measures to find central nodes in networks in the fully dynamic setting: Katz centrality, betweenness centrality and closeness centrality. The only two theoretical fully dynamic results that we are aware of are due to Pontecorvi and Ramachandran [191], who achieve amortized $\mathcal{O}(\nu^* \cdot \log^2 n)$ update time for betweenness centrality where ν^* bounds the number of distinct edges that lie on shortest paths through any single vertex, and a result due to van den Brand and Nanongkai [233], who present a $(1 + \epsilon)$ -approximate fully-dynamic algorithm for closeness centrality with $\mathcal{O}(n^{1.823})$ update time. This is an obvious area for future work.

Katz Centrality

Katz centrality is a centrality metric that measure the relation between vertices by counting weighted walks between them. Nathan and Bader [177] were the first to look at the problem in a dynamic setting. At that time, static algorithms mostly used linear algebra-based techniques to compute Katz scores. The dynamic version of their algorithm incrementally updates the scores by exploiting properties of iterative solvers, i.e. Jacobi iterations. Their algorithm achieved speedups of over two orders of magnitude over the simple algorithms that perform static recomputation every time the graph changes. Later, they improved their algorithm [176] to handle updates by using an alternative, agglomerative method of calculating Katz scores. While their static algorithm is already several orders of magnitude

faster than typical linear algebra approaches, their dynamic algorithm is also faster than pure static recomputation every time the graph changes. A drawback of the algorithms by Nathan and Bader is that they are unable to reproduce the exact Katz ranking after dynamic updates. Van der Grinten et al. [235] fixed this problem by presenting a dynamic algorithm that iteratively improves upper and lower bounds on the centrality scores. The computed scores are approximate, but the bounds guarantee the correct ranking. The dynamic algorithm improves over the static recomputation of the Katz rankings as long as the size of the batches in the update sequence is smaller than 10 000.

Betweenness Centrality

Given a graph and a vertex v in the graph, the betweenness centrality measure is defined to be $c(v) = \sum_{u,w, u \neq w} \frac{\sigma_{u,w}(v)}{\sigma_{u,w}}$, where $\sigma_{u,w}$ is the number of shortest paths between u and w and $\sigma_{u,w}(v)$ is the number of shortest paths between u and w that include v . Statically computing betweenness centrality involves solving the all-pairs shortest path problem. Dynamically maintaining betweenness centrality is challenging as the insertion or deletion of a single edge can lead to changes of many shortest paths in the graph. The QUBE algorithm [156] was the first to provide a non-trivial update routine. The key idea is to perform the betweenness computation on a reduced set of vertices, i.e. the algorithm first finds vertices whose centrality index might have changed. Betweenness centrality is then only computed on the first set of vertices. However, QUBE is limited to the insertion and deletion of non-bridge edges. Lee et al. [155] extended the QUBE algorithm [156] to be able to insert and delete non-bridge edges. Moreover, the authors reduced the number of shortest paths that need to be recomputed and thus gained additional speedups over QUBE. Kourtellis et al. [152, 151] contributed an algorithm that maintains both vertex and edge betweenness centrality. Their algorithm needs less space than the algorithm by Green et al. [102] as it avoids storing predecessor lists. Their method can be parallelized and runs on top of parallel data processing engines such as Hadoop. Bergamini et al. [35] presented an incremental *approximation* algorithm for the problem which is based on the first theory result that is asymptotically faster than recomputing everything from scratch due to Nasre et al. [175]. As a building block of their algorithm, the authors used an asymptotically faster algorithm for the dynamic single-source shortest path problem and additionally sample shortest paths. Experiments indicate that the algorithm can be up to four orders of magnitude faster compared to restarting the static approximation algorithm by Riondato and Kornaropoulos [197]. In the same year, the authors extended their algorithm to become a fully dynamic approximation algorithm for the problem [33, 34]. In addition to dynamic single-source shortest paths, the authors also employed an approximation of the vertex diameter that is needed to compute the number of shortest paths that need to be sampled as a function of a given error guarantee that should be achieved. Hayashi et al. [112] provided a fully dynamic approximation algorithm that is also based on sampling. In contrast to Bergamini et al. [35, 33, 34], which selects between each pair of sampled vertices, the authors save all the paths between each sampled pair of vertices. Moreover, the shortest paths are represented in a data structure called hypergraph sketch. To further reduce the running time when handling unreachable pairs, the authors maintain a reachability index. Gil-Pons [190] focused on exact betweenness in incremental graphs. The author presented a space-efficient algorithm with linear space complexity. Lastly, Chehreghani et al. [61] focused on the special case in which the betweenness of a single node has to be maintained under updates.

Closeness Centrality

Given a graph and a vertex v , the harmonic closeness centrality measure is defined as $\text{clo}(v) = \sum_{u \in V, u \neq v} \frac{1}{d(u,v)}$ where $d(u,v)$ is the distance between u and v . Roughly speaking, it is the sum of the reciprocal length of the shortest path between the node and all other nodes in the graph. Baevla's definition of closeness centrality is similarly $\sum_{v \in V} \frac{|V|-1}{d(u,v)}$. Kas et al. [144] were the first to give an *fully dynamic* algorithm for the problem. As computing closeness centrality depends on the all-pairs shortest path problem, the authors extended an existing dynamic all-pairs shortest path algorithm [193] for their problem. As the algorithm stores pairwise distances between nodes it has quadratic memory requirement. Sariyuce et al. [209] provided an algorithm that can handle insertions and deletions. In contrast to Kas et al. [144], the authors used static single-source shortest paths from each vertex. The algorithm does not need to store pairwise distances and hence requires only a linear amount of memory. Moreover, the authors observed that in scale-free networks the diameter grows proportional to the logarithm of the number of nodes, i.e. the diameter is typically small. When the graph is modified with minor updates, the diameter also tends to stay small. This can be used to limit the number of vertices that need to be updated. In particular, the authors showed that recomputation of closeness can be skipped for vertices s such that $|d(s,u) - d(s,v)| = 1$ where u, v are the endpoints of the newly inserted edge. Lastly, the authors used data reduction rules to filter vertices, i.e. real-life networks can contain nodes that have the same or similar neighborhood structure that can be merged. Later, Sariyuce et al. [210, 211] proposed a distributed memory-parallel algorithm for the problem. Yen et al. [247] proposed the fully dynamic algorithm **CENDY** which can reduce the number of internal updates to a few single-source shortest path computations necessary by using breadth-first searches. The main idea is that given an augmented rooted BFS tree of an unweighted network, edges that are inserted or deleted within the same level of the tree do not change the distances from the root to all other vertices. Putman et al. [192] provided a faster algorithm for fully dynamic harmonic closeness. The authors also used a filtering method to heavily reduce the number of computations for each incremental update. The filtering method is an extension of level-based filtering to directed and weighted networks. The dynamic algorithm by Shao et al. [219] maintains closeness centrality by efficiently detecting all affected shortest paths based on articulation points. The main observation is that a graph can be divided into a series of biconnected components which are connected by articulation points – the distances between two arbitrary vertices in the graph can be expressed as multiple distances between different biconnected components.

Bisenius et al. [49] contributed an algorithm to maintain top- k harmonic closeness in fully dynamic graphs. The algorithm is not required to compute closeness centrality for the initial graph and the memory footprint of their algorithm is linear. Their algorithm also tries to skip recomputations of vertices that are unaffected by the modifications of the graph by running breadth-first searches. Crescenzi et al. [63] gave a fully dynamic *approximation* algorithm for top- k harmonic closeness. The algorithm is based on sampling paths and a backward dynamic breadth-first search algorithm.

3.13 Graph Partitioning

Typically the graph partitioning problem asks for a partition of a graph into k blocks of about equal size such that there are few edges between them. More formally, given a graph $G = (V, E)$, we are looking for disjoint *blocks* of vertices V_1, \dots, V_k that partition V , i.e., $V_1 \cup \dots \cup V_k = V$. A *balancing constraint* demands that all blocks have weight

$c(V_i) \leq (1 + \epsilon) \lceil \frac{c(V)}{k} \rceil$ for some imbalance parameter ϵ . The most used objective is to minimize the total *cut* $\omega(E \cap \bigcup_{i < j} V_i \times V_j)$. The problem is known to be NP-hard and no constant-factor approximation algorithms exist. Thus heuristic algorithms are mostly used in practice. Dynamic graph partitioning algorithms are also known in the community as *repartitioning algorithms*. As the problem is typically not solved to optimality in practice, repartitioning involves a tradeoff between the *quality*, i.e. the number of edges in different sets of the partitioning, and the amount of vertices that need to change their block as they cause communication when physically moved between the processors as the partition is adopted. The latter is especially important when graph partitioning is used in adaptive numerical simulations. In these simulations, the main goal is to partition a model of computation and communication in which nodes model computation and edges model communication. The blocks of the partition are then fixed to a specific processing element. When the dynamic graph partitioning algorithm decides to change the blocks due to changes in the graph topology, nodes that are moved to a different block create communication in the simulation system as the underlying data needs to be moved between the corresponding processors.

Hendrikson et al. [114] tackled the repartitioning problem by introducing k virtual vertices. Each of the virtual vertices is connected to all nodes of its corresponding block. The edges get a weight α which is proportional to the migration cost of a vertex and the vertex weights of the virtual vertices are set to zero. Then an updated partition can be computed using a static partitioning algorithm since the model accounts for migration costs and edge cut size at the same time.

Schloegel et al. [213] presented heuristics to control the tradeoff between edge-cut size and vertex migration costs. The most simple algorithm is to compute a completely new partition and then determine a mapping between the blocks of the old and the new partition that minimizes migration. The more sophisticated algorithm of [213] is a multilevel algorithm based on a simple process, i.e. nodes are moved from blocks that contain too many vertices to blocks that contain not enough vertices. However, this approach often yields partitions that cut a large number of edges. The result has been improved later by combining the two approaches in the parallel partitioning tool *ParMetis* [214]. Schloegel et al. [215] later extended their algorithm to be able to handle multiple balance constraints. Hu and Blake [129] noted that diffusion processes can suffer from slow convergence and improved the performance of diffusion through the use of Chebyshev polynomials. More precisely, the diffusion process in their paper is a directed diffusion that computes a diffusion solution by solving a so-called head conduction equation while minimizing the data movement. Walshaw et al. [240] integrated a repartitioning algorithm into their parallel (meanwhile uncontinued) tool *Jostle*. The algorithm is a directed diffusion process based on the solver proposed by Hu and Blake [129]. Rotaru and Năgeli [203] extended previous diffusion-based algorithms to be able to handle heterogeneous systems. These approaches, however, have certain weaknesses: For example, in numerical applications the maximum number of boundary nodes of a block is often a better estimate of the occurring communication in the simulation than the number of edges cut. Meyerhenke and Gehweiler [165, 166] explored a disturbed diffusion process that is able to overcome some of the issues of the previous approaches. To do so, Meyerhenke adapted DIBAP, a previously developed algorithm that aims at computing well-shaped partitions. A diffusion process is called disturbed if its convergence state does not result in a balanced distribution. These processes can be helpful to find densely connected regions in the graph.

There has been also work that tackles slightly different problem formulations. Kiefer et al. [146] noted that performance in applications usually does not scale linearly with the amount of work per block due to contention on different compute components. Their

algorithm uses a simplified penalized resource consumption model. Roughly speaking, the authors introduced a penalized block weight and modified the graph partitioning problem accordingly. More precisely, a positive, monotonically increasing penalty function p is used to penalize the weight of a block based on the partition cardinality. Vaquero et al. [237] looked at the problem for distributed graph processing systems. Their approach is based on iterative vertex migration based on label propagation. More precisely, a vertex has a list of candidate blocks where the highest number of its neighbors are located. However, initial partitions are computed using hashing which does not yield high quality partitions since it completely ignores the structure of the graph. The authors did not compare their work against other state-of-the-art repartitioning algorithms, so it is unclear how well the algorithm performs compared to other algorithms. Xu et al. [245] and Nicoara et al. [180] also presented dynamic algorithms specifically designed for graph processing systems. Other approaches have focused on the edge partitioning problem [204, 131, 83] or the special case of road networks [52].

4 Dynamic Graph Systems

The methodology of the previous two sections is to engineer algorithms for specific dynamic graph problems. In contrast to this, there are also approaches that try to engineer dynamic graph systems that can be applied to a wide range of dynamic graph problems. Alberts et al. [11] started this effort and presented a software library of dynamic graph algorithms. The library is written in C++ and is an extension of the well known LEDA library of efficient data types and algorithms. The library contains algorithms for connectivity, spanning trees, single-source shortest paths and transitive closure.

A decade later Weigert et al. [243] presented a system that is able to deal with dynamic distributed graphs, i.e. in settings in which a graph is too large for the memory of a single machine and, thus, needs to be distributed over multiple machines. A user can implement a query function to implement graph queries. Based on their experiments, the system appears to scale well to large distributed graphs. Ediger et al. [77] engineered **STINGER** which is short for Spatio-Temporal Interaction Networks and Graphs Extensible Representation. **STINGER** provides fast insertions, deletions, and updates on semantic graphs that have a skewed degree distribution. The authors showed in their experiments that the system can handle 3 million updates per second on a scale-free graph with 537 million edges on a Cray XMT machine. The authors already implemented a variety of algorithms on **STINGER** including community detection, k -core extraction, and many more. Later, Feng et al. [84] presented **DISTINGER** which has the same goals as **STINGER**, but focuses on the distributed memory case, i.e. the authors presented a distributed graph representation. Vaquero et al. [236] presented a dynamic graph processing system that uses adaptive partitioning to update the graph distribution over the processors over time. This speeds up queries as a better graph distribution significantly reduces communication overhead. Experiments showed that the repartitioning heuristic (also explained in Section 3.13) improves computation performance in their system up to 50% for an algorithm that computes the estimated diameter in a graph. Sengupta et al. [217] introduced a dynamic graph analytics framework called **GraphIn**. Part of **GraphIn** is a new programming model based on the gather-apply-scatter programming paradigm that allows users to implement a wide range of graph algorithms that run in parallel. Compared to **STINGER**, the authors reported a 6.6-fold speedup. Iwabuchi et al. [136] presented an even larger speedup over **STINGER**. Their dynamic graph data store is, like **STINGER**, designed for scale-free networks. The system uses compact hash tables with high data locality. In their experiments, their system called **DegAwareRHH**, is a factor 206.5 faster than **STINGER**.

Another line of research focuses on graph analytic frameworks and data structures for GPUs. Green and Bader [84] presented **cuSTINGER**, which is a GPU extension of **STINGER** and targets NVIDIA’s CUDA supported GPUs. One drawback of **cuSTINGER** is that the system has to perform restarts after a large number of edge updates. Busato et al. fixed this issue in their system, called **Hornet**, and, thus, outperform **cuSTINGER**. Moreover, **Hornet** uses a factor of 5 to 10 less memory than **cuSTINGER**. In contrast to previous approaches, **faimGraph** due to Winter et al. [244] is able to deal with a changing number of vertices. Awad et al. [22] noted that the experiments performed by Busato et al. are missing true dynamism that is expected in real world scenarios and proposed a dynamic graph structure that uses one hash table per vertex to store adjacency lists. The system achieves speedups between 5.8 to 14.8 compared to **Hornet** and 3.4 to 5.4 compared to **faimGraph** for batched edge insertions (and slightly smaller speedups for batched edge deletions). The algorithm also supports vertex deletions, as does **faimGraph**.

5 Methodology

Currently there is a limited amount of real-world *fully* dynamic networks publicly available. There are repositories that feature a lot of real-world insertions only instances such as SNAP³ and KONECT⁴. However, since the fully dynamic instances are rarely available at the moment, we start a new graph repository that provides fully dynamic graph instances⁵. Currently, there is also very limited work on dynamic graph generators. A generator for clustered dynamic random networks has been proposed by Görke et al. [97]. Another approach is due to Sengupta [218] to generate networks for dynamic overlapping communities in networks. A generative model for dynamic networks with community structure can be found in [30]. This is a widely open topic for future work, both in terms of oblivious adversaries as well as adaptive adversaries. To still be able to evaluate fully dynamic algorithms in practice, research uses a wide range of models at the moment to turn static networks into dynamic ones. We give a brief overview over the most important ones. In *undo-based* approaches, edges of a static network are inserted in some order until all edges are inserted. In the end, $x\%$ of the last insertions are undone. The intuition here is that one wants undo changes that happened to a network and to recreate a previous state of the data structure. In *window-based* approaches, edges are inserted and have a predefined lifetime. That means an edge is deleted after a given number d of new edges have been inserted. In *remove and add* based approaches, a small fraction of random edges from a static network is removed and later on reinserted. In practice, researchers use a single edge as well as whole batches of edges. In *morphing-based* approaches, one takes two related networks and creates a sequence of edge updates such that the second network obtained after the update sequence has been applied to the first network.

³ <https://snap.stanford.edu/>

⁴ <http://konect.cc/>

⁵ <https://DynGraphLab.github.io>

References

- 1 Amir Abboud, Raghavendra Addanki, Fabrizio Grandoni, Debmalya Panigrahi, and Barna Saha. Dynamic set cover: improved algorithms and lower bounds. In Moses Charikar and Edith Cohen, editors, *Proc. of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 114–125. ACM, 2019. doi:10.1145/3313276.3316376.
- 2 Amir Abboud and Søren Dahlgaard. Popular conjectures as a barrier for dynamic planar graph algorithms. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 477–486. IEEE, 2016.
- 3 Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 434–443. IEEE, 2014. doi:10.1109/FOCS.2014.53.
- 4 Ittai Abraham, Shiri Chechik, and Sebastian Krininger. Fully dynamic all-pairs shortest paths with worst-case update-time revisited. In Philip N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 440–452. SIAM, 2017. doi:10.1137/1.9781611974782.28.
- 5 Ittai Abraham, Shiri Chechik, and Kunal Talwar. Fully dynamic all-pairs shortest paths: Breaking the $o(n)$ barrier. In Klaus Jansen, José D. P. Rolim, Nikhil R. Devanur, and Cristopher Moore, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2014, September 4-6, 2014, Barcelona, Spain*, volume 28 of *LIPIcs*, pages 1–16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014. doi:10.4230/LIPIcs.APPROX-RANDOM.2014.1.
- 6 Umut A. Acar, Daniel Anderson, Guy E. Blelloch, and Laxman Dhulipala. Parallel batch-dynamic graph connectivity. In Christian Scheideler and Petra Berenbrink, editors, *The 31st ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA 2019, Phoenix, AZ, USA, June 22-24, 2019*, pages 381–392. ACM, 2019. doi:10.1145/3323165.3323196.
- 7 Umut A. Acar, Daniel Anderson, Guy E. Blelloch, Laxman Dhulipala, and Sam Westrick. Parallel batch-dynamic trees via change propagation. In Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders, editors, *28th Annual European Symposium on Algorithms, ESA 2020, September 7-9, 2020, Pisa, Italy (Virtual Conference)*, volume 173 of *LIPIcs*, pages 2:1–2:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ESA.2020.2.
- 8 Nesreen K. Ahmed, Nick G. Duffield, Theodore L. Willke, and Ryan A. Rossi. On sampling from massive graph streams. *Proc. VLDB Endow.*, 10(11):1430–1441, 2017. doi:10.14778/3137628.3137651.
- 9 Takuya Akiba, Yoichi Iwata, and Yuichi Yoshida. Dynamic and historical shortest-path distance queries on large evolving networks by pruned landmark labeling. In Chin-Wan Chung, Andrei Z. Broder, Kyuseok Shim, and Torsten Suel, editors, *23rd International World Wide Web Conference, WWW '14, Seoul, Republic of Korea, April 7-11, 2014*, pages 237–248. ACM, 2014. doi:10.1145/2566486.2568007.
- 10 David Alberts, Giuseppe Cattaneo, and Giuseppe F. Italiano. An empirical study of dynamic graph algorithms. *ACM J. Exp. Algorithmics*, 2:5, 1997. doi:10.1145/264216.264223.
- 11 David Alberts, Giuseppe Cattaneo, Giuseppe F. Italiano, Umberto Nanni, and Christos Zaroliagis. A software library of dynamic graph algorithms. In *Proc. Workshop on Algorithms and Experiments*, pages 129–136. Citeseer, 1998.
- 12 Bowen Alpern, Roger Hoover, Barry K. Rosen, Peter F. Sweeney, and F. Kenneth Zadeck. Incremental evaluation of computational circuits. In David S. Johnson, editor, *Proc. of the First Annual ACM-SIAM Symposium on Discrete Algorithms, 22-24 January 1990, San Francisco, California, USA*, pages 32–42. SIAM, 1990. URL: <http://dl.acm.org/citation.cfm?id=320176.320180>.

- 13 Hamidreza Alvari, Alireza Hajibagheri, and Gita Reese Sukthankar. Community detection in dynamic social networks: A game-theoretic approach. In Xindong Wu, Martin Ester, and Guandong Xu, editors, *2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2014, Beijing, China, August 17-20, 2014*, pages 101–107. IEEE Computer Society, 2014. doi:10.1109/ASONAM.2014.6921567.
- 14 Abhash Anand, Surender Baswana, Manoj Gupta, and Sandeep Sen. Maintaining approximate maximum weighted matching in fully dynamic graphs. In Deepak D’Souza, Telikepalli Kavitha, and Jaikumar Radhakrishnan, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2012, December 15-17, 2012, Hyderabad, India*, volume 18 of *LIPIcs*, pages 257–266. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012. doi:10.4230/LIPIcs.FSTTCS.2012.257.
- 15 Bertie Ancona, Monika Henzinger, Liam Roditty, Virginia Vassilevska Williams, and Nicole Wein. Algorithms and hardness for diameter in dynamic graphs. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPIcs*, pages 13:1–13:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ICALP.2019.13.
- 16 Eugenio Angriman, Henning Meyerhenke, Christian Schulz, and Bora Uçar. Fully-dynamic weighted matching approximation in practice. *CoRR*, abs/2104.13098, 2021. arXiv:2104.13098.
- 17 Moab Arar, Shiri Chechik, Sarel Cohen, Cliff Stein, and David Wajc. Dynamic matching: Reducing integral algorithms to approximately-maximal fractional algorithms. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018*, pages 7:1–7:16, 2018. doi:10.4230/LIPIcs.ICALP.2018.7.
- 18 Sabeur Aridhi, Martin Brugnara, Alberto Montresor, and Yannis Velegrakis. Distributed k-core decomposition and maintenance in large dynamic graphs. In *Proc. of the 10th ACM International Conference on Distributed and Event-based Systems*, pages 161–168, 2016. doi:10.1145/2933267.2933299.
- 19 Sepehr Assadi, Krzysztof Onak, Baruch Schieber, and Shay Solomon. Fully dynamic maximal independent set with sublinear update time. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proc. of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 815–826. ACM, 2018. doi:10.1145/3188745.3188922.
- 20 Sepehr Assadi, Krzysztof Onak, Baruch Schieber, and Shay Solomon. Fully dynamic maximal independent set with sublinear in n update time. In Timothy M. Chan, editor, *Proc. of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1919–1936. SIAM, 2019. doi:10.1137/1.9781611975482.116.
- 21 G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation: Combinatorial Optimization Problems and their Approximability Properties*. Springer Science & Business Media, 2012.
- 22 Muhammad A. Awad, Saman Ashkiani, Serban D. Porumbescu, and John D. Owens. Dynamic graphs on the GPU. In *2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS), New Orleans, LA, USA, May 18-22, 2020*, pages 739–748. IEEE, 2020. doi:10.1109/IPDPS47924.2020.00081.
- 23 Thomas Aynaud and Jean-Loup Guillaume. Static community detection algorithms for evolving networks. In Lavy Libman Ariel Orda, Nidhi Hegde, editor, *8th International Symposium on Modeling and Optimization in Mobile, Ad-Hoc and Wireless Networks (WiOpt 2010), May 31 - June 4, 2010, University of Avignon, Avignon, France*, pages 513–519. IEEE, 2010. URL: <http://ieeexplore.ieee.org/document/5520221/>.

- 24 Shweta Bansal, Sanjukta Bhowmick, and Prashant Paymal. Fast community detection for dynamic complex networks. In Luciano da F. Costa, Alexandre G. Evsukoff, Giuseppe Mangioni, and Ronaldo Menezes, editors, *Complex Networks - Second International Workshop, CompleNet 2010, Rio de Janeiro, Brazil, October 13-15, 2010, Revised Selected Papers*, volume 116 of *Communications in Computer and Information Science*, pages 196–207. Springer, 2010. doi:10.1007/978-3-642-25501-4_20.
- 25 Surender Baswana, Shreejit Ray Chaudhury, Keerti Choudhary, and Shahbaz Khan. Dynamic DFS in undirected graphs: Breaking the $o(m)$ barrier. *SIAM J. Comput.*, 48(4):1335–1363, 2019. doi:10.1137/17M114306X.
- 26 Surender Baswana, Manoj Gupta, and Sandeep Sen. Fully dynamic maximal matching in $O(\log n)$ update time. *SIAM J. Comput.*, 44(1):88–113, 2015. doi:10.1137/16M1106158.
- 27 Surender Baswana, Shiv Kumar Gupta, and Ayush Tulsyan. Fault tolerant and fully dynamic DFS in undirected graphs: Simple yet efficient. In Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen, editors, *44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, August 26-30, 2019, Aachen, Germany*, volume 138 of *LIPIcs*, pages 65:1–65:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.MFCS.2019.65.
- 28 Reinhard Bauer, Daniel Delling, Peter Sanders, Dennis Schieferdecker, Dominik Schultes, and Dorothea Wagner. Combining hierarchical and goal-directed speed-up techniques for dijkstra’s algorithm. *ACM J. Exp. Algorithmics*, 15, 2010. doi:10.1145/1671970.1671976.
- 29 Reinhard Bauer and Dorothea Wagner. Batch dynamic single-source shortest-path algorithms: An experimental study. In Jan Vahrenhold, editor, *Experimental Algorithms, 8th International Symposium, SEA 2009, Dortmund, Germany, June 4-6, 2009. Proc.*, volume 5526 of *Lecture Notes in Computer Science*, pages 51–62. Springer, 2009. doi:10.1007/978-3-642-02011-7_7.
- 30 F. Becker. Generative Model for Dynamic Networks with Community Structures. Master’s Thesis, Heidelberg University, 2020.
- 31 Soheil Behnezhad, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Cliff Stein, and Madhu Sudan. Fully dynamic maximal independent set with polylogarithmic update time. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 382–405. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00032.
- 32 Richard Bellman. On a routing problem. *Quarterly of applied mathematics*, 16(1):87–90, 1958.
- 33 Elisabetta Bergamini and Henning Meyerhenke. Fully-dynamic approximation of betweenness centrality. In Nikhil Bansal and Irene Finocchi, editors, *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proc.*, volume 9294 of *Lecture Notes in Computer Science*, pages 155–166. Springer, 2015. doi:10.1007/978-3-662-48350-3_14.
- 34 Elisabetta Bergamini and Henning Meyerhenke. Approximating betweenness centrality in fully dynamic networks. *Internet Math.*, 12(5):281–314, 2016. doi:10.1080/15427951.2016.1177802.
- 35 Elisabetta Bergamini, Henning Meyerhenke, and Christian Staudt. Approximating betweenness centrality in large evolving networks. In Ulrik Brandes and David Eppstein, editors, *Proc. of the Seventeenth Workshop on Algorithm Engineering and Experiments, ALENEX 2015, San Diego, CA, USA, January 5, 2015*, pages 133–146. SIAM, 2015. doi:10.1137/1.9781611973754.12.
- 36 Aaron Bernstein, Sebastian Forster, and Monika Henzinger. A deamortization approach for dynamic spanner and dynamic maximal matching. In Timothy M. Chan, editor, *Proc. of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1899–1918. SIAM, 2019. doi:10.1137/1.9781611975482.115.
- 37 Aaron Bernstein and Cliff Stein. Faster fully dynamic matchings with small approximation ratios. In *Proc. of the 27th Symposium on Discrete Algorithms SODA*, pages 692–711. SIAM, 2016. doi:10.1137/1.9781611974331.ch50.

- 38 Emanuele Berrettini, Gianlorenzo D'Angelo, and Daniel Delling. Arc-flags in dynamic graphs. In Jens Clausen and Gabriele Di Stefano, editors, *ATMOS 2009 - 9th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems*, IT University of Copenhagen, Denmark, September 10, 2009, volume 12 of *OASICS*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, 2009. URL: <http://drops.dagstuhl.de/opus/volltexte/2009/2149>.
- 39 Sayan Bhattacharya, Deeparnab Chakrabarty, and Monika Henzinger. Deterministic dynamic matching in $O(1)$ update time. *Algorithmica*, 82(4):1057–1080, 2020. doi:10.1007/s00453-019-00630-4.
- 40 Sayan Bhattacharya, Monika Henzinger, and Giuseppe F. Italiano. Deterministic fully dynamic data structures for vertex cover and matching. *SIAM J. Comput.*, 47(3):859–887, 2018. doi:10.1137/140998925.
- 41 Sayan Bhattacharya, Monika Henzinger, and Giuseppe F. Italiano. Dynamic algorithms via the primal-dual method. *Inf. Comput.*, 261:219–239, 2018. doi:10.1016/j.ic.2018.02.005.
- 42 Sayan Bhattacharya, Monika Henzinger, and Danupon Nanongkai. New deterministic approximation algorithms for fully dynamic matching. In *Proc. of the 48th Annual Symposium on Theory of Computing*, pages 398–411. ACM, 2016. doi:10.1145/2897518.2897568.
- 43 Sayan Bhattacharya, Monika Henzinger, and Danupon Nanongkai. Fully dynamic approximate maximum matching and minimum vertex cover in $O(\log^3 n)$ worst case update time. In Philip N. Klein, editor, *Proc. of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms SODA*, pages 470–489. SIAM, 2017. doi:10.1137/1.9781611973730.54.
- 44 Sayan Bhattacharya, Monika Henzinger, and Danupon Nanongkai. A new deterministic algorithm for dynamic set cover. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 406–423. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00033.
- 45 Sayan Bhattacharya, Monika Henzinger, Danupon Nanongkai, and Charalampos E. Tsourakakis. Space- and time-efficient algorithm for maintaining dense subgraphs on one-pass dynamic streams. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *Proc. of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 173–182. ACM, 2015. doi:10.1145/2746539.2746592.
- 46 Sayan Bhattacharya, Monika Henzinger, Danupon Nanongkai, and Xiaowei Wu. An improved algorithm for dynamic set cover. *CoRR*, abs/2002.11171, 2020. arXiv:2002.11171.
- 47 Sayan Bhattacharya and Janardhan Kulkarni. Deterministically maintaining a $(2 + \epsilon)$ -approximate minimum vertex cover in $o(1/\epsilon^2)$ amortized update time. In Timothy M. Chan, editor, *Proc. of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1872–1885. SIAM, 2019. doi:10.1137/1.9781611975482.113.
- 48 Sujoy Bhore, Guangping Li, and Martin Nöllenburg. An algorithmic study of fully dynamic independent sets for map labeling. In Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders, editors, *28th Annual European Symposium on Algorithms, ESA 2020, September 7-9, 2020, Pisa, Italy (Virtual Conference)*, volume 173 of *LIPIcs*, pages 19:1–19:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ESA.2020.19.
- 49 Patrick Bisenius, Elisabetta Bergamini, Eugenio Angriman, and Henning Meyerhenke. Computing top- k closeness centrality in fully-dynamic graphs. In Rasmus Pagh and Suresh Venkatasubramanian, editors, *Proc. of the Twentieth Workshop on Algorithm Engineering and Experiments, ALENEX 2018, New Orleans, LA, USA, January 7-8, 2018*, pages 21–35. SIAM, 2018. doi:10.1137/1.9781611975055.3.
- 50 Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(9):1124–1137, 2004. doi:10.1109/TPAMI.2004.60.

- 51 U. Brandes, D. Delling, M. Gaertler, R. Gorke, M. Hoefer, Z. Nikoloski, and D. Wagner. On Modularity Clustering. *IEEE Transactions on Knowledge and Data Engineering*, 20(2):172–188, 2008. doi:10.1109/TKDE.2007.190689.
- 52 Valentin Buchhold, Daniel Delling, Dennis Schieferdecker, and Michael Wegner. Fast and stable repartitioning of road networks. In *18th International Symposium on Experimental Algorithms (SEA 2020)*, volume 160 of *LIPIcs*, pages 26:1–26:15. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.SEA.2020.26.
- 53 Laurent Bulteau, Vincent Froese, Konstantin Kutzkov, and Rasmus Pagh. Triangle counting in dynamic graph streams. *Algorithmica*, 76(1):259–278, 2016. doi:10.1007/s00453-015-0036-4.
- 54 Luciana S. Buriol, Gereon Frahling, Stefano Leonardi, Alberto Marchetti-Spaccamela, and Christian Sohler. Counting triangles in data streams. In Stijn Vansummeren, editor, *Proc. of the Twenty-Fifth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 26-28, 2006, Chicago, Illinois, USA*, pages 253–262. ACM, 2006. doi:10.1145/1142351.1142388.
- 55 Luciana S. Buriol, Mauricio G. C. Resende, and Mikkel Thorup. Speeding up dynamic shortest-path algorithms. *INFORMS J. Comput.*, 20(2):191–204, 2008. doi:10.1287/ijoc.1070.0231.
- 56 Giuseppe Cattaneo, Pompeo Faruolo, Umberto Ferraro Petrillo, and Giuseppe F. Italiano. Maintaining dynamic minimum spanning trees: An experimental study. In David M. Mount and Clifford Stein, editors, *Algorithm Engineering and Experiments, 4th International Workshop, ALENEX 2002, San Francisco, CA, USA, January 4-5, 2002, Revised Papers*, volume 2409 of *Lecture Notes in Computer Science*, pages 111–125. Springer, 2002. doi:10.1007/3-540-45643-0_9.
- 57 Giuseppe Cattaneo, Pompeo Faruolo, Umberto Ferraro Petrillo, and Giuseppe F. Italiano. Maintaining dynamic minimum spanning trees: An experimental study. *Discret. Appl. Math.*, 158(5):404–425, 2010. doi:10.1016/j.dam.2009.10.005.
- 58 Edward P. F. Chan and Yaya Yang. Shortest path tree computation in dynamic graphs. *IEEE Trans. Computers*, 58(4):541–557, 2009. doi:10.1109/TC.2008.198.
- 59 Moses Charikar and Shay Solomon. Fully dynamic almost-maximal matching: Breaking the polynomial worst-case time barrier. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018*, volume 107 of *LIPIcs*, pages 33:1–33:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ICALP.2018.33.
- 60 Shiri Chechik and Tianyi Zhang. Fully dynamic maximal independent set in expected poly-log update time. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 370–381. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00031.
- 61 Mostafa Haghir Chehreghani, Albert Bifet, and Talel Abdessalem. Dybed: An efficient algorithm for updating betweenness centrality in directed dynamic graphs. In Naoki Abe, Huan Liu, Calton Pu, Xiaohua Hu, Nesreen K. Ahmed, Mu Qiao, Yang Song, Donald Kossmann, Bing Liu, Kisung Lee, Jiliang Tang, Jingrui He, and Jeffrey S. Saltz, editors, *IEEE International Conference on Big Data, Big Data 2018, Seattle, WA, USA, December 10-13, 2018*, pages 2114–2123. IEEE, 2018. doi:10.1109/BigData.2018.8622452.
- 62 Li Chen, Gramoz Goranci, Monika Henzinger, Richard Peng, and Thatchaphol Saranurak. Fast dynamic cuts, distances and effective resistances via vertex sparsifiers. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 1135–1146. IEEE, 2020. doi:10.1109/FOCS46700.2020.00109.
- 63 Pierluigi Crescenzi, Clémence Magnien, and Andrea Marino. Finding top-k nodes for temporal closeness in large temporal graphs. *Algorithms*, 13(9):211, 2020. doi:10.3390/a13090211.
- 64 Michael Crouch and Daniel S. Stubbs. Improved streaming algorithms for weighted matching, via unweighted matching. In Klaus Jansen, José D. P. Rolim, Nikhil R. Devanur, and Cristopher Moore, editors, *Approximation, Randomization, and Combinatorial Optimization*.

- Algorithms and Techniques, APPROX/RANDOM 2014, September 4-6, 2014, Barcelona, Spain*, volume 28 of *LIPIcs*, pages 96–104. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014. doi:10.4230/LIPIcs.APPROX-RANDOM.2014.96.
- 65 Annalisa D’Andrea, Mattia D’Emidio, Daniele Frigioni, Stefano Leucci, and Guido Proietti. Dynamic maintenance of a shortest-path tree on homogeneous batches of updates: New algorithms and experiments. *ACM J. Exp. Algorithmics*, 20:1.5:1.1–1.5:1.33, 2015. doi:10.1145/2786022.
 - 66 Gianlorenzo D’Angelo, Mattia D’Emidio, and Daniele Frigioni. Fully dynamic update of arc-flags. *Networks*, 63(3):243–259, 2014. doi:10.1002/net.21542.
 - 67 Gianlorenzo D’Angelo, Mattia D’Emidio, and Daniele Frigioni. Fully dynamic 2-hop cover labeling. *ACM J. Exp. Algorithmics*, 24(1):1.6:1–1.6:36, 2019. doi:10.1145/3299901.
 - 68 Daniel Delling, Andrew V. Goldberg, Thomas Pajor, and Renato Fonseca F. Werneck. Customizable route planning. In Panos M. Pardalos and Steffen Rebennack, editors, *Experimental Algorithms - 10th International Symposium, SEA 2011, Kolimpari, Chania, Crete, Greece, May 5-7, 2011. Proc.*, volume 6630 of *Lecture Notes in Computer Science*, pages 376–387. Springer, 2011. doi:10.1007/978-3-642-20662-7_32.
 - 69 Daniel Delling and Dorothea Wagner. Landmark-based routing in dynamic graphs. In Camil Demetrescu, editor, *Experimental Algorithms, 6th International Workshop, WEA 2007, Rome, Italy, June 6-8, 2007, Proc.*, volume 4525 of *Lecture Notes in Computer Science*, pages 52–65. Springer, 2007. doi:10.1007/978-3-540-72845-0_5.
 - 70 Camil Demetrescu. *Fully Dynamic Algorithms for Path Problems on Directed Graphs*. PhD thesis, Universtita Degli Studi Di Roma, 2001. URL: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.21.8921>.
 - 71 Camil Demetrescu, Daniele Frigioni, Alberto Marchetti-Spaccamela, and Umberto Nanni. Maintaining shortest paths in digraphs with arbitrary arc weights: An experimental study. In Stefan Näher and Dorothea Wagner, editors, *Algorithm Engineering, 4th International Workshop, WAE 2000, Saarbrücken, Germany, September 5-8, 2000, Proc.*, volume 1982 of *Lecture Notes in Computer Science*, pages 218–229. Springer, 2000. doi:10.1007/3-540-44691-5_19.
 - 72 Camil Demetrescu and Giuseppe F. Italiano. A new approach to dynamic all pairs shortest paths. *J. ACM*, 51(6):968–992, 2004. doi:10.1145/1039488.1039492.
 - 73 Camil Demetrescu and Giuseppe F. Italiano. Experimental analysis of dynamic all pairs shortest path algorithms. *ACM Trans. Algorithms*, 2(4):578–601, 2006. doi:10.1145/1198513.1198519.
 - 74 Laxman Dhulipala, Quanquan C. Liu, and Julian Shun. Parallel batch-dynamic k-clique counting. *CoRR*, abs/2003.13585, 2020. arXiv:2003.13585.
 - 75 Edsger W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959. doi:10.1007/BF01386390.
 - 76 Christof Doll, Tanja Hartmann, and Dorothea Wagner. Fully-dynamic hierarchical graph clustering using cut trees. In *12th Intl. Symp. on Algorithms and Data Structures, WADS’11*, volume 6844 of *LNCS*, pages 338–349, 2011. doi:10.1007/978-3-642-22300-6_29.
 - 77 David Ediger, Robert McColl, E. Jason Riedy, and David A. Bader. STINGER: high performance data structure for streaming graphs. In *IEEE Conference on High Performance Extreme Computing, HPEC 2012, Waltham, MA, USA, September 10-12, 2012*, pages 1–5. IEEE, 2012. doi:10.1109/HPEC.2012.6408680.
 - 78 Jack R. Edmonds and Richard M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM*, 19(2):248–264, 1972. doi:10.1145/321694.321699.
 - 79 David Eppstein, Zvi Galil, Giuseppe F. Italiano, and Thomas H. Spencer. Separator based sparsification. i. planary testing and minimum spanning trees. *J. Comput. Syst. Sci.*, 52(1):3–27, 1996. doi:10.1006/jcss.1996.0002.
 - 80 David Eppstein, Zvi Galil, Giuseppe F. Italiano, and Thomas H. Spencer. Separator-based sparsification II: edge and vertex connectivity. *SIAM J. Comput.*, 28(1):341–381, 1998. doi:10.1137/S0097539794269072.

- 81 David Eppstein, Michael T. Goodrich, Darren Strash, and Lowell Trott. Extended dynamic subgraph statistics using h-index parameterized data structures. *Theor. Comput. Sci.*, 447:44–52, 2012. doi:10.1016/j.tcs.2011.11.034.
- 82 David Eppstein and Emma S. Spiro. The h-index of a graph and its application to dynamic subgraph statistics. *J. Graph Algorithms Appl.*, 16(2):543–567, 2012. doi:10.7155/jgaa.00273.
- 83 Wenfei Fan, Muyang Liu, Chao Tian, Ruiqi Xu, and Jingren Zhou. Incrementalization of graph partitioning algorithms. *Proc. VLDB Endow.*, 13(8):1261–1274, 2020. doi:10.14778/3389133.3389142.
- 84 Guoyao Feng, Xiao Meng, and Khaled Ammar. Distingr: A distributed graph data structure for massive dynamic graph processing. In *2015 IEEE International Conference on Big Data (Big Data)*, pages 1814–1822. IEEE, 2015. doi:10.1109/BigData.2015.7363954.
- 85 Greg N. Frederickson. Data structures for on-line updating of minimum spanning trees, with applications. *SIAM J. Comput.*, 14(4):781–798, 1985. doi:10.1137/0214055.
- 86 Greg N. Frederickson. Ambivalent data structures for dynamic 2-edge-connectivity and k smallest spanning trees. *SIAM J. Comput.*, 26(2):484–538, 1997. doi:10.1137/S0097539792226825.
- 87 Daniele Frigioni, Mario Ioffreda, Umberto Nanni, and Giulio Pasqualone. Experimental analysis of dynamic algorithms for the single-source shortest-path problem. *ACM J. Exp. Algorithmics*, 3:5, 1998. doi:10.1145/297096.297147.
- 88 Daniele Frigioni, Alberto Marchetti-Spaccamela, and Umberto Nanni. Fully dynamic output bounded single source shortest path problem (extended abstract). In Éva Tardos, editor, *Proc. of the Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, 28-30 January 1996, Atlanta, Georgia, USA*, pages 212–221. ACM/SIAM, 1996. URL: <http://dl.acm.org/citation.cfm?id=313852.313926>.
- 89 Daniele Frigioni, Tobias Miller, Umberto Nanni, and Christos D. Zaroliagis. An experimental study of dynamic algorithms for transitive closure. *ACM J. Exp. Algorithmics*, 6:9, 2001. doi:10.1145/945394.945403.
- 90 Andrew V. Goldberg and Chris Harrelson. Computing the shortest path: A search meets graph theory. In *Proc. of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005, Vancouver, British Columbia, Canada, January 23-25, 2005*, pages 156–165. SIAM, 2005. URL: <http://dl.acm.org/citation.cfm?id=1070432.1070455>.
- 91 Andrew V. Goldberg, Sagi Hed, Haim Kaplan, Pushmeet Kohli, Robert Endre Tarjan, and Renato F. Werneck. Faster and more dynamic maximum flow by incremental breadth-first search. In Nikhil Bansal and Irene Finocchi, editors, *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proc.*, volume 9294 of *Lecture Notes in Computer Science*, pages 619–630. Springer, 2015. doi:10.1007/978-3-662-48350-3_52.
- 92 Andrew V. Goldberg and Robert E. Tarjan. A new approach to the maximum-flow problem. *Journal of the ACM*, 35(4):921–940, 1988.
- 93 Andrew V. Goldberg and Robert Endre Tarjan. A new approach to the maximum-flow problem. *J. ACM*, 35(4):921–940, 1988. doi:10.1145/48014.61051.
- 94 Gramoz Goranci, Monika Henzinger, and Mikkel Thorup. Incremental exact min-cut in polylogarithmic amortized update time. *ACM Transactions on Algorithms (TALG)*, 14(2):1–21, 2018.
- 95 Gramoz Goranci, Harald Räcke, Thatchaphol Saranurak, and Zihan Tan. The expander hierarchy and its applications to dynamic graph algorithms. *CoRR*, abs/2005.02369, 2020. arXiv:2005.02369.
- 96 Robert Görke, Tanja Hartmann, and Dorothea Wagner. Dynamic graph clustering using minimum-cut trees. *J. Graph Algorithms Appl.*, 16(2):411–446, 2012. doi:10.7155/jgaa.00269.

- 97 Robert Görke, Roland Kluge, Andrea Schumm, Christian Staudt, and Dorothea Wagner. An efficient generator for clustered dynamic random networks. In Guy Even and Dror Rawitz, editors, *Design and Analysis of Algorithms - First Mediterranean Conference on Algorithms, MedAlg 2012, Kibbutz Ein Gedi, Israel, December 3-5, 2012. Proc.*, volume 7659 of *Lecture Notes in Computer Science*, pages 219–233. Springer, 2012. doi:10.1007/978-3-642-34862-4_16.
- 98 Robert Görke, Pascal Maillard, Andrea Schumm, Christian Staudt, and Dorothea Wagner. Dynamic graph clustering combining modularity and smoothness. *ACM J. Exp. Algorithmics*, 18, 2013. doi:10.1145/2444016.2444021.
- 99 Robert Görke, Pascal Maillard, Christian Staudt, and Dorothea Wagner. Modularity-driven clustering of dynamic graphs. In Paola Festa, editor, *Experimental Algorithms, 9th International Symposium, SEA 2010, Ischia Island, Naples, Italy, May 20-22, 2010. Proc.*, volume 6049 of *Lecture Notes in Computer Science*, pages 436–448. Springer, 2010. doi:10.1007/978-3-642-13193-6_37.
- 100 Fabrizio Grandoni, Stefano Leonardi, Piotr Sankowski, Chris Schwiegelshohn, and Shay Solomon. $(1 + \epsilon)$ -approximate incremental matching in constant deterministic amortized time. In Timothy M. Chan, editor, *Proc. of the 20th Symposium on Discrete Algorithms*, pages 1886–1898. SIAM, 2019. doi:10.1137/1.9781611975482.114.
- 101 Sergio Greco, Cristian Molinaro, Chiara Pulice, and Ximena Quintana. Incremental maximum flow computation on evolving networks. In Ahmed Seffah, Birgit Penzenstadler, Carina Alves, and Xin Peng, editors, *Proc. of the Symposium on Applied Computing, SAC 2017, Marrakech, Morocco, April 3-7, 2017*, pages 1061–1067. ACM, 2017. doi:10.1145/3019612.3019816.
- 102 Oded Green, Robert McColl, and David A. Bader. A fast algorithm for streaming betweenness centrality. In *2012 International Conference on Privacy, Security, Risk and Trust, PASSAT 2012, and 2012 International Conference on Social Computing, SocialCom 2012, Amsterdam, Netherlands, September 3-5, 2012*, pages 11–20. IEEE Computer Society, 2012. doi:10.1109/SocialCom-PASSAT.2012.37.
- 103 Manoj Gupta and Shahbaz Khan. Simple dynamic algorithms for maximal independent set and other problems. In *4th Symposium on Simplicity in Algorithms, SOSA@SODA 2021*, to appear, 2021. arXiv:1804.01823.
- 104 Manoj Gupta and Richard Peng. Fully dynamic $(1 + \epsilon)$ -approximate matchings. In *54th Symposium on Foundations of Computer Science, FOCS*, pages 548–557. IEEE Computer Society, 2013. URL: <https://ieeexplore.ieee.org/xpl/conhome/6685222/proceeding>.
- 105 Manoj Gupta and Richard Peng. Fully dynamic $(1 + \epsilon)$ -approximate matchings. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 548–557. IEEE Computer Society, 2013.
- 106 Maximilian Probst Gutenberg and Christian Wulff-Nilsen. Fully-dynamic all-pairs shortest paths: Improved worst-case time and space bounds. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2562–2574. SIAM, 2020. doi:10.1137/1.9781611975994.156.
- 107 Guyue Han and Harish Sethu. Edge sample and discard: A new algorithm for counting triangles in large dynamic graphs. In Jana Diesner, Elena Ferrari, and Guandong Xu, editors, *Proc. of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017, Sydney, Australia, July 31 - August 03, 2017*, pages 44–49. ACM, 2017. doi:10.1145/3110025.3110061.
- 108 Kathrin Hanauer, Monika Henzinger, and Qi Cheng Hua. Fully dynamic four-vertex subgraph counting. *CoRR*, abs/2106.15524, 2021. arXiv:2106.15524.
- 109 Kathrin Hanauer, Monika Henzinger, and Christian Schulz. Faster fully dynamic transitive closure in practice. In Simone Faro and Domenico Cantone, editors, *18th International Symposium on Experimental Algorithms, SEA 2020, June 16-18, 2020, Catania, Italy*, volume 160 of *LIPIcs*, pages 14:1–14:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.SEA.2020.14.

- 110 Kathrin Hanauer, Monika Henzinger, and Christian Schulz. Fully dynamic single-source reachability in practice: An experimental study. In Guy E. Blelloch and Irene Finocchi, editors, *Proc. of the Symposium on Algorithm Engineering and Experiments, ALENEX 2020, Salt Lake City, UT, USA, January 5-6, 2020*, pages 106–119. SIAM, 2020. doi:10.1137/1.9781611976007.9.
- 111 Takanori Hayashi, Takuya Akiba, and Ken-ichi Kawarabayashi. Fully dynamic shortest-path distance query acceleration on massive networks. In Snehasis Mukhopadhyay, ChengXiang Zhai, Elisa Bertino, Fabio Crestani, Javed Mostafa, Jie Tang, Luo Si, Xiaofang Zhou, Yi Chang, Yunyao Li, and Parikshit Sondhi, editors, *Proc. of the 25th ACM International Conference on Information and Knowledge Management, CIKM 2016, Indianapolis, IN, USA, October 24-28, 2016*, pages 1533–1542. ACM, 2016. doi:10.1145/2983323.2983731.
- 112 Takanori Hayashi, Takuya Akiba, and Yuichi Yoshida. Fully dynamic betweenness centrality maintenance on massive networks. *Proc. VLDB Endow.*, 9(2):48–59, 2015. doi:10.14778/2850578.2850580.
- 113 Keith Henderson and Tina Eliassi-Rad. Applying latent dirichlet allocation to group discovery in large graphs. In Sung Y. Shin and Sascha Ossowski, editors, *Proc. of the 2009 ACM Symposium on Applied Computing (SAC), Honolulu, Hawaii, USA, March 9-12, 2009*, pages 1456–1461. ACM, 2009. doi:10.1145/1529282.1529607.
- 114 Bruce Hendrickson, Robert W. Leland, and Rafael Van Driessche. Enhancing data locality by using terminal propagation. In *29th Annual Hawaii International Conference on System Sciences (HICSS-29), January 3-6, 1996, Maui, Hawaii, USA*, pages 565–574. IEEE Computer Society, 1996. doi:10.1109/HICSS.1996.495507.
- 115 Monika Henzinger. The state of the art in dynamic graph algorithms. In *44th Intl. Conf. on Current Trends in Theory and Practice of Computer Science, SOFSEM'18*, volume 10706 of *LNCS*, pages 40–44. Springer, 2018. doi:10.1007/978-3-319-73117-9_3.
- 116 Monika Henzinger, Shahbaz Khan, Richard Paul, and Christian Schulz. Dynamic matching algorithms in practice. In Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders, editors, *28th Annual European Symposium on Algorithms, ESA 2020, September 7-9, 2020, Pisa, Italy (Virtual Conference)*, volume 173 of *LIPIcs*, pages 58:1–58:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ESA.2020.58.
- 117 Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *Proc. of the forty-seventh annual ACM symposium on Theory of computing*, pages 21–30, 2015. doi:10.1145/2746539.2746609.
- 118 Monika Henzinger, Alexander Noe, and Christian Schulz. Practical fully dynamic minimum cut algorithms. *CoRR*, abs/2101.05033, 2021. arXiv:2101.05033.
- 119 Monika R Henzinger and Valerie King. Maintaining minimum spanning trees in dynamic graphs. In *International Colloquium on Automata, Languages, and Programming*, pages 594–604. Springer, 1997.
- 120 Monika Rauch Henzinger. Approximating minimum cuts under insertions. In *International Colloquium on Automata, Languages, and Programming*, pages 280–291. Springer, 1995.
- 121 Monika Rauch Henzinger. Fully dynamic biconnectivity in graphs. *Algorithmica*, 13(6):503–538, 1995. doi:10.1007/BF01189067.
- 122 Monika Rauch Henzinger. Improved data structures for fully dynamic biconnectivity. *SIAM J. Comput.*, 29(6):1761–1815, 2000. doi:10.1137/S0097539794263907.
- 123 Monika Rauch Henzinger and Michael L. Fredman. Lower bounds for fully dynamic connectivity problems in graphs. *Algorithmica*, 22(3):351–362, 1998. doi:10.1007/PL00009228.
- 124 Monika Rauch Henzinger and Valerie King. Randomized dynamic graph algorithms with polylogarithmic time per operation. In Frank Thomson Leighton and Allan Borodin, editors, *Proc. of the Twenty-Seventh Annual ACM Symposium on Theory of Computing, 29 May-1 June 1995, Las Vegas, Nevada, USA*, pages 519–527. ACM, 1995. doi:10.1145/225058.225269.

- 125 Monika Rauch Henzinger and Valerie King. Randomized fully dynamic graph algorithms with polylogarithmic time per operation. *J. ACM*, 46(4):502–516, 1999. doi:10.1145/320211.320215.
- 126 Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. In Jeffrey Scott Vitter, editor, *Proc. of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 79–89. ACM, 1998. doi:10.1145/276698.276715.
- 127 Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM*, 48(4):723–760, 2001. doi:10.1145/502090.502095.
- 128 Jacob Holm, Eva Rotenberg, and Christian Wulff-Nilsen. Faster fully-dynamic minimum spanning forest. In Nikhil Bansal and Irene Finocchi, editors, *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proc.*, volume 9294 of *Lecture Notes in Computer Science*, pages 742–753. Springer, 2015. doi:10.1007/978-3-662-48350-3_62.
- 129 Y.F. Hu and R.J. Blake. An improved diffusion algorithm for dynamic load balancing. *Parallel Computing*, 25(4):417–444, 1999. doi:10.1016/S0167-8191(99)00002-2.
- 130 Qiang-Sheng Hua, Yuliang Shi, Dongxiao Yu, Hai Jin, Jiguo Yu, Zhipeng Cai, Xiuzhen Cheng, and Hanhua Chen. Faster parallel core maintenance algorithms in dynamic graphs. *IEEE Trans. Parallel Distributed Syst.*, 31(6):1287–1300, 2020. doi:10.1109/TPDS.2019.2960226.
- 131 Jiewen Huang and Daniel Abadi. LEOPARD: lightweight edge-oriented partitioning and replication for dynamic graphs. *Proc. VLDB Endow.*, 9(7):540–551, 2016. doi:10.14778/2904483.2904486.
- 132 Shang-En Huang, Dawei Huang, Tsvi Kopelowitz, and Seth Pettie. Fully dynamic connectivity in $O(\log n(\log \log n)^2)$ amortized expected time. In Philip N. Klein, editor, *Proc. of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 510–520. SIAM, 2017. doi:10.1137/1.9781611974782.32.
- 133 Giuseppe Amato II, Giuseppe Cattaneo, and Giuseppe F. Italiano. Experimental analysis of dynamic minimum spanning tree algorithms (extended abstract). In Michael E. Saks, editor, *Proc. of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, 5-7 January 1997, New Orleans, Louisiana, USA*, pages 314–323. ACM/SIAM, 1997. URL: <http://dl.acm.org/citation.cfm?id=314161.314314>.
- 134 Giuseppe F. Italiano. Fully dynamic higher connectivity. In *Encyclopedia of Algorithms*, pages 797–800. Springer, 2016. doi:10.1007/978-1-4939-2864-4_154.
- 135 Zoran Ivkovic and Errol L. Lloyd. Fully dynamic maintenance of vertex cover. In *19th International Workshop Graph-Theoretic Concepts in Computer Science*, volume 790 of *LNCS*, pages 99–111, 1993.
- 136 Keita Iwabuchi, Scott Sallinen, Roger Pearce, Brian Van Essen, Maya Gokhale, and Satoshi Matsuoka. Towards a distributed large-scale dynamic graph data store. In *2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 892–901. IEEE, 2016. doi:10.1109/IPDPSW.2016.189.
- 137 Raj Iyer, David R. Karger, Hariharan Rahul, and Mikkel Thorup. An experimental study of polylogarithmic, fully dynamic, connectivity algorithms. *ACM J. Exp. Algorithmics*, 6:4, 2001. doi:10.1145/945394.945398.
- 138 Hai Jin, Na Wang, Dongxiao Yu, Qiang-Sheng Hua, Xuanhua Shi, and Xia Xie. Core maintenance in dynamic graphs: A parallel approach based on matching. *IEEE Trans. Parallel Distributed Syst.*, 29(11):2416–2428, 2018. doi:10.1109/TPDS.2018.2835441.
- 139 Wenyu Jin and Xiaorui Sun. Fully dynamic c-edge connectivity in subpolynomial time. *CoRR*, abs/2004.07650, 2020. arXiv:2004.07650.

- 140 Hossein Jowhari and Mohammad Ghodsi. New streaming algorithms for counting triangles in graphs. In Lusheng Wang, editor, *Computing and Combinatorics, 11th Annual International Conference, COCOON 2005, Kunming, China, August 16-29, 2005, Proc.*, volume 3595 of *Lecture Notes in Computer Science*, pages 710–716. Springer, 2005. doi:10.1007/11533719_72.
- 141 Bruce M. Kapron, Valerie King, and Ben Mountjoy. Dynamic graph connectivity in poly-logarithmic worst case time. In Sanjeev Khanna, editor, *Proc. of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1131–1142. SIAM, 2013. doi:10.1137/1.9781611973105.81.
- 142 Ahmet Kara, Hung Q. Ngo, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. Counting triangles under updates in worst-case optimal time. In Pablo Barceló and Marco Calautti, editors, *22nd International Conference on Database Theory, ICDT 2019, March 26-28, 2019, Lisbon, Portugal*, volume 127 of *LIPICs*, pages 4:1–4:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.ICDT.2019.4.
- 143 Ahmet Kara, Hung Q. Ngo, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. Maintaining triangle queries under updates. *ACM Trans. Database Syst.*, 45(3):11:1–11:46, 2020. doi:10.1145/3396375.
- 144 Miray Kas, Kathleen M. Carley, and L. Richard Carley. Incremental closeness centrality for dynamically changing social networks. In Jon G. Rokne and Christos Faloutsos, editors, *Advances in Social Networks Analysis and Mining 2013, ASONAM '13, Niagara, ON, Canada - August 25 - 29, 2013*, pages 1250–1258. ACM, 2013. doi:10.1145/2492517.2500270.
- 145 Shahbaz Khan. Near optimal parallel algorithms for dynamic DFS in undirected graphs. *ACM Trans. Parallel Comput.*, 6(3):18:1–18:33, 2019. doi:10.1145/3364212.
- 146 Tim Kiefer, Dirk Habich, and Wolfgang Lehner. Penalized graph partitioning for static and dynamic load balancing. In Pierre-François Dutot and Denis Trystram, editors, *Euro-Par 2016: Parallel Processing - 22nd International Conference on Parallel and Distributed Computing, Grenoble, France, August 24-26, 2016, Proc.*, volume 9833 of *Lecture Notes in Computer Science*, pages 146–158. Springer, 2016. doi:10.1007/978-3-319-43659-3_11.
- 147 Valerie King. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In *40th Annual Symposium on Foundations of Computer Science, FOCS '99, 17-18 October, 1999, New York, NY, USA*, pages 81–91. IEEE Computer Society, 1999. doi:10.1109/SFFCS.1999.814580.
- 148 Valerie King and Mikkel Thorup. A space saving trick for directed dynamic transitive closure and shortest path algorithms. In Jie Wang, editor, *Computing and Combinatorics, 7th Annual International Conference, COCOON 2001, Guilin, China, August 20-23, 2001, Proc.*, volume 2108 of *Lecture Notes in Computer Science*, pages 268–277. Springer, 2001. doi:10.1007/3-540-44679-6_30.
- 149 Pushmeet Kohli and Philip H. S. Torr. Dynamic graph cuts for efficient inference in markov random fields. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(12):2079–2088, 2007. doi:10.1109/TPAMI.2007.1128.
- 150 Pushmeet Kohli and Philip H. S. Torr. Dynamic graph cuts and their applications in computer vision. In Roberto Cipolla, Sebastiano Battiato, and Giovanni Maria Farinella, editors, *Computer Vision: Detection, Recognition and Reconstruction*, volume 285 of *Studies in Computational Intelligence*, pages 51–108. Springer, 2010. doi:10.1007/978-3-642-12848-6_3.
- 151 Nicolas Kourtellis, Gianmarco De Francisci Morales, and Francesco Bonchi. Scalable online betweenness centrality in evolving graphs. *IEEE Trans. Knowl. Data Eng.*, 27(9):2494–2506, 2015. doi:10.1109/TKDE.2015.2419666.
- 152 Nicolas Kourtellis, Gianmarco De Francisci Morales, and Francesco Bonchi. Scalable online betweenness centrality in evolving graphs. In *32nd IEEE International Conference on Data Engineering, ICDE 2016, Helsinki, Finland, May 16-20, 2016*, pages 1580–1581. IEEE Computer Society, 2016. doi:10.1109/ICDE.2016.7498421.

- 153 Ioannis Krommidas and Christos D. Zaroliagis. An experimental study of algorithms for fully dynamic transitive closure. *ACM J. Exp. Algorithmics*, 12:1.6:1–1.6:22, 2008. doi:10.1145/1227161.1370597.
- 154 S. Kumar and P. Gupta. An incremental algorithm for the maximum flow problem. *J. Math. Model. Algorithms*, 2(1):1–16, 2003. doi:10.1023/A:1023607406540.
- 155 Min-Joong Lee, Sunghee Choi, and Chin-Wan Chung. Efficient algorithms for updating betweenness centrality in fully dynamic graphs. *Inf. Sci.*, 326:278–296, 2016. doi:10.1016/j.ins.2015.07.053.
- 156 Min-Joong Lee, Jungmin Lee, Jaimie Yejean Park, Ryan Hyun Choi, and Chin-Wan Chung. QUBE: a quick algorithm for updating betweenness centrality. In Alain Mille, Fabien L. Gandon, Jacques Misselis, Michael Rabinovich, and Steffen Staab, editors, *Proc. of the 21st World Wide Web Conference 2012, WWW 2012, Lyon, France, April 16-20, 2012*, pages 351–360. ACM, 2012. doi:10.1145/2187836.2187884.
- 157 Rong-Hua Li, Jeffrey Xu Yu, and Rui Mao. Efficient core maintenance in large dynamic graphs. *IEEE Trans. Knowl. Data Eng.*, 26(10):2453–2465, 2014. doi:10.1109/TKDE.2013.158.
- 158 Yongsub Lim, Minsoo Jung, and U Kang. Memory-efficient and accurate sampling for counting local triangles in graph streams: From simple to multigraphs. *ACM Trans. Knowl. Discov. Data*, 12(1):4:1–4:28, 2018. doi:10.1145/3022186.
- 159 Yongsub Lim and U Kang. MASCOT: memory-efficient and accurate sampling for counting local triangles in graph streams. In Longbing Cao, Chengqi Zhang, Thorsten Joachims, Geoffrey I. Webb, Dragos D. Margineantu, and Graham Williams, editors, *Proc. of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Sydney, NSW, Australia, August 10-13, 2015*, pages 685–694. ACM, 2015. doi:10.1145/2783258.2783285.
- 160 Paul Liu, Austin R. Benson, and Moses Charikar. Sampling methods for counting temporal motifs. In J. Shane Culpepper, Alistair Moffat, Paul N. Bennett, and Kristina Lerman, editors, *Proc. of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM 2019, Melbourne, VIC, Australia, February 11-15, 2019*, pages 294–302. ACM, 2019. doi:10.1145/3289600.3290988.
- 161 Shangqi Lu and Yufei Tao. Towards optimal dynamic indexes for approximate (and exact) triangle counting. In Ke Yi and Zhewei Wei, editors, *24th International Conference on Database Theory, ICDT 2021, March 23-26, 2021, Nicosia, Cyprus*, volume 186 of *LIPICs*, pages 6:1–6:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.ICDT.2021.6.
- 162 Amgad Madkour, Walid G Aref, Faizan Ur Rehman, Mohamed Abdur Rahman, and Saleh Basalamah. A survey of shortest-path algorithms. *arXiv preprint*, 2017. arXiv:1705.02044.
- 163 Devavret Makkar, David A. Bader, and Oded Green. Exact and parallel triangle counting in dynamic graphs. In *24th IEEE International Conference on High Performance Computing, HiPC 2017, Jaipur, India, December 18-21, 2017*, pages 2–12. IEEE Computer Society, 2017. doi:10.1109/HiPC.2017.00011.
- 164 Alberto Marchetti-Spaccamela, Umberto Nanni, and Hans Rohnert. Maintaining a topological order under edge insertions. *Inf. Process. Lett.*, 59(1):53–58, 1996. doi:10.1016/0020-0190(96)00075-0.
- 165 Henning Meyerhenke. Dynamic load balancing for parallel numerical simulations based on repartitioning with disturbed diffusion. In *15th International Conference on Parallel and Distributed Systems*, pages 150–157. IEEE, 2009. doi:10.1109/ICPADS.2009.114.
- 166 Henning Meyerhenke and Joachim Gehweiler. On dynamic graph partitioning and graph clustering using diffusion. In Giuseppe F. Italiano, David S. Johnson, Petra Mutzel, and Peter Sanders, editors, *Algorithm Engineering, 27.06. - 02.07.2010*, volume 10261 of *Dagstuhl Seminar Proc.* Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany, 2010. URL: <http://drops.dagstuhl.de/opus/volltexte/2010/2798/>.

- 167 Kurt T Miller and Tina Eliassi-Rad. Continuous time group discovery in dynamic graphs. In *Notes of the 2009 NIPS Workshop on Analyzing Networks and Learning with Graphs*, Whistler, BC, Canada, 2009.
- 168 Peter Bro Miltersen, Sairam Subramanian, Jeffrey Scott Vitter, and Roberto Tamassia. Complexity models for incremental computation. *Theor. Comput. Sci.*, 130(1):203–236, 1994. doi:10.1016/0304-3975(94)90159-7.
- 169 Daniele Miorandi and Francesco De Pellegrini. K-shell decomposition for dynamic complex networks. In Lavy Libman Ariel Orda, Nidhi Hegde, editor, *8th International Symposium on Modeling and Optimization in Mobile, Ad-Hoc and Wireless Networks (WiOpt 2010)*, May 31 - June 4, 2010, University of Avignon, Avignon, France, pages 488–496. IEEE, 2010. URL: <http://ieeexplore.ieee.org/document/5520231/>.
- 170 Sudip Misra and B. John Oommen. Dynamic algorithms for the shortest path routing problem: Learning automata-based solutions. *IEEE Trans. Syst. Man Cybern. Part B*, 35(6):1179–1192, 2005. doi:10.1109/TSMCB.2005.850180.
- 171 Kingshuk Mukherjee, Md Mahmudul Hasan, Christina Boucher, and Tamer Kahveci. Counting motifs in dynamic networks. *BMC Syst. Biol.*, 12(1):1–12, 2018. doi:10.1186/s12918-018-0533-6.
- 172 Kengo Nakamura and Kunihiro Sadakane. Space-efficient fully dynamic DFS in undirected graphs. *Algorithms*, 12(3):52, 2019. doi:10.3390/a12030052.
- 173 Paolo Narváez, Kai-Yeung Siu, and Hong-Yi Tzeng. New dynamic algorithms for shortest path tree computation. *IEEE/ACM Trans. Netw.*, 8(6):734–746, 2000. doi:10.1109/90.893870.
- 174 Paolo Narváez, Kai-Yeung Siu, and Hong-Yi Tzeng. New dynamic SPT algorithm based on a ball-and-string model. *IEEE/ACM Trans. Netw.*, 9(6):706–718, 2001. doi:10.1109/90.974525.
- 175 Meghana Nasre, Matteo Pontecorvi, and Vijaya Ramachandran. Betweenness centrality - incremental and faster. In Erzsébet Csuhaj-Varjú, Martin Dietzfelbinger, and Zoltán Ésik, editors, *Mathematical Foundations of Computer Science 2014 - 39th International Symposium, MFCS 2014, Budapest, Hungary, August 25-29, 2014. Proceedings, Part II*, volume 8635 of *Lecture Notes in Computer Science*, pages 577–588. Springer, 2014. doi:10.1007/978-3-662-44465-8_49.
- 176 Eisha Nathan and David A. Bader. Approximating personalized katz centrality in dynamic graphs. In Roman Wyrzykowski, Jack J. Dongarra, Ewa Deelman, and Konrad Karczewski, editors, *Parallel Processing and Applied Mathematics - 12th International Conference, PPAM 2017, Lublin, Poland, September 10-13, 2017, Revised Selected Papers, Part I*, volume 10777 of *Lecture Notes in Computer Science*, pages 290–302. Springer, 2017. doi:10.1007/978-3-319-78024-5_26.
- 177 Eisha Nathan and David A. Bader. A dynamic algorithm for updating katz centrality in graphs. In Jana Diesner, Elena Ferrari, and Guandong Xu, editors, *Proc. of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017, Sydney, Australia, July 31 - August 03, 2017*, pages 149–154. ACM, 2017. doi:10.1145/3110025.3110034.
- 178 Ofer Neiman and Shay Solomon. Simple deterministic algorithms for fully dynamic maximal matching. *ACM Trans. Algorithms*, 12(1):7:1–7:15, 2016. doi:10.1145/2700206.
- 179 M. EJ Newman and M. Girvan. Finding and Evaluating Community Structure in Networks. *Physical review E*, 69(2):026113, 2004. doi:10.1103/PhysRevE.69.026113.
- 180 Daniel Nicoara, Shahin Kamali, Khuzaima Daudjee, and Lei Chen. Hermes: Dynamic partitioning for distributed social network graph databases. In *EDBT*, pages 25–36, 2015. doi:10.5441/002/edbt.2015.04.
- 181 Krzysztof Onak and Ronitt Rubinfeld. Maintaining a large matching and a small vertex cover. In *STOC*, pages 457–464, 2010. doi:10.1145/1806689.1806753.
- 182 James B. Orlin. Max flows in $o(nm)$ time, or better. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 765–774. ACM, 2013. doi:10.1145/2488608.2488705.

- 183 Mihai Patrascu. Towards polynomial lower bounds for dynamic problems. In *Proc. of the 42nd ACM Symposium on Theory of Computing, STOC*, pages 603–610. ACM, 2010. doi:10.1145/1806689.1806772.
- 184 Mihai Patrascu and Erik D. Demaine. Logarithmic lower bounds in the cell-probe model. *SIAM J. Comput.*, 35(4):932–963, 2006. doi:10.1137/S0097539705447256.
- 185 A. Pavan, Kanat Tangwongsan, Srikanta Tirthapura, and Kun-Lung Wu. Counting and sampling triangles from a graph stream. *Proc. VLDB Endow.*, 6(14):1870–1881, 2013. doi:10.14778/2556549.2556569.
- 186 D. Pearce and P. Kelly. A batch algorithm for maintaining a topological order. In *ACSC*, 2010. doi:10.5555/1862199.1862208.
- 187 David J. Pearce and Paul H. J. Kelly. A dynamic algorithm for topologically sorting directed acyclic graphs. In Celso C. Ribeiro and Simone L. Martins, editors, *Experimental and Efficient Algorithms, Third International Workshop, WEA 2004, Angra dos Reis, Brazil, May 25-28, 2004, Proc.*, volume 3059 of *Lecture Notes in Computer Science*, pages 383–398. Springer, 2004. doi:10.1007/978-3-540-24838-5_29.
- 188 David J. Pearce and Paul H. J. Kelly. A dynamic topological sort algorithm for directed acyclic graphs. *ACM J. Exp. Algorithmics*, 11, 2006. doi:10.1145/1187436.1210590.
- 189 David J. Pearce, Paul H. J. Kelly, and Chris Hankin. Online cycle detection and difference propagation: Applications to pointer analysis. *Softw. Qual. J.*, 12(4):311–337, 2004. doi:10.1023/B:SQJ0.0000039791.93071.a2.
- 190 Reynaldo Gil Pons. Space efficient incremental betweenness algorithm for directed graphs. In Rubén Vera-Rodríguez, Julian Fierrez, and Aythami Morales, editors, *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications - 23rd Iberoamerican Congress, CIARP 2018, Madrid, Spain, November 19-22, 2018, Proc.*, volume 11401 of *Lecture Notes in Computer Science*, pages 262–270. Springer, 2018. doi:10.1007/978-3-030-13469-3_31.
- 191 Matteo Pontecorvi and Vijaya Ramachandran. Fully dynamic betweenness centrality. In Khaled M. Elbassioni and Kazuhisa Makino, editors, *Algorithms and Computation - 26th International Symposium, ISAAC 2015, Nagoya, Japan, December 9-11, 2015, Proceedings*, volume 9472 of *Lecture Notes in Computer Science*, pages 331–342. Springer, 2015. doi:10.1007/978-3-662-48971-0_29.
- 192 K. (Lynn) Putman, Hanjo D. Boekhout, and Frank W. Takes. Fast incremental computation of harmonic closeness centrality in directed weighted networks. In Francesca Spezzano, Wei Chen, and Xiaokui Xiao, editors, *ASONAM '19: International Conference on Advances in Social Networks Analysis and Mining, Vancouver, British Columbia, Canada, 27-30 August, 2019*, pages 1018–1025. ACM, 2019. doi:10.1145/3341161.3344829.
- 193 G Ramalingam and Thomas Reps. On the computational complexity of incremental algorithms. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 1991.
- 194 G. Ramalingam and Thomas W. Reps. An incremental algorithm for a generalization of the shortest-path problem. *J. Algorithms*, 21(2):267–305, 1996. doi:10.1006/jagm.1996.0046.
- 195 G. Ramalingam and Thomas W. Reps. On the computational complexity of dynamic graph problems. *Theor. Comput. Sci.*, 158(1&2):233–277, 1996. doi:10.1016/0304-3975(95)00079-8.
- 196 Celso C. Ribeiro and Rodrigo F. Toso. Experimental analysis of algorithms for updating minimum spanning trees on graphs subject to changes on edge weights. In Camil Demetrescu, editor, *Experimental Algorithms, 6th International Workshop, WEA 2007, Rome, Italy, June 6-8, 2007, Proc.*, volume 4525 of *Lecture Notes in Computer Science*, pages 393–405. Springer, 2007. doi:10.1007/978-3-540-72845-0_30.
- 197 Matteo Riondato and Evgenios M. Kornaropoulos. Fast approximation of betweenness centrality through sampling. In Ben Carterette, Fernando Diaz, Carlos Castillo, and Donald Metzler, editors, *Seventh ACM International Conference on Web Search and Data Mining, WSDM 2014, New York, NY, USA, February 24-28, 2014*, pages 413–422. ACM, 2014. doi:10.1145/2556195.2556224.

- 198 Liam Roditty. A faster and simpler fully dynamic transitive closure. *ACM Trans. Algorithms*, 4(1):6:1–6:16, 2008. doi:10.1145/1328911.1328917.
- 199 Liam Roditty and Uri Zwick. On dynamic shortest paths problems. *Algorithmica*, 61(2):389–401, 2011. doi:10.1007/s00453-010-9401-5.
- 200 Liam Roditty and Uri Zwick. Dynamic approximate all-pairs shortest paths in undirected graphs. *SIAM J. Comput.*, 41(3):670–683, 2012. doi:10.1137/090776573.
- 201 Liam Roditty and Uri Zwick. A fully dynamic reachability algorithm for directed graphs with an almost linear update time. *SIAM Journal on Computing*, 45(3):712–733, 2016. doi:10.1137/13093618X.
- 202 Giulio Rossetti and Rémy Cazabet. Community discovery in dynamic networks: A survey. *ACM Comput. Surv.*, 51(2):35:1–35:37, 2018. doi:10.1145/3172867.
- 203 Tiberiu Rotaru and Hans-Heinrich Nägele. Dynamic load balancing by diffusion in heterogeneous systems. *Journal of Parallel and Distributed Computing*, 64(4):481–497, 2004. doi:10.1016/j.jpdc.2004.02.001.
- 204 Chayma Sakouhi, Sabeur Aridhi, Alessio Guerrieri, Salma Sassi, and Alberto Montresor. Dynamicdfep: A distributed edge partitioning approach for large dynamic graphs. In Evan Desai, Bipin C. Desai, Motomichi Toyama, and Jorge Bernardino, editors, *Proc. of the 20th International Database Engineering & Applications Symposium, IDEAS 2016, Montreal, QC, Canada, July 11-13, 2016*, pages 142–147. ACM, 2016. doi:10.1145/2938503.2938506.
- 205 Piotr Sankowski. Dynamic transitive closure via dynamic matrix inverse (extended abstract). In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proc.*, pages 509–517. IEEE Computer Society, 2004. doi:10.1109/FOCS.2004.25.
- 206 Piotr Sankowski. Subquadratic algorithm for dynamic shortest distances. In Lusheng Wang, editor, *Computing and Combinatorics, 11th Annual International Conference, COCOON 2005, Kunming, China, August 16-29, 2005, Proceedings*, volume 3595 of *Lecture Notes in Computer Science*, pages 461–470. Springer, 2005. doi:10.1007/11533719_47.
- 207 Piotr Sankowski. Faster dynamic matchings and vertex connectivity. In *SODA*, pages 118–126, 2007. doi:10.1145/1283383.1283397.
- 208 Ahmet Erdem Sariyüce, Bugra Gedik, Gabriela Jacques-Silva, Kun-Lung Wu, and Ümit V. Çatalyürek. Incremental k-core decomposition: algorithms and evaluation. *VLDB J.*, 25(3):425–447, 2016. doi:10.1007/s00778-016-0423-8.
- 209 Ahmet Erdem Sariyüce, Kamer Kaya, Erik Saule, and Ümit V. Çatalyürek. Incremental algorithms for closeness centrality. In Xiaohua Hu, Tsau Young Lin, Vijay V. Raghavan, Benjamin W. Wah, Ricardo Baeza-Yates, Geoffrey C. Fox, Cyrus Shahabi, Matthew Smith, Qiang Yang, Rayid Ghani, Wei Fan, Ronny Lempel, and Raghunath Nambiar, editors, *Proc. of the 2013 IEEE International Conference on Big Data, 6-9 October 2013, Santa Clara, CA, USA*, pages 487–492. IEEE Computer Society, 2013. doi:10.1109/BigData.2013.6691611.
- 210 Ahmet Erdem Sariyüce, Erik Saule, Kamer Kaya, and Ümit V. Çatalyürek. STREAMER: A distributed framework for incremental closeness centrality computation. In *2013 IEEE International Conference on Cluster Computing, CLUSTER 2013, Indianapolis, IN, USA, September 23-27, 2013*, pages 1–8. IEEE Computer Society, 2013. doi:10.1109/CLUSTER.2013.6702680.
- 211 Ahmet Erdem Sariyüce, Erik Saule, Kamer Kaya, and Ümit V. Çatalyürek. Incremental closeness centrality in distributed memory. *Parallel Comput.*, 47:3–18, 2015. doi:10.1016/j.parco.2015.01.003.
- 212 Benjamin Schiller, Sven Jager, Kay Hamacher, and Thorsten Strufe. Stream - A stream-based algorithm for counting motifs in dynamic graphs. In Adrian-Horia Dediu, Francisco Hernández Quiroz, Carlos Martín-Vide, and David A. Rosenblueth, editors, *Algorithms for Computational Biology - Second International Conference, AICoB 2015, Mexico City, Mexico, August 4-5, 2015, Proc.*, volume 9199 of *Lecture Notes in Computer Science*, pages 53–67. Springer, 2015. doi:10.1007/978-3-319-21233-3_5.

- 213 Kirk Schloegel, George Karypis, and Vipin Kumar. Multilevel diffusion schemes for re-partitioning of adaptive meshes. *J. Parallel Distributed Comput.*, 47(2):109–124, 1997. doi:10.1006/jpdc.1997.1410.
- 214 Kirk Schloegel, George Karypis, and Vipin Kumar. A unified algorithm for load-balancing adaptive scientific simulations. In Jed Donnelley, editor, *Proc. Supercomputing 2000, November 4-10, 2000, Dallas, Texas, USA. IEEE Computer Society, CD-ROM*, page 59. IEEE Computer Society, 2000. doi:10.1109/SC.2000.10035.
- 215 Kirk Schloegel, George Karypis, and Vipin Kumar. Parallel static and dynamic multi-constraint graph partitioning. *Concurr. Comput. Pract. Exp.*, 14(3):219–240, 2002. doi:10.1002/cpe.605.
- 216 Dominik Schultes and Peter Sanders. Dynamic highway-node routing. In Camil Demetrescu, editor, *Experimental Algorithms, 6th International Workshop, WEA 2007, Rome, Italy, June 6-8, 2007, Proc.*, volume 4525 of *Lecture Notes in Computer Science*, pages 66–79. Springer, 2007. doi:10.1007/978-3-540-72845-0_6.
- 217 Dipanjan Sengupta, Narayanan Sundaram, Xia Zhu, Theodore L Willke, Jeffrey Young, Matthew Wolf, and Karsten Schwan. Graphin: An online high performance incremental graph processing framework. In *European Conference on Parallel Processing*, pages 319–333. Springer, 2016. doi:10.1007/978-3-319-43659-3_24.
- 218 Neha Sengupta, Michael Hamann, and Dorothea Wagner. Benchmark generator for dynamic overlapping communities in networks. In Vijay Raghavan, Srinivas Aluru, George Karypis, Lucio Miele, and Xindong Wu, editors, *2017 IEEE International Conference on Data Mining, ICDM 2017, New Orleans, LA, USA, November 18-21, 2017*, pages 415–424. IEEE Computer Society, 2017. doi:10.1109/ICDM.2017.51.
- 219 Zhenzhen Shao, Na Guo, Yu Gu, Zhigang Wang, Fangfang Li, and Ge Yu. Efficient closeness centrality computation for dynamic graphs. In Yunmook Nah, Bin Cui, Sang-Won Lee, Jeffrey Xu Yu, Yang-Sae Moon, and Steven Euijong Whang, editors, *Database Systems for Advanced Applications - 25th International Conference, DASFAA 2020, Jeju, South Korea, September 24-27, 2020, Proc., Part II*, volume 12113 of *Lecture Notes in Computer Science*, pages 534–550. Springer, 2020. doi:10.1007/978-3-030-59416-9_32.
- 220 Y. Shiloach and S. Even. An on-line edge-deletion problem. *Journal of the ACM*, 28(1):1–4, 1981. doi:10.1145/322234.322235.
- 221 Dharendra Singh and Nilay Khare. Parallel batch dynamic single source shortest path algorithm and its implementation on GPU based machine. *Int. Arab J. Inf. Technol.*, 16(2):217–225, 2019. URL: http://iajit.org/index.php?option=com_content&task=blogcategory&id=137&Itemid=469.
- 222 Daniel Dominic Sleator and Robert Endre Tarjan. A data structure for dynamic trees. In *Proc. of the 13th Annual ACM Symposium on Theory of Computing, May 11-13, 1981, Milwaukee, Wisconsin, USA*, pages 114–122. ACM, 1981. doi:10.1145/800076.802464.
- 223 Shay Solomon. Fully dynamic maximal matching in constant update time. In *57th Symposium on Foundations of Computer Science FOCS*, pages 325–334, 2016. doi:10.1109/FOCS.2016.43.
- 224 Lorenzo De Stefani, Alessandro Epasto, Matteo Riondato, and Eli Upfal. Triest: Counting local and global triangles in fully dynamic streams with fixed memory size. *ACM Trans. Knowl. Discov. Data*, 11(4):43:1–43:50, 2017. doi:10.1145/3059194.
- 225 Daniel Stubbs and Virginia Vassilevska Williams. Metatheorems for dynamic weighted matching. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, volume 67 of *LIPICs*, pages 58:1–58:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICs.ITCS.2017.58.
- 226 Binta Sun, T.-H. Hubert Chan, and Mauro Sozio. Fully dynamic approximate k-core decomposition in hypergraphs. *ACM Trans. Knowl. Discov. Data*, 14(4), May 2020. doi:10.1145/3385416.

- 227 Robert Endre Tarjan and Renato Fonseca F. Werneck. Dynamic trees in practice. *ACM J. Exp. Algorithmics*, 14, 2009. doi:10.1145/1498698.1594231.
- 228 Mikkel Thorup. Fully-dynamic all-pairs shortest paths: Faster and allowing negative cycles. In Torben Hagerup and Jyrki Katajainen, editors, *Algorithm Theory - SWAT 2004, 9th Scandinavian Workshop on Algorithm Theory, Humlebaek, Denmark, July 8-10, 2004, Proceedings*, volume 3111 of *Lecture Notes in Computer Science*, pages 384–396. Springer, 2004. doi:10.1007/978-3-540-27810-8_33.
- 229 Mikkel Thorup. Worst-case update times for fully-dynamic all-pairs shortest paths. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 112–119. ACM, 2005. doi:10.1145/1060590.1060607.
- 230 Mikkel Thorup. Fully-dynamic min-cut. *Comb.*, 27(1):91–127, 2007. doi:10.1007/s00493-007-0045-2.
- 231 Mikkel Thorup and David R. Karger. Dynamic graph algorithms with applications. In Magnús M. Halldórsson, editor, *Algorithm Theory - SWAT 2000, 7th Scandinavian Workshop on Algorithm Theory, Bergen, Norway, July 5-7, 2000, Proceedings*, volume 1851 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 2000. doi:10.1007/3-540-44985-X_1.
- 232 Thomas Tseng, Laxman Dhulipala, and Guy E. Blelloch. Batch-parallel euler tour trees. In Stephen G. Kobourov and Henning Meyerhenke, editors, *Proceedings of the Twenty-First Workshop on Algorithm Engineering and Experiments, ALENEX 2019, San Diego, CA, USA, January 7-8, 2019*, pages 92–106. SIAM, 2019. doi:10.1137/1.9781611975499.8.
- 233 Jan van den Brand and Danupon Nanongkai. Dynamic approximate shortest paths and beyond: Subquadratic and worst-case update time. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 436–455. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00035.
- 234 Jan van den Brand, Danupon Nanongkai, and Thatchaphol Saranurak. Dynamic matrix inverse: Improved algorithms and matching conditional lower bounds. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 456–480. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00036.
- 235 Alexander van der Grinten, Elisabetta Bergamini, Oded Green, David A. Bader, and Henning Meyerhenke. Scalable katz ranking computation in large static and dynamic graphs. In Yossi Azar, Hannah Bast, and Grzegorz Herman, editors, *26th Annual European Symposium on Algorithms, ESA 2018, August 20-22, 2018, Helsinki, Finland*, volume 112 of *LIPIcs*, pages 42:1–42:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ESA.2018.42.
- 236 Luis M. Vaquero, Félix Cuadrado, Dionysios Logothetis, and Claudio Martella. xdgp: A dynamic graph processing system with adaptive partitioning. *CoRR*, abs/1309.1049, 2013. arXiv:1309.1049.
- 237 Luis M. Vaquero, Félix Cuadrado, Dionysios Logothetis, and Claudio Martella. Adaptive partitioning for large-scale dynamic graphs. In *IEEE 34th International Conference on Distributed Computing Systems, ICDCS 2014, Madrid, Spain, June 30 - July 3, 2014*, pages 144–153. IEEE Computer Society, 2014. doi:10.1109/ICDCS.2014.23.
- 238 Tanmay Verma and Dhruv Batra. Maxflow revisited: An empirical comparison of maxflow algorithms for dense vision problems. In Richard Bowden, John P. Collomosse, and Krystian Mikolajczyk, editors, *British Machine Vision Conference, BMVC 2012, Surrey, UK, September 3-7, 2012*, pages 1–12. BMVA Press, 2012. doi:10.5244/C.26.61.
- 239 Dorothea Wagner, Thomas Willhalm, and Christos D. Zaroliagis. Geometric containers for efficient shortest-path computation. *ACM J. Exp. Algorithmics*, 10, 2005. doi:10.1145/1064546.1103378.

- 240 Chris Walshaw, Mark Cross, and Martin G. Everett. Parallel dynamic graph partitioning for adaptive unstructured meshes. *J. Parallel Distributed Comput.*, 47(2):102–108, 1997. doi:10.1006/jpdc.1997.1407.
- 241 Jingjing Wang, Yanhao Wang, Wenjun Jiang, Yuchen Li, and Kian-Lee Tan. Efficient sampling algorithms for approximate temporal motif counting. In Mathieu d’Aquin, Stefan Dietze, Claudia Hauff, Edward Curry, and Philippe Cudré-Mauroux, editors, *CIKM ’20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19–23, 2020*, pages 1505–1514. ACM, 2020. doi:10.1145/3340531.3411862.
- 242 Na Wang, Dongxiao Yu, Hai Jin, Chen Qian, Xia Xie, and Qiang-Sheng Hua. Parallel algorithm for core maintenance in dynamic graphs. In Kisung Lee and Ling Liu, editors, *37th IEEE International Conference on Distributed Computing Systems, ICDCS 2017, Atlanta, GA, USA, June 5–8, 2017*, pages 2366–2371. IEEE Computer Society, 2017. doi:10.1109/ICDCS.2017.288.
- 243 Stefan Weigert, Matti Hiltunen, and Christof Fetzer. Mining large distributed log data in near real time. In *Managing Large-scale Systems via the Analysis of System Logs and the Application of Machine Learning Techniques*, pages 1–8. Association for Computing Machinery, 2011. doi:10.1145/2038633.2038638.
- 244 Martin Winter, Daniel Mlakar, Rhaleb Zayer, Hans-Peter Seidel, and Markus Steinberger. faimgraph: high performance management of fully-dynamic graphs under tight memory constraints on the GPU. In *Proc. of the International Conference for High Performance Computing, Networking, Storage, and Analysis, SC 2018, Dallas, TX, USA, November 11–16, 2018*, pages 60:1–60:13. IEEE / ACM, 2018. URL: <http://dl.acm.org/citation.cfm?id=3291736>.
- 245 Ning Xu, Lei Chen, and Bin Cui. Loggp: a log-based dynamic graph partitioning method. *Proc. of the VLDB Endowment*, 7(14):1917–1928, 2014. doi:10.14778/2733085.2733097.
- 246 Bohua Yang, Dong Wen, Lu Qin, Ying Zhang, Xubo Wang, and Xuemin Lin. Fully dynamic depth-first search in directed graphs. *Proc. VLDB Endow.*, 13(2):142–154, 2019. doi:10.14778/3364324.3364329.
- 247 Chia-Chen Yen, Mi-Yen Yeh, and Ming-Syan Chen. An efficient approach to updating closeness centrality and average path length in dynamic networks. In Hui Xiong, George Karypis, Bhavani M. Thuraisingham, Diane J. Cook, and Xindong Wu, editors, *2013 IEEE 13th International Conference on Data Mining, Dallas, TX, USA, December 7–10, 2013*, pages 867–876. IEEE Computer Society, 2013. doi:10.1109/ICDM.2013.135.
- 248 Anita Zakrzewska and David A. Bader. Fast incremental community detection on dynamic graphs. In Roman Wyrzykowski, Ewa Deelman, Jack Dongarra, Konrad Karczewski, Jacek Kitowski, and Kazimierz Wiatr, editors, *Parallel Processing and Applied Mathematics*, pages 207–217, Cham, 2016. Springer International Publishing. doi:10.1007/978-3-319-32149-3_20.
- 249 Christos D Zaroliagis. Implementations and experimental studies of dynamic graph algorithms. In *Experimental algorithmics*, pages 229–278. Springer, 2002. doi:10.1007/3-540-36383-1_11.
- 250 Yikai Zhang, Jeffrey Xu Yu, Ying Zhang, and Lu Qin. A fast order-based approach for core maintenance. In *33rd IEEE International Conference on Data Engineering, ICDE 2017, San Diego, CA, USA, April 19–22, 2017*, pages 337–348. IEEE Computer Society, 2017. doi:10.1109/ICDE.2017.93.
- 251 Weiguo Zheng, Chengzhi Piao, Hong Cheng, and Jeffrey Xu Yu. Computing a near-maximum independent set in dynamic graphs. In *35th IEEE International Conference on Data Engineering, ICDE 2019, Macao, China, April 8–11, 2019*, pages 76–87. IEEE, 2019. doi:10.1109/ICDE.2019.00016.
- 252 Weiguo Zheng, Qichen Wang, Jeffrey Xu Yu, Hong Cheng, and Lei Zou. Efficient computation of a near-maximum independent set over evolving graphs. In *34th IEEE International Conference on Data Engineering, ICDE 2018, Paris, France, April 16–19, 2018*, pages 869–880. IEEE Computer Society, 2018. doi:10.1109/ICDE.2018.00083.

- 253 Lei Zhu, Shaoning Pang, Abdolhossein Sarrafzadeh, Tao Ban, and Daisuke Inoue. Incremental and decremental max-flow for online semi-supervised learning. *IEEE Trans. Knowl. Data Eng.*, 28(8):2115–2127, 2016. doi:10.1109/TKDE.2016.2550042.
- 254 Di Zhuang, Morris J Chang, and Mingchen Li. Dynamo: Dynamic community detection by incrementally maximizing modularity. *IEEE Transactions on Knowledge and Data Engineering*, 2019. doi:10.1109/TKDE.2019.2951419.