


Dynamic Matching Algorithms in Practice

Monika Henzinger 


University of Vienna, Faculty of Computer Science, Austria
monika.henzinger@univie.ac.at

Shahbaz Khan 

Department of Computer Science, University of Helsinki, Finland
shahbaz.khan@helsinki.fi

Richard Paul 

University of Vienna, Faculty of Computer Science, Austria
richard.paul@univie.ac.at

Christian Schulz 

University of Vienna, Faculty of Computer Science, Austria
christian.schulz@univie.ac.at

Abstract

In recent years, significant advances have been made in the design and analysis of fully dynamic maximal matching 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. In this paper, we attempt to bridge the gap between theory and practice that is currently observed for the *fully dynamic maximal matching problem*. We engineer several algorithms and empirically study those algorithms on an extensive set of dynamic instances.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms

Keywords and phrases Matching, Dynamic Matching, Blossom Algorithm

Digital Object Identifier 10.4230/LIPIcs.ESA.2020.58

Supplementary Material Source code and instances are available at <https://github.com/schulzchristian/DynMatch>.

Funding The research leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) /ERC grant agreement No. 340506.

1 Introduction

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. These problems can be solved in polynomial time. For example, Micali and Vazirani [31] compute a maximum cardinality matching in $O(m\sqrt{n})$ time. For the weighted case, the fastest algorithm is by Galil et. al [19] requiring $O(mn \log n)$ time which improves the $O(n^3)$ time algorithm [18] for sparse graphs.

However, often the underlying graphs change over time, e.g., edges are inserted or deleted in the graph as the time progresses. For example, new relations between objects of a network may be created or removed over time (for example [30]). Even though the matching problem can be solved in polynomial time, computing a new matching from scratch every time the graph changes is an expensive task on huge networks, as this ignores the previously computed



© Monika Henzinger, Shahbaz Khan, Richard Paul, and Christian Schulz;
licensed under Creative Commons License CC-BY
28th Annual European Symposium on Algorithms (ESA 2020).

Editors: Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders; Article No. 58; pp. 58:1–58:20
Leibniz International Proceedings in Informatics



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

information on the given instance. Hence, in the recent years significant advances have been made in the design and analysis of fully dynamic maximal matching algorithms. These theoretical algorithmic ideas have received very little attention from the practical perspective. Only a few of the dynamic algorithms are implemented and tested on real datasets, and hence their practical potential is far from being understood.

Contribution and Outline. In this paper, we start to bridge the gap between theory and practice that is currently observed for the *fully dynamic maximal matching problem*. We engineer several dynamic maximal matching algorithms as well as an algorithm that is able to maintain the maximum matching. To this end, we look at an algorithm due to Baswana, Gupta and Sen [4], which performs edge updates in $O(\sqrt{n})$ time and maintains a 2-approximate maximum matching, the algorithm of Neiman and Solomon [33], which takes $O(\sqrt{n+m})$ time to maintain a $3/2$ -approximate maximum matching, as well as two *novel* dynamic algorithms: a random walk-based algorithm as well as a dynamic algorithm that searches for augmenting paths using a (depth bounded) blossom algorithm. Without depth bound, the latter algorithm is able to maintain a maximum matching. We perform extensive experiments comparing the performance of these algorithms on the real-world and artificially generated instances. Experiments indicate that maintaining optimum matchings can be done much more efficiently than the naive algorithm that recomputes maximum matchings from scratch (more than an order of magnitude faster). Second, all non-optimum dynamic algorithms that we consider in this work are able to maintain near-optimum matchings in practice while being multiple orders of magnitudes faster than the naive optimum dynamic algorithm.

2 Preliminaries

2.1 Basic Concepts

Let $G = (V = \{0, \dots, n-1\}, E)$ be an *undirected graph* without parallel edges and self-loops. We set $n = |V|$, and $m = |E|$, $N(v) := \{u : \{v, u\} \in E\}$ denotes the *neighbors* of v . The degree of a vertex v is $d(v) := |N(v)|$. A matching $\mathcal{M} \subseteq E$ in a graph is a set of edges without common vertices. The *cardinality* or *size* of a matching is simply the cardinality of the edge subset \mathcal{M} . We call a matching *maximal*, if there is no edge in E that can be added to \mathcal{M} . A *maximum cardinality matching* \mathcal{M}_{opt} is a matching that contains the largest possible number of edges of all matchings. An α -*approximate* maximum matching is a matching, that contains at least $\frac{|\mathcal{M}_{\text{opt}}|}{\alpha}$ edges. A vertex is called *free* or *unmatched* if it is not incident to an edge of the matching. Otherwise, we call it *unfree* or *matched*. For a matched vertex u with $\{u, v\} \in \mathcal{M}$, we call vertex v the *mate* of u , which we denote as $\text{mate}(u) = v$. For an unmatched vertex u , we define $\text{mate}(u) = \perp$. An *augmenting path* is defined as a cycle-free path in the graph G , that starts and ends on a *free* vertex and where edges from \mathcal{M} alternate with edges from $E \setminus \mathcal{M}$. The *trivial augmenting path* is a single edge, that has both its endpoints unmatched. Throughout this paper, we call such an edge a *free* edge. If we take an augmenting path and resolve it by matching every unmatched edge and unmatching every matched edge, we increase the cardinality of the matching by one. Any matching without *augmenting paths* is a maximum matching [5] and any matching with no augmenting paths of length at most $2k-3$ is a $(k/(k-1))$ -approximate maximum matching [23]. Hence, a maximal matching having no augmenting paths of length one (or free edges) is a 2-approximate maximum matching. Throughout the paper, we omit the inverse Ackermann function from complexity statements.

Our focus in this paper are *fully dynamic graphs*, where the number of vertices is fixed, but edges can be added and removed. All the algorithms evaluated can handle edge insertions as well as edge deletions. In the following, Δ denotes the maximum degree that can be found in any state of the dynamic graph.

2.2 Related Work

Computing large or maximum matchings in graphs is a well researched topic. Edmonds [17] gave an algorithm that can compute a maximum cardinality matching in a static graph in time $O(mn^2)$. This result was later improved to $O(mn^{0.5})$ by Micali and Vazirani [31]. Recently, algorithms use simple data reductions rules such as [25] to speed up computations, or shrink-trees instead of blossoms [16] to speed up computations in static graphs. In practice, these algorithms can still be time consuming for many applications involving large graphs. Hence, several near linear time approximation algorithms exist in practice such as the local max algorithm [11], the path growing algorithm [15] and the global paths algorithm [28]. As the focus of this work are dynamic graphs, we refer the reader to the quite extensive related work section of [16] for more recent static matching algorithms.

In the dynamic setting, the maximum matching problem has been prominently studied ensuring α -approximate guarantees. A major exception is the randomized algorithm by Sankowski [36] which maintains a maximum matching in $O(n^{1.495})$ update time. One can trivially maintain a maximal (2-approximate) matching in $O(n)$ update time by resolving all trivial augmenting paths of length one. Ivković and Llyod [24] designed the first fully dynamic algorithm to improve this bound to $O((n+m)^{\sqrt{2}/2})$ update time. Later, Onak and Rubinfeld [34] presented a randomized algorithm for maintaining a $O(1)$ -approximate matching in a dynamic graph that takes $O(\log^2 n)$ expected amortized time for each edge update. This result led to a flurry of results in this area. Baswana, Gupta and Sen [4] improved the approximation ratio of [34] from $O(1)$ to 2 and the amortized update time to $O(\log n)$. Further, Solomon [38] improved the update time of [4] from amortized $O(\log n)$ to *constant*. However, the first deterministic data structure improving [24] was given by Bhattacharya et al. [8] maintaining $(3 + \epsilon)$ approximate matching in $\tilde{O}(\min(\sqrt{n}, m^{1/3}/\epsilon^2))$ amortized update time, which was further improved to $(2 + \epsilon)$ requiring $O(\log n/\epsilon^2)$ update time by Bhattacharya et al. [9]. Recently, Bhattacharya et al. [7] achieved the first $O(1)$ amortized update time for a deterministic algorithm but for a weaker approximation guarantee of $O(1)$. For worst-case bounds, the best results are by Gupta and Peng [21] requiring $O(\sqrt{m}/\epsilon)$ update time for $(1 + \epsilon)$ approximation, Neiman and Solomon [33] requiring $O(\sqrt{m})$ update time for $3/2$ approximation, Bernstein and Stein [6] requiring $m^{1/4}/\epsilon^{2.5}$ for $(3/2 + \epsilon)$ approximation. Recently, Charikar and Solomon [12], and Arar et al. [2] (using [10]), independently presented the first algorithms requiring $O(\text{poly log } n)$ worst-case update time both maintaining $(2 + \epsilon)$ approximation. Recently, Grandoni et al. [20] gave an incremental matching algorithm that achieves a $(1 + \epsilon)$ -approximate matching in constant deterministic amortized time. Despite this variety of different algorithms, to the best of our knowledge, there has been no effort made so far, to engineer and evaluate these algorithms on real-world instances.

3 Algorithms

We now present the fully dynamic algorithms for the maximal matching problem under consideration. We implemented and tested a variety of simple, combinatorial algorithms that seemed likely to work well in practice. We begin with random walk based dynamic algorithms, followed by dynamic algorithms based on (bounded) augmenting path search and finally

review the algorithms by Baswana, Gupta and Sen [4] and Neiman and Solomon [33]. All of the algorithms are fully dynamic. Throughout this section, we provide a brief description of these algorithms and their implementation. In each case, we explain how we handle initialization, edge insertions and edge deletions separately.

3.1 Random Walk-based Algorithms

In general finding long augmenting paths in networks is an expensive step. The main idea of the random walk based methods proposed in this section is to use random walks in order to detect augmenting paths, and hence to improve the size of the matching. We start by explaining the general idea to use random walks for finding augmenting paths and then explain how we handle edge insertions and deletions.

3.1.1 Random Walks For Augmenting Paths

The algorithm works as follows: we start at a free vertex u and randomly choose a neighbour w of u . If this neighbour is free, then we match the edge (u, w) and our random walk stops. If w is matched, we unmatch $(w, \text{mate}(w))$ and match (u, w) . Note that $u \neq \text{mate}(w)$ since u is free in the beginning and therefore $\text{mate}(u) = \perp$, but $\text{mate}(\text{mate}(w)) = w$ and $w \neq \perp$. Afterwards, the previous mate of w is free. Hence, we continue our random walk at this vertex. Our random walk performs $O(\frac{1}{\epsilon})$ steps (see below). Here, ϵ basically controls the quality of the matching (see below). Since picking a random neighbor can be done in constant time, the overall time for the random walk update algorithm is $O(\frac{1}{\epsilon})$. Note that the length of the random walk is a natural parameter of the algorithm that we will investigate in the experimental evaluation.

Also note that if the algorithm does not end by matching two free vertices, the matching may not be maximal even if it was initially – this can be the case if the vertex freed last is incident to a free vertex. There are multiple possibilities to fix this. Our default is to undo all changes that have been done in this case. The overall running time of a random walk is then $O(1/\epsilon)$. Another possibility is Δ -*settling*: The algorithm tries to settle visited vertices. The algorithm scans through their neighbors to find a free vertex and stops if once it was successful or the number of steps exceeds $1/\epsilon$. If the random walk was not successful, the algorithm tries to match the last vertex touched by the random walk by scanning its neighbors instead of undoing all changes. This also ensures that the matching is maximal but requires $O(\Delta)$ additional time per visited vertex. The running time of the Δ -settling random walk is then $O(\Delta/\epsilon)$.

We now explain how we perform edge insertions and deletions.

Edge Insertion. Our algorithm handles edge insertions as follows: when inserting an edge (u, v) , if both the endpoints are free, we match it. Note that the simple algorithm stops here if at least one of the endpoints is not free. The random walk based algorithms try to improve insertion by doing the following: If both endpoints are matched, thus prohibiting to match the inserted edge, we do nothing. If only one of the endpoints is matched, w.l.o.g let this be u , we unmatch u and $w := \text{mate}(u)$ and match (u, v) . We then start a random walk as described above to find augmenting paths from w . If the random walk is unsuccessful to further increase the size of the matching, we undo all changes and restore the matching to the state before we unmatched u and w .

Edge Deletion. Deleting a matched edge (u, v) leaves the two endpoints u and v free. If possible, our algorithm matches them in $O(\Delta)$ time by scanning their neighbors in order to maintain a maximal matching. If u and v cannot be matched and the matching before edge deletion was maximal, then the matching remains maximal. However, a free vertex may be a starting point for an augmenting path of arbitrary length. Hence, we start a random walk as described above from u if it is free and do the same for v .

3.1.2 Analysis

The algorithm can maintain a $(1 + \epsilon)$ -approximation, if the random walks are of appropriate length and repeated sufficiently often. More precisely, if the algorithm uses random paths of length $2/\epsilon - 1$ and the process is repeated until successful or $\Delta^{2/\epsilon-1} \log n$ times, then with high probability the matching is a $(1 + \epsilon)$ approximation of the maximum matching (at each point in time).

► **Lemma 1.** *The random walk based algorithm maintains a $(1 + \epsilon)$ -approximate maximum matching if the length of the walk is $2/\epsilon - 1$ and the walks are repeated $\Delta^{2/\epsilon-1} \log n$ times.*

Proof. If no augmenting path of length $\leq 2/\epsilon - 1$ exists, then the matching is a $(\frac{1/\epsilon+1}{1/\epsilon}) = (1 + \epsilon)$ -approximate maximum matching. To see this, rewrite the length of the path to $2(1/\epsilon + 1) - 3$ and set $k = 1/\epsilon + 1$ in the approximation lemma above. If there is such a path from a free node, then the probability of finding it is $\geq (\frac{1}{\Delta})^{2/\epsilon-1}$ since one possibility is the that random walker makes the “correct” decision at every vertex of the path. The probability that λ random walks of length $2/\epsilon - 1$ do not find an augmenting path of length $2/\epsilon - 1$ is $\leq (1 - \frac{1}{\Delta^{2/\epsilon-1}})^\lambda \leq e^{-\frac{\lambda}{\Delta^{2/\epsilon-1}}}$. Thus for $\lambda \geq \Delta^{2/\epsilon-1} \log n$ the probability is $\leq 1/n$. ◀

Parallelization. Note that multiple repetitions of the random walks can be easily parallelized as they are completely independent if changes are made thread-local. If one random walker finds an augmenting path, it is accepted and the other random walkers can be stopped.

3.2 Blossom-based (Optimum) Algorithms

Note that the random walk algorithm also yields a static $(1 + \epsilon)$ -approximate maximum matching algorithm: use a simple greedy algorithm as initialization and then run the random walks as stated above from the remaining free nodes. However, the amount of repetitions to achieve the approximation is fairly high. Simply, running a modified BFS to find augmenting paths bounded in depth by $2/\epsilon - 1$ from a free node has a theoretically faster running time $O(\Delta^{2/\epsilon-1})$ per free node. Note however that the theoretical bound for the dynamic random walk algorithm is fairly pessimistic: our algorithm stops as soon as *one* augmenting path has been found – this path can also be shorter or in practice there may be multiple possibilities for augmenting paths so that the probability of finding it increases. So the natural question arises, whether a bounded augmenting path search is superior over random walk based methods stated above. Hence, we propose the following dynamic algorithms for the dynamic matching problem.

In most implementations (such as Boost [37]) finding an augmenting path starting from a free node takes $\Omega(n + m)$ running time due to initialization of the data structures of the modified BFS. These data structures are initialized every time an augmenting path search is started. Hence, the observed performance of Edmonds blossom algorithm to find an optimum matching in libraries such as Boost is $\Theta(n(n + m))$ if no algorithm to initialize the matching is used and $O(F(n + m))$ if some greedy algorithm is used as initialization and

F is the number of remaining free nodes after greedy initialization. The later is the reason why in practice greedy initialization strategies generally help to find optimum matchings. However, finding an augmenting path can easily be implemented such that it a) stops as soon as an augmenting path is found, and b) has running time $\Theta(n' + m')$, where n' and m' refers to the number of nodes and edges touched by the augmenting path search modified BFS [39, 29]. The first augmenting path search needs time $O(n + m)$ to initialize the typical data structures. All searches then do book keeping of the changes they made in the data structures and undo them afterwards. Note that this clearly changes the behaviour of the algorithm in practice: if there are many short augmenting paths the algorithm will run much faster than $\Theta(n(n + m))$. The implementation does not change the worst-case complexity, but improves the best case to $O(m)$ [29]. In fact, in our experience the static version of our implementation scales close to linear in m in practice (as there are many short augmenting paths in real world instances). In the following, we always use this variant of augmenting path search and each of the dynamic operations does book keeping to be able to quickly search for augmenting paths.

Edge Insertion. Let (u, v) be the inserted edge. If u and v are free, then we match that edge directly. Otherwise, we start an augmenting path search from u if u is free and from v if v is free. If both u and v are not free, then we perform a breadth first search from u to find a free node reachable via an alternating path. From this node we start an augmenting path search. Note that an augmenting path must use (u, v) as both connected components did not contain an augmenting path with the component before as the algorithm maintains a maximum matching. Also note that the last case will be an expensive step in practice as the algorithm tries to maintain a maximum matching, newly inserted edges will often not result in a new augmenting path and hence the augmenting path search takes $\Theta(n + m)$ time. Without the third case of the algorithm, we call it *unsafe*. That is in case both u and v are not free, the unsafe configuration of the algorithm does nothing.

Not using the unsafe option, the algorithm maintains a maximum matching. This is due to the fact that if the graph did not contain an augmenting path before insertion, the only way we can create one is due to the insertion of the new edge. The first and second case are obvious. In the third case, after finding a single free node, the augmenting path search must use the newly inserted edge (u, v) (which is not matched, but both endpoints are non-free). Hence, it is sufficient to find a single free node. After running the augmenting path search, the matching size has either increased by one, or there was no augmenting path. Hence, the matching must be maximum. Lastly, note that the third case is only necessary if both endpoints of the inserted edge are in different connected components.

Note that when considering insertions only, the algorithm is more expensive than just running the static algorithm. This is due to the fact that the static algorithm runs an augmenting path search from each free *node* once, while our dynamic algorithm does try to find augmenting paths every time we insert an edge (since the graph may have changed at other places not close to the inserted edge). The overall worst case complexity in this case is $O(m(n + m))$ compared to $O(n(n + m))$ for the static algorithm. In our experiments, this effect is especially noticeable if we start a search from a node where a previous augmenting path search has been unsuccessful.

Hence, besides using the unsafe option which drops the property that the matching is maximum, we propose the following optimization called *lazy augmenting path search*. Here, we start an augmenting path search from u and v only if at least $m'/2$ edges have been inserted or deleted since the last augmenting path search from u or v or no augmenting path

search has been started. Note that this effectively amortizes the cost for the augmenting path search, yielding amortized constant time per edge. Our experiments indicate that this speeds up the overall time of the algorithm drastically, while being only slightly worse than the optimum algorithm. Our third optimization limits the search depth of the augmenting path search to $2/\epsilon - 1$. This ensures that there is no augmenting path of length $2/\epsilon - 1$ and hence is a deterministic $(1 + \epsilon)$ -approximate matching algorithm (if the deletion part algorithm ensures this as well, and the algorithm is run with the safe option). Note that the worst case complexity of the optimum version of the insertion operation is $O(n + m)$, but in practice augmenting paths (if present) are much shorter. The bounded version of our algorithm has, however, worse case complexity of $O(\Delta^{2/\epsilon-1})$.

Edge Deletion. Let (u, v) be the deleted edge. After the deletion we start an augmenting path search from any free endpoint u or v . Depending on the configuration of the algorithm this either does a full run for an augmenting path or stops when the augmenting path search reached depth $2/\epsilon - 1$. In the first case, this guarantees that the matching is maximum if it was maximum before and in the latter case, our algorithm maintains an $(1 + \epsilon)$ approximate maximum matching. If case we use lazy augmenting path search, we start an (depth bounded) augmenting path search from u and v only if at least $m'/2$ edges have been inserted or deleted since the last augmenting path search from u or v . Otherwise, we limit augmenting path search from u and v to augmenting paths of length $\min(3, 2/\epsilon - 1)$.

3.3 Baswana, Gupta and Sen Algorithm

Baswana, Gupta and Sen (BGS) presented a randomized algorithm in [4], that maintains a maximal matching in a dynamic graph in *amortized* $O(\sqrt{n})$ update time with high probability. They also present a multi-level variant that runs in $O(\log n)$ amortized time. To be self contained, we briefly review the main concepts of the algorithm and follow their description closely.

Levels and Ownership of Edges. The algorithm uses the concept of *ownership* for edges. More precisely, based upon the number of edges that a vertex *owns*, the algorithm partitions the set of vertices into two levels 0 and 1. An edge is always owned by at least one of its endpoints. If both endpoints are at level 0, then both vertices own the edge. If only one endpoint is at level 1, then this endpoint owns the edge. If both endpoints are at level 1, then exactly one endpoint, namely the first mentioned vertex owns the edge. If a new edge (u, v) with $level(u) = level(v) = 1$ is inserted, it will therefore be owned by the vertex u .

■ **Algorithm 1** RANDOM-SETTLE(u): find a random edge (u, v) from the set of owned edges of u , matches it and returns the previous mate of v .

Let (u, v) be a uniformly randomly selected edge from \mathcal{O}_u

forall $(v, w) \in \mathcal{O}_u$ **do**

 └ remove (v, w) from \mathcal{O}_w

if v is matched **then**

 └ $x \leftarrow \text{mate}(v)$; $M \leftarrow M \setminus \{(v, x)\}$

else

 └ $x \leftarrow \text{NULL}$

$M \leftarrow M \cup \{(u, v)\}$, LEVEL(u) \leftarrow 1, LEVEL(v) \leftarrow 1

return x

The set \mathcal{O}_u denotes the set of edges owned by a vertex u . The level of an edge is defined by $level(e = \{u, v\}) = \max(level(u), level(v))$. BGS maintains the following invariants: (1) Every vertex on level 1 is matched. (2) Every free vertex on level 0 has all neighbours matched. (3) Every vertex on level 0 owns less than \sqrt{n} edges (at any moment of time). (4) Both endpoints of each matched edge are on same level.

Edge Insertion. Let (u, v) be the edge being inserted. If either u or v are at level 1, then there is no violation of any invariant. The algorithm adds (u, v) to \mathcal{O}_u if $level(u) = 1$ and to \mathcal{O}_v otherwise. If both endpoints of (u, v) are at level 0, then the algorithm proceeds as follows: If both endpoints are free, the edge is added to the matching. Adding the edge (u, v) to the sets \mathcal{O}_u and \mathcal{O}_v increases the number of edges owned by u and v . If at least one set \mathcal{O}_u or \mathcal{O}_v exceeds the threshold of \sqrt{n} in size, the vertex with the higher number of owned edges will be *repaired*. Let u be that vertex. Repairing a vertex u is done by calling the procedure RANDOM-SETTLE on u . As a result, u moves to level 1 and gets matched to some vertex y selected randomly uniformly from the set of owned edges \mathcal{O}_u . The vertex y is also moved to level 1 to satisfy invariant 4. If w and x were the earlier mates of u and y at level 0, respectively, then matching u and y has rendered w and x free. The algorithm tries to settle each of those by scanning their set of owned edges for free vertices.

Edge Deletion. Let (u, v) be an edge that is deleted. If the edge has not been matched, then after removing the edge from the graph all invariants still hold. If it has been matched, then u and v are now free. Therefore, the first invariant may be violated. If (u, v) is at level 0, then the algorithm tries to settle both endpoints by scanning their sets of owned edges. If (u, v) is at level 1, then u the algorithm does the following: First, u disowns all its edges whose other endpoint is at level 1. If \mathcal{O}_u is still greater than or equal to \sqrt{n} , then u stays at level 1 and executes RANDOM-SETTLE(u). If u owns less than \sqrt{n} edges, it moves to level 0 and tries to settle it by scanning its set of owned edges. The transition of u from level 1 to 0 leads to an increase in the number of edges owned by each of its neighbors at level 0. This may violate the size constraint of owned edges for those neighbors. Hence, the algorithm calls RANDOM-SETTLE for each neighbor that violates the constraint, which moves it to level 1.

3.4 Neiman and Solomon Algorithm

In contrast to the BGS algorithm [4], which is randomized, Neiman and Solomon (NS) [33] present a deterministic algorithm for maintaining a maximal matching in a dynamic graph. Their approach guarantees, that the maintained matching is a $3/2$ -approximate maximum matching and that update time is $O(\sqrt{m})$ in *worst case*, where m denotes the number of edges present in the graph in the moment of the update. NS maintains the following invariants: There are no augmenting paths of length ≤ 3 , ensuring $3/2$ -approx matching. All free vertices have degree at most $\sqrt{2n + 2m}$.

► **Lemma 2** ([33]). *Any free vertex of degree larger than \sqrt{m} can always be matched, so as to generate a free vertex with degree less than \sqrt{m} . This can be achieved in $O(\sqrt{m})$ time.*

Edge Insertion. Let (u, v) be the edge being inserted. If both the endpoints are free, the edge is simply added to the matching. Also, if both endpoints are matched it does not entail any further processing. However, if exactly one endpoint of the edge, say u , is matched, they try to remove a possible augmenting path of length 3 as follows. The neighbours of the mate of u say $u' = \text{mate}(u)$, are scanned for a free vertex, say x . If such a free vertex exists, an augmenting path of length 3 has been found, which is augmented increasing the matching size.

Edge Deletion. Let (u, v) be the edge being deleted. If the edge was unmatched, its deletion cannot create any new augmenting paths. However, if it was a matched edge, both the endpoints become free after the edge deletion. First, the algorithm checks for both freed vertices whether they have free neighbours and if so matches the freed vertices with those free neighbours. Now, in order to eliminate augmenting paths of length 3 starting from a free vertex, say u , all neighbours w of u are scanned checking if $\text{mate}(w)$ has a free neighbour. By providing appropriate data structures, this can be done in $O(\sqrt{n})$ time. If an augmenting path has been found, it is augmented increasing the size of the matching by one. If no augmenting path has been found, vertex u remains free, but only if its degree is at most $\sqrt{2m}$. If the degree of u exceeds $\sqrt{2m}$, using Lemma 2 a surrogate can be found in $O(\sqrt{m})$.

The overall update time of the algorithm is bound by the bounded degree of all free vertices, making any linear search through the neighbourhood $N(u)$ of a vertex u cost at most $O(\sqrt{n+m})$. Bounding the degree can further be achieved in $O(\sqrt{m})$ time using Lemma 2.

4 Experimental Evaluation

Implementation and System

We implemented the algorithms described in the previous section. The codes are written in C++ and have been compiled using g++-7.3.0 with flags `-O3`. All codes are sequential. We plan to further improve the codes and then to release them to make it available to a larger audience. Our experiments are conducted on one core of a machine with AMD Opteron Processors 6174 with 2.2GHz and 256GB of RAM. *Dynamic Graph Data Structure*: our algorithms use the following dynamic graph data structure. For each node v , we maintain a vector L_v of adjacent nodes, and a hash table \mathcal{H}_v that maps a vertex u that is incident to v to its position in L_v . This data structure allows for expected constant time insertion and deletion as well as a constant time operation to select a random neighbor of v . The deletion operation on (v, u) is implemented as follows: get the position of u in L_v via a lookup in $\mathcal{H}_v(u)$. Swap the element in L_v with the last element w in the vector and update the position of w in \mathcal{H}_v . Finally, pop the last element (now u) from L_v and delete its entry from \mathcal{H}_v .

Instances and Methodology

By default we perform ten repetitions per instance. We measure the total time taken to compute all edge insertions and deletions and generally use the *geometric mean* when averaging over different instances in order to give every instance a comparable influence on the final result. In order to compare different algorithms, we use *performance profiles* [14]. These plots relate the matching size / running time of all algorithms to the corresponding matching size / running time produced / consumed by each algorithm. More precisely, the y -axis shows $\#\{\text{objective} \geq \tau * \text{best}\} / \#\text{instances}$, where *objective* corresponds to the result of an algorithm on an instance and *best* refers to the best result of any algorithm shown within the plot. When we look at running time, the y -axis shows $\#\{t \leq \text{fastest} / \tau\} / \#\text{instances}$ (as a function of the parameter τ), where t corresponds to the time of an algorithm on an instance and *fastest* refers to the time of the fastest algorithm on that instance. The parameter $\tau \leq 1$ in this equation is plotted on the x -axis. For each algorithm, this yields a non-decreasing, piecewise constant function. Thus, if we are interested in the number of instances where an algorithm is the best/fastest, we only need to look at $\tau = 1$.

Instances


We evaluate our algorithms on a number of large graphs. These graphs are collected from [3, 13, 27, 26, 35]. Table 3 summarizes the main properties of the benchmark set. Our benchmark set includes a number of graphs from numeric simulations as well as complex networks. These include static graphs as well as real dynamic graphs. As our algorithms do only handle undirected graphs, we consider all input graphs to be undirected by ignoring edge directions and we remove self-loops and parallel edges. We perform *two different types* of experiments. First, we use the algorithms using insertions only, i.e. we start with an empty graph and insert all edges of the static graph in a random order. We do this with all graphs from Table 3. Second, we use real dynamic instances from Table 4. Most of these instances, however, only feature insertions (with the exception being `dewiki` and `wiki-simple-en` which have both real insert and real delete operations). Hence, we perform additional experiments with fully dynamic graphs from these inputs, by undoing x percent of the update operations performed last.

4.1 Random Walk and Blossom-based Algorithms

In this section, we use our algorithms with random insertions only. More precisely, we use the static graphs from Table 3. For each experiment, we start with an empty graph and insert edges of the static input in random order until all edges are inserted compare the result of our dynamic algorithms the maximum matching on the final graph Edmond [17].

Random Walk-Based Algorithms. We start with random walk-based algorithms. Preliminary experiments have shown that decreasing ϵ is more effective in getting better solutions than performing more repeated random walks at the start node. Hence, we exclude algorithms that perform multiple repetitions of random walks per insert operation here from the evaluation and focus on the different values of ϵ . We vary $\epsilon \in \{0.1, 0.25, 0.5, 1, 2\}$. Recall that the path length of a single random walk is then bounded by $2/\epsilon - 1$. If *all* paths of that length were explored, the algorithms would be guaranteed to give a $(1 + \epsilon)$ -approximation. Figure 1 summarizes the result. It is not surprising that the algorithm needs more running time for smaller ϵ , but also yields better results with increasing path lengths. On average, the algorithm is 2.4%, 3.2%, 4.2%, 5.5%, 11.5% percent away from the optimum for $\epsilon = 0.1, 0.25, 0.5, 1, 2$ respectively. Thus, even though the algorithms are not guaranteed to explore all paths of length $2/\epsilon - 1$, they achieve in practice an approximation that is much better than the theoretical bound for algorithms that explore all such paths. The strongest configuration ($\epsilon = 0.1$) is at most 1% away from the optimum matching size in 50% of the cases. Note that the random walk algorithm does not achieve the guarantee of 1% approximation as claimed by Lemma 1 since we did not perform the vast amount of repetitions necessary to get the result in expectation – instead we performed a single repetition of the random walker for each insertion. As expected the running time does increase with decreasing ϵ . However, due to random walks that can finish early because they managed to match an edge, the effect is less visible than theory expects. The running time increase over the random walk using $\epsilon = 2$ (which is essentially a random walker not allowed to move, and hence boils down to the very simple greedy algorithm), is 12%, 17%, 21%, 27% for $\epsilon = 1, 0.5, 0.25, 0.1$, respectively.

Enabling Δ -settling generally improves the result. On average, the random walk with Δ -settling is now 1.1%, 1.4%, 1.8%, 2.2%, 3.7% away from the optimum matching for $\epsilon = 0.1, 0.25, 0.5, 1, 2$, respectively. On average in our experiments using Δ -settling has a negligible impact on running time. Hence, we recommend to use Δ -settling when using random walk-based algorithms and do so in the following unless otherwise mentioned.



■ **Figure 1** Performance profile for matching size $|\mathcal{M}|$ for Random Walk and for Unsafe Dyn Opt.

(Optimum) Blossom-Based Algorithms. We now consider dynamic blossom-based algorithms from Section 3.2. We start this section with the version of the algorithm that maintains the optimum matching, and compare it to the naive dynamic optimum matching algorithm that recomputes a maximum matching from scratch each time an edge is inserted. Since the running time of the naive optimum algorithm is fairly excessive, we run it only on the graphs of our benchmark set having less than 25k nodes. First of all, our dynamic algorithm that maintains the optimum matching is more than an order of magnitude faster than the naive optimum algorithm (roughly a factor 12). We expect that the difference will be even more pronounced if even larger graphs are used. Running our dynamic optimum algorithm with the unsafe option indeed significantly speeds up the algorithm – the lazy augmenting path search configuration is more than two orders of magnitude over the safe version of our algorithm (roughly a factor 115). The improvements in running time stem

from the fact that our algorithms try to maintain a very large matching. Hence, the case that is executed by the safe option often does not find an augmenting path which implies that the augmenting path search has to look at the overall network and hence reaches its worst-case complexity. Of course, the unsafe option does not have a guarantee on optimality anymore. In our experiments, the unsafe option computes matchings that are 0.02% worse than the optimum on average. We conclude that the algorithm maintains near-optimum matchings while being three orders of magnitude faster than the naive optimum dynamic algorithm. Henceforth, we only consider the unsafe version of our algorithm.

We now switch our set of graphs back to all of our benchmark graphs from Table 3. Using lazy augmenting path search in the unsafe algorithm additionally speedups up computations. Unsafe+LP is on average 20.5 faster than the unsafe algorithm without lazy augmenting path search – again at the cost of solution quality. The algorithm is already only 30% slower than running the static algorithm a *single* time on the final graph that contains all edges. On the other hand, the unsafe dynamic algorithm using lazy augmenting path search computes 0.6% worse matchings than the unsafe algorithm without lazy augmenting path search.


Lastly, we focus on the third variation of the algorithm, which is to bound the depth of the augmenting path search to that is done during update operations. The depth is bounded to $2/\epsilon - 1$ so that given ϵ and running the safe option of the algorithm would maintain a $1 + \epsilon$ approximate matching. We, however, only consider the unsafe version of the algorithm. We use same values of $\epsilon = 0.1, 0.25, 0.5, 1$ as in the random walk-based algorithms section, but do not consider $\epsilon = 2$, since this is again essentially the very simple greedy algorithm. Moreover, we run the algorithm with and without the lazy augmenting path search.

First of all, running without lazy augmenting path search, the algorithm indeed maintains the approximation guarantee. On average, the algorithm is 0.1%, 0.4%, 1.5% and 3.6% worse than the optimum algorithm for $\epsilon = 0.1, 0.25, 0.5, 1$, respectively. Using the lazy augmenting path search, speeds up to algorithm by a factor of 5.52, 2.65, 2.03, 1.76 for $\epsilon = 0.1, 0.25, 0.5, 1$, respectively. With lazy augmenting path search, the algorithm is 0.6%, 1.00%, 2.2%, 4.3% worse than the optimum for $\epsilon = 0.1, 0.25, 0.5, 1$, respectively (and hence still achieves the approximation guarantee). The algorithm using $\epsilon = 0.1$ is only 0.2% worse than the algorithm not bounding the depth. However, the algorithm is also not much faster. On average, bounding the search depth with $\epsilon = 0.1$ improves running time by 6%. Figure 1 shows a summarizing performance profile.

4.2 Comparison of Algorithms

Dynamic Sequences from Static Graphs. We now compare all of the different non-optimal algorithms against each other for the insertion-only case. For random-walks, we always enable Δ -settling, for blossom-based algorithm always use the unsafe option and with and without lazy augmenting path search. Table 1 shows average results for matching size and running time after all edges and operations have been performed. Figure 2 shows performance profiles for running time and for matching size.

First of all, both the blossom-based (with lazy augmenting path search) and random walk-based algorithms dominate the algorithms by Neiman Solomon and Baswana Gupta Sen (BGS). The algorithms find consistently larger matchings and do so in less time. However, note that the real-world instances we look at rarely have nodes with more than \sqrt{n} neighbors, so that the BGS algorithm is roughly similar to the simple greedy algorithm. We also try to use $c \cdot \sqrt{n}$ as a threshold for different values of c , but this always resulted in worse matching sizes.



■ **Figure 2** Performance profile for matching size $|\mathcal{M}|$ and time for all algorithms. In all cases, if an algorithm has a curve closer to the upper left corner, then the algorithm is better.

In general, performance differences in running time are not very big (except if we don't use lazy augmenting path search in the blossom-based algorithms). Secondly, for the same values of ϵ the blossom-based algorithms compute slightly better results than their random walk-based counter parts. This is not surprising as the blossom-based algorithms explore larger subgraphs for each edge that has been inserted. We conclude here that both types of algorithms are feasible in practice and have an advantage in solution quality over Neiman Solomon and Baswana Gupta Sen on graphs with random insertions. Moreover, both of these algorithm yield a clear trade-off between running time and solution quality via the ϵ parameter. On the other hand, all of the algorithms considered here are roughly five orders of magnitude faster than the naive dynamic optimum algorithm (only considering instances having less than 25k nodes) and except Baswana Gupta Sen, all of these algorithms are within a 4% range of the optimum matching size.

Real-World Dynamic Instances. We now switch to the real-world dynamic instances. As already mentioned, most of these instances are insertion-only. Hence, we perform additional experiments with fully dynamic graphs from these inputs, by undoing x percent of the update operations performed last (call them \mathcal{O}_x). More precisely, we perform the operations in \mathcal{O}_x in reverse order. More precisely, if an edge operation was an insertion in \mathcal{O}_x , we perform a delete operation and if it was a delete operation we insert it. As before, we compute the update on the graph after each removal/insertion. The connection to practice in this case, is that with undoing operations, we want to restore a previous state. Table 2 summarizes the results of the experiment and Figure 3 compares the algorithms on the two real-world dynamic graphs dewiki and wiki_simple_en. Overall, the situation is similar to experiments with random insertions that we have seen before. The random walk with Δ -settling and $\epsilon = 0.5$ dominates Baswana, Gupta, Sen and Neiman Solomon in terms of running time *and* matching size for every number of undo operations. The blossom-based algorithm with lazy path search, however, yields smaller matchings if no operations are undone. We believe that this is due to the edges not being inserted randomly and hence the lazy augmenting path search heuristic is less effective, and misses augmenting paths that have been created over time. If operations are undone, the blossom-based algorithms outperform Neiman Solomon in terms of matching size, but are also considerably slower as the deletion operations search for augmenting paths of lengths three (except for $\epsilon = 1$). The blossom-based algorithm without lazy augmenting path search get very close to the optimum solutions. The best algorithm here is blossom-based algorithm without lazy augmenting path search for $\epsilon = 1/3$. On average, it computes solutions that are $< 0.6\%$ away from the optimum (for every amount of undo operations done).

In general, all algorithms improve quality relative to the optimum matching size, if we undo operations. This is due to the fact that the matching may have changed over time and hence new (short) augmenting paths may be found. In case of random-walks this is also simply due to the fact that additional work is performed and the likelihood to find an augmenting path is increased by running additional random walks. Summing up, all of the algorithms, except Baswana, Gupta, Sen, compute/maintain very large matchings. Blossom-based and random walk-based algorithm are highly flexible and are able to trade solution quality for time. Overall, random walk-based algorithms seem to be the method of choice in practice.

■ **Table 1** Random insertions from static graphs: mean of the matching size relative to optimum after all operations have been done as well as mean increase in running time over Random Walk, $\Delta, \epsilon = 0.25$.

algorithm		mean $ \mathcal{M} / \mathcal{M}_{\text{opt}} $	rel. time
BGS		0.885	32%
Neiman Solomon		0.964	28%
Unsafe Dyn Opt+LP	$\epsilon = 0.1$	0.994	27%
Unsafe Dyn Opt+LP	$\epsilon = 0.25$	0.990	11%
Unsafe Dyn Opt	$\epsilon = 0.1$	0.999	613%
Unsafe Dyn Opt	$\epsilon = 0.25$	0.996	192%
Unsafe Dyn Opt	$\epsilon = 0.5$	0.985	101%
Unsafe Dyn Opt	$\epsilon = 1$	0.964	67%
Random Walk, Δ	$\epsilon = 0.1$	0.989	5%
Random Walk, Δ	$\epsilon = 0.25$	0.986	1


■ **Table 2** Real-world dynamic instances: mean of the matching size relative to optimum after all operations have been done as well as the mean increase in running time over Random Walk, $\Delta, \epsilon = 0.5$.

# undo op		0	5%	10%	25%
algorithm		mean $ \mathcal{M} / \mathcal{M}_{\text{opt}} $			
BGS		0.845	0.847	0.848	0.851
Neiman Solomon		0.968	0.971	0.973	0.976
Unsafe Dyn Opt+LP	$\epsilon = 0.1$	0.947	0.985	0.990	0.996
Unsafe Dyn Opt+LP	$\epsilon = 0.25$	0.942	0.982	0.988	0.993
Unsafe Dyn Opt	$\epsilon = 0.\overline{33}$	0.994	0.996	0.997	0.998
Unsafe Dyn Opt	$\epsilon = 0.5$	0.988	0.991	0.992	0.994
Unsafe Dyn Opt	$\epsilon = 1$	0.968	0.971	0.973	0.976
Random Walk, Δ	$\epsilon = 0.1$	0.982	0.984	0.985	0.986
Random Walk, Δ	$\epsilon = 0.25$	0.981	0.983	0.984	0.985
Random Walk, Δ	$\epsilon = 0.5$	0.978	0.980	0.981	0.982
algorithm		rel. time			
BGS		4%	14%	13%	18%
Neiman Solomon		64%	82%	92%	112%
Unsafe Dyn Opt+LP	$\epsilon = 0.1$	82%	306%	383%	633%
Unsafe Dyn Opt+LP	$\epsilon = 0.25$	23%	149%	200%	346%
Unsafe Dyn Opt	$\epsilon = 0.\overline{33}$	1 551%	1 551%	1 598%	1814%
Unsafe Dyn Opt	$\epsilon = 0.5$	679%	682%	713%	800%
Unsafe Dyn Opt	$\epsilon = 1$	210%	212%	223%	250%
Random Walk, Δ	$\epsilon = 0.1$	25%	26%	24%	24%
Random Walk, Δ	$\epsilon = 0.25$	10%	11%	9%	11%
Random Walk, Δ	$\epsilon = 0.5$	1	1	1	1

5 Conclusion

We looked at several dynamic matching algorithms including Baswana, Gupta and Sen [4], Neiman and Solomon [33], as well as random walk-based algorithms and blossom-based algorithms. We performed extensive experiments comparing the performance of these algorithms on the real-world and artificially generated instances. In terms of results, first we have shown that maintaining optimum matchings can be done much more efficiently than the naive algorithm that recomputes maximum matchings from scratch. Second, we have seen that all non-optimum dynamic algorithms that we considered in this work are able to maintain near-optimum matchings in practice while being multiple orders of magnitudes faster than the naive optimum dynamic algorithm. In practice, random walk-based algorithms with Δ -settling will be the method of choice.

In future work, it may be interesting to transfer results to the weighted case, and to combine our algorithms with simple data reductions rules such as [25]. It could be interesting to use these dynamic matching algorithms to derive dynamic multilevel algorithms for example for graph partitioning [32, 1]. Another direction will be to explore the parallelization potential of random walk-based algorithms. Lastly, it may be interesting to incorporate dynamic transitive closure algorithms (e.g. [22]) into the delete operation of the dynamic optimum matching algorithm to further reduce the number of augmenting path searches.



■ **Figure 3** Matching size over time compared to Baswana, Gupta, Sen on the two real dynamic instances dewiki and wiki_simple_en.

References

- 1 Yaroslav Akhremtsev, Peter Sanders, and Christian Schulz. High-quality shared-memory graph partitioning. In Marco Aldinucci, Luca Padovani, and Massimo Torquati, editors, *Euro-Par 2018: Parallel Processing - 24th International Conference on Parallel and Distributed Computing, Turin, Italy, August 27-31, 2018, Proceedings*, volume 11014 of *Lecture Notes in Computer Science*, pages 659–671. Springer, 2018. doi:10.1007/978-3-319-96983-1_47.
- 2 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.
- 3 D. Bader, A. Kappes, H. Meyerhenke, P. Sanders, C. Schulz, and D. Wagner. Benchmarking for Graph Clustering and Partitioning. In *Encyclopedia of Social Network Analysis and Mining*. Springer, 2014.

- 4 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.
- 5 Claude Berge. Two theorems in graph theory. *Proceedings of the National Academy of Sciences*, 43(9):842–844, 1957. doi:10.1073/pnas.43.9.842.
- 6 Aaron Bernstein and Cliff Stein. Faster fully dynamic matchings with small approximation ratios. In *Proceedings of the 27th Symposium on Discrete Algorithms SODA*, pages 692–711. SIAM, 2016. doi:10.1137/1.9781611974331.ch50.
- 7 Sayan Bhattacharya, Deeparnab Chakrabarty, and Monika Henzinger. Deterministic fully dynamic approximate vertex cover and fractional matching in $O(1)$ amortized update time. In *19th International Conf. on Integer Programming and Combinatorial Optimization IPCO*, pages 86–98, 2017.
- 8 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.
- 9 Sayan Bhattacharya, Monika Henzinger, and Danupon Nanongkai. New deterministic approximation algorithms for fully dynamic matching. In *Proceedings of the 48th Annual Symposium on Theory of Computing*, pages 398–411. ACM, 2016.
- 10 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, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms SODA*, pages 470–489. SIAM, 2017.
- 11 Marcel Birn, Vitaly Osipov, Peter Sanders, Christian Schulz, and Nodari Sitchinava. Efficient parallel and external matching. In *Euro-Par 2013*, volume 8097 of *LNCS*, pages 659–670. Springer, 2013. doi:10.1007/978-3-642-40047-6_66.
- 12 Moses Charikar and Shay Solomon. Fully dynamic almost-maximal matching: Breaking the polynomial worst-case time barrier. In *45th International Colloquium on Automata, Languages, and Programming ICALP*, pages 33:1–33:14, 2018.
- 13 T. Davis. The University of Florida Sparse Matrix Collection, <http://www.cise.ufl.edu/research/sparse/matrices>, 2008. URL: <http://www.cise.ufl.edu/research/sparse/matrices/>.
- 14 Elizabeth D. Dolan and Jorge J. Moré. Benchmarking optimization software with performance profiles. *Math. Program.*, 91(2):201–213, 2002. doi:10.1007/s101070100263.
- 15 D. Drake and S. Hougardy. A Simple Approximation Algorithm for the Weighted Matching Problem. *Information Processing Letters*, 85:211–213, 2003.
- 16 Andre Droschinsky, Petra Mutzel, and Erik Thorndsen. Shrinking trees not blossoms: A recursive maximum matching approach. In *Proceedings of the Symposium on Algorithm Engineering and Experiments, ALENEX 2020*, pages 146–160. SIAM, 2020. doi:10.1137/1.9781611976007.12.
- 17 Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17(3):449–467, 1965.
- 18 Harold Neil Gabow. *Implementation of Algorithms for Maximum Matching on Nonbipartite Graphs*. PhD thesis, Stanford University, Stanford, CA, USA, 1974.
- 19 Zvi Galil, Silvio Micali, and Harold N. Gabow. An $O(|E||V| \log |V|)$ algorithm for finding a maximal weighted matching in general graphs. *SIAM Journal Computing*, 15(1):120–130, 1986.
- 20 Fabrizio Grandoni, Stefano Leonardi, Piotr Sankowski, Chris Schwegelshohn, and Shay Solomon. $(1 + \epsilon)$ -approximate incremental matching in constant deterministic amortized time. In *Proceedings of the 20th Symposium on Discrete Algorithms*, pages 1886–1898. SIAM, 2019. doi:10.1137/1.9781611975482.114.
- 21 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. doi:10.1109/FOCS.2013.65.

- 22 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.
- 23 J. E. Hopcroft and R. M. Karp. A $n^{5/2}$ algorithm for maximum matchings in bipartite. In *12th Annual Symposium on Switching and Automata Theory (SWAT)*, pages 122–125, 1971. doi:10.1109/SWAT.1971.1.
- 24 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.
- 25 Viatcheslav Korenwein, André Nichterlein, Rolf Niedermeier, and Philipp Zschoche. Data reduction for maximum matching on real-world graphs: Theory and experiments. In *26th European Symposium on Algorithms ESA*, volume 112 of *LIPICs*, pages 53:1–53:13. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018. doi:10.4230/LIPICs.ESA.2018.53.
- 26 Jérôme Kunegis. KONECT: the koblenz network collection. In Leslie Carr, Alberto H. F. Laender, Bernadette Farias Lóscio, Irwin King, Marcus Fontoura, Denny Vrandečić, Lora Aroyo, José Palazzo M. de Oliveira, Fernanda Lima, and Erik Wilde, editors, *22nd World Wide Web Conference, WWW '13*, pages 1343–1350. International World Wide Web Conferences Steering Committee / ACM, 2013. doi:10.1145/2487788.2488173.
- 27 J. Lescovec. Stanford Network Analysis Package (SNAP). <http://snap.stanford.edu/index.html>.
- 28 J. Maue and P. Sanders. Engineering Algorithms for Approximate Weighted Matching. In *Proceedings of the 6th Workshop on Experimental Algorithms (WEA'07)*, volume 4525 of *LNCS*, pages 242–255. Springer, 2007. doi:10.1007/978-3-540-72845-0_19.
- 29 Kurt Mehlhorn and Stefan Näher. *LEDA: A Platform for Combinatorial and Geometric Computing*. Cambridge University Press, 1999. URL: <http://www.mpi-sb.mpg.de/%7Emehlhorn/LEDAbook.html>.
- 30 Aranyak Mehta, Amin Saberi, Umesh V. Vazirani, and Vijay V. Vazirani. Adwords and generalized on-line matching. In *46th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 264–273. IEEE Computer Society, 2005. doi:10.1109/SFCS.2005.12.
- 31 Silvio Micali and Vijay V. Vazirani. An $O(\sqrt{|V||E|})$ algorithm for finding maximum matching in general graphs. In *21st Symposium on Foundations of Computer Science*, pages 17–27. IEEE Computer Society, 1980. doi:10.1109/SFCS.1980.12.
- 32 Orlando Moreira, Merten Popp, and Christian Schulz. Evolutionary multi-level acyclic graph partitioning. In Hernán E. Aguirre and Keiki Takadama, editors, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2018, Kyoto, Japan, July 15-19, 2018*, pages 332–339. ACM, 2018. doi:10.1145/3205455.3205464.
- 33 Ofer Neiman and Shay Solomon. Simple deterministic algorithms for fully dynamic maximal matching. *ACM Trans. Algorithms*, 12(1):7:1–7:15, 2016.
- 34 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.
- 35 Julia Preusse, Jérôme Kunegis, Matthias Thimm, Thomas Gottron, and Steffen Staab. Structural dynamics of knowledge networks. In *Proc. Int. Conf. on Weblogs and Social Media*, 2013.
- 36 Piotr Sankowski. Faster dynamic matchings and vertex connectivity. In *SODA*, pages 118–126, 2007. doi:10.1145/1283383.1283397.
- 37 Jeremy G. Siek, Lie-Quan Lee, and Andrew Lumsdaine. *The Boost Graph Library - User Guide and Reference Manual*. C++ in-depth series. Pearson / Prentice Hall, 2002.
- 38 Shay Solomon. Fully dynamic maximal matching in constant update time. In *57th Symposium on Foundations of Computer Science FOCS*, pages 325–334, 2016.
- 39 Robert Endre Tarjan. *Data structures and network algorithms*, volume 44 of *CBMS-NSF regional conference series in applied mathematics*. SIAM, 1983. doi:10.1137/1.9781611970265.

A Instances

■ **Table 3** Basic properties of the benchmark set of static graphs obtained from [3, 13, 27].

graph	n	m	graph	n	m
144	144 649	1 074 393	eu-2005	862 664	16 138 468
3elt	4 720	13 722	fe_4elt2	11 143	32 818
4elt	15 606	45 878	fe_body	45 087	163 734
598a	110 971	741 934	fe_ocean	143 437	409 593
add20	2 395	7 462	fe_pwt	36 519	144 794
add32	4 960	9 462	fe_rotor	99 617	662 431
amazon-2008	735 323	3 523 472	fe_sphere	16 386	49 152
as-22july06	22 963	48 436	fe_tooth	78 136	452 591
as-skitter	554 930	5 797 663	finan512	74 752	261 120
auto	448 695	3 314 611	in-2004	1 382 908	13 591 473
bcstkt29	13 992	302 748	loc-brightkite_edges	56 739	212 945
bcstkt30	28 924	1 007 284	loc-gowalla_edges	196 591	950 327
bcstkt31	35 588	572 914	m14b	214 765	1 679 018
bcstkt32	44 609	985 046	memplus	17 758	54 196
bcstkt33	8 738	291 583	p2p-Gnutella04	6 405	29 215
brack2	62 631	366 559	PGPgiantcompo	10 680	24 316
citationCiteseer	268 495	1 156 647	rgg_n_2_15_s0	32 768	160 240
cnr-2000	325 557	2 738 969	soc-Slashdot0902	28 550	379 445
coAuthorsCiteseer	227 320	814 134	t60k	60 005	89 440
coAuthorsDBLP	299 067	977 676	uk	4 824	6 837
coPapersCiteseer	434 102	16 036 720	vibrobox	12 328	165 250
coPapersDBLP	540 486	15 245 729	wave	156 317	1 059 331
crack	10 240	30 380	web-Google	356 648	2 093 324
cs4	22 499	43 858	whitaker3	9 800	28 989
cti	16 840	48 232	wiki-Talk	232 314	1 458 806
data	2 851	15 093	wing	62 032	121 544
email-EuAll	16 805	60 260	wing_nodal	10 937	75 488
enron	69 244	254 449	wordassociation-2011	10 617	63 788

■ **Table 4** Basic properties of the benchmark set of dynamic graphs with number of update operations \mathcal{O} . Most of the graphs only feature insertions. The only two exceptions are marked with a *. All of these graphs have been obtained from the KONECT graph database [35].

graph	n	\mathcal{O}
amazon-ratings	2 146 058	5 838 041
citeulike_ui	731 770	2 411 819
dewiki*	2 166 670	86 337 879
dnc-temporalGraph	2 030	39 264
facebook-wosn-wall	46 953	876 993
flickr-growth	2 302 926	33 140 017
haggle	275	28 244
lastfm_band	174 078	19 150 868
lkml-reply	63 400	1 096 440
movielens10m	69 879	10 000 054
munmun_digg	30 399	87 627
proper_loans	89 270	3 394 979
sociopatterns-infections	411	17 298
stackexchange-stackoverflow	545 197	1 301 942
topology	34 762	171 403
wikipedia-growth	1 870 710	39 953 145
wiki_simple_en*	100 313	1 627 472
youtube-u-growth	3 223 590	9 375 374