

Fully Dynamic MIS in Uniformly Sparse Graphs

Krzysztof Onak

IBM Research, TJ Watson Research Center, Yorktown Heights, New York, USA

Baruch Schieber

IBM Research, TJ Watson Research Center, Yorktown Heights, New York, USA

Shay Solomon¹

IBM Research, TJ Watson Research Center, Yorktown Heights, New York, USA

Nicole Wein²

Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

Abstract

We consider the problem of maintaining a maximal independent set (MIS) in a dynamic graph subject to edge insertions and deletions. Recently, Assadi, Onak, Schieber and Solomon (STOC 2018) showed that an MIS can be maintained in *sublinear* (in the dynamically changing number of edges) amortized update time. In this paper we significantly improve the update time for *uniformly sparse graphs*. Specifically, for graphs with arboricity α , the amortized update time of our algorithm is $O(\alpha^2 \cdot \log^2 n)$, where n is the number of vertices. For low arboricity graphs, which include, for example, minor-free graphs as well as some classes of “real world” graphs, our update time is polylogarithmic. Our update time improves the result of Assadi *et al.* for all graphs with arboricity bounded by $m^{3/8-\epsilon}$, for any constant $\epsilon > 0$. This covers much of the range of possible values for arboricity, as the arboricity of a general graph cannot exceed $m^{1/2}$.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms, Theory of computation → Graph algorithms analysis, Theory of computation → Dynamic graph algorithms

Keywords and phrases dynamic graph algorithms, independent set, sparse graphs, graph arboricity

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.92

Related Version A full version of the paper is available at [23].

1 Introduction

The importance of the maximal independent set (MIS) problem is hard to overstate. In general, MIS algorithms constitute a useful subroutine for locally breaking symmetry between several choices. The MIS problem has intimate connections to a plethora of fundamental combinatorial optimization problems such as maximum matching, minimum vertex cover, and graph coloring. As a prime example, MIS is often used in the context of graph coloring, as all vertices in an independent set can be assigned the same color. As another important example, one can compute a large matching (approximating the maximum matching to within a factor arbitrarily close to 1) by applying maximal independent sets of longer and longer augmenting paths as observed by Hopcroft and Karp [11]. The seminal papers of Luby [18] and Linial [17] discuss additional applications of MIS. A non-exhaustive list of further direct and indirect

¹ Supported by the IBM Herman Goldstine Postdoctoral Fellowship.

² Supported by an NSF Graduate Fellowship.



© Nicole Wein and IBM Research;

licensed under Creative Commons License CC-BY

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).

Editors: Ioannis Chatzigiannakis, Christos Kaklamani, Dániel Marx, and Donald Sannella;

Article No. 92; pp. 92:1–92:14



Leibniz International Proceedings in Informatics

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



applications of MIS includes resource allocation [25], leader election [7], the construction of network backbones [16, 12], and sublinear-time approximation algorithms [22].

In the 1980s, questions concerning the computational complexity of the MIS problem spurred a line of research that led to the celebrated parallel algorithms of Luby [18], and Alon, Babai, and Itai [1]. These algorithms find an MIS in $O(\log n)$ rounds without global coordination. More recently, Fischer and Noever [8] gave an even simpler greedy algorithm that considers vertices in random order and takes $O(\log n)$ rounds with high probability (see also an earlier result of Blelloch, Fineman, and Shun [4]).

In this work we continue the study of the MIS problem by considering the *dynamic setting*, where the underlying graph is not fixed, but rather evolves over time via edge updates. Formally, a *dynamic graph* is a graph sequence $\mathcal{G} = (G_0, G_1, \dots, G_M)$ on n fixed vertices, where the initial graph is $G_0 = (V, \emptyset)$ and each graph $G_i = (V, E_i)$ is obtained from the previous graph G_{i-1} in the sequence by either adding or deleting a single edge. The basic goal in this context is to maintain an MIS in time significantly faster than it takes to recompute it from scratch following every edge update.

In STOC'18, Assadi, Onak, Schieber, and Solomon [2] gave the first sub-linear (amortized) update time fully dynamic algorithm for maintaining a MIS. Their amortized update time is $\min\{m^{3/4}, \Delta\}$, where m is the (dynamically changing) number of edges, and Δ is a fixed bound on the maximum degree of the graph. For graphs of high maximum degree, the update time of the algorithm of [2] decreases as the graph becomes sparser.

1.1 Our contribution

We focus on graphs that are “uniformly sparse” or “sparse everywhere”, as opposed to the previous work by Assadi et al. [2] that considers unrestricted sparse graphs. We aim to improve the update time of [2] as a function of the “uniform sparsity” of the graph. This fundamental property of graphs has been studied in various contexts and names over the years, one of which is via the notion of *arboricity* [19, 20, 24]:

► **Definition 1.1.** The arboricity α of a graph $G = (V, E)$ is defined as $\alpha = \max_{U \subset V} \lceil \frac{|E(U)|}{|U|-1} \rceil$, where $E(U) = \{(u, v) \in E \mid u, v \in U\}$.

Thus a graph has bounded arboricity if every induced subgraph has bounded density. The family of low arboricity graphs contains, among others, bounded-degree graphs, all minor-closed graph classes (e.g., planar graphs, graphs with bounded treewidth), and randomly generated preferential attachment graphs. Moreover, it is believed that many real-world graphs such as the world wide web graph and social networks also have low arboricity [9].

A dynamic graph of *arboricity* α is a dynamic graph such that all graphs G_i have arboricity bounded by α . We prove the following result.

► **Theorem 1.2.** *For any dynamic n -vertex graph of arboricity α , an MIS can be maintained deterministically in $O(\alpha^2 \log^2 n)$ amortized update time.*

Theorem 1.2 improves the result of Assadi et al. for all graphs with arboricity bounded by $m^{3/8-\epsilon}$, for any constant $\epsilon > 0$. This covers much of the range of possible values for arboricity, as the arboricity of a general graph cannot exceed \sqrt{m} . Furthermore, we obtain polylogarithmic update time for graphs of polylogarithmic arboricity; in particular, for the family of constant arboricity graphs the update time is $O(\log^2 n)$.

1.2 Our and previous techniques

1.2.1 The dynamic edge orientation problem

Our algorithm utilizes two properties of arboricity α graphs: a) every subgraph contains a vertex of degree at most 2α , and b) there exists an orientation of the edges so that every vertex has out-degree at most α . The first property follows from Definition 1.1 and the second property is due to an alternate definition [20].

Brodal and Fagerberg [5] initiated the study of the dynamic edge orientation problem and gave an algorithm that maintains an $O(\alpha)$ out-degree orientation in amortized $O(\alpha + \log n)$ time; our algorithm uses the algorithm of [5] as a *black box*. (Refer to [14, 10, 13, 3] for additional results on the dynamic edge orientation problem.)

1.2.2 A comparison to Assadi et al.

As noted already by Assadi et al. [2], a central obstacle in maintaining an MIS in a dynamic graph is the maintenance of a detailed 2-hop neighborhood of a vertex. We need this information to update the MIS. Consider the “hard” case that an edge is added to the graph and as a result, one of its endpoints v is removed from the MIS. In this case, to maintain the maximality of the MIS we need to identify and add to the MIS all neighbors of v that are not adjacent to a vertex in the MIS. This means that for each of v ’s neighbors, we need to know whether it has a neighbor the MIS (other than v). Alas, dynamically maintaining the 2-neighborhood of each vertex *explicitly* is prohibitive.

To overcome this hurdle, we build upon the approach of Assadi et al. [2] and maintain an *incomplete* 2-hop neighborhood. Consequently, we may err by adding vertices to the MIS even though they are adjacent to MIS vertices. To fix this, we need to remove vertices from the MIS. The important property that we maintain, following [2], is that the number of added vertices is significantly higher than the number of removed vertices. Like Assadi et al., we use a potential function defined as the number of vertices not in the MIS. To amortize the update time, if an update takes a long time, then the size of the MIS increases substantially as a result. On the other hand, the size of the MIS can only decrease by one in each update. Our algorithm deviates significantly from Assadi et al. in several respects, described next.

I: An underlying bounded out-degree orientation. We apply the algorithm of [5] for efficiently maintaining a bounded out-degree orientation. Given such an orientation, we can maintain 2-hop neighborhoods with respect to *outgoing edges* explicitly, which helps significantly in the maintenance of the MIS. More specifically, the usage of a bounded out-degree orientation enables us to reduce the problem to a single nontrivial case.

II: An intricate “chain reaction”. To handle the nontrivial case efficiently, we develop an intricate “chain reaction” process, initiated by adding vertices to the MIS that violate the independence property, which then forces the removal of other vertices from the MIS, which in turn forces the addition of other vertices, and so forth. Such a process may take a long time. However, a novel partition of a subset of the in-neighborhood of each vertex in the MIS into a logarithmic number of “buckets” (see below) together with a careful analysis allow for limiting this chain reaction to a logarithmic number of phases; upper bounding the number of phases is crucial for achieving a low update time. We note that such a careful analysis was not required in [2], where the chain reaction was handled by a simple recursive treatment.

III: A precise bucketing scheme. In order to achieve a sufficient increase in the size of the MIS, we need to carefully choose which vertices to add to the MIS so as to guarantee that at every step of the aforementioned chain reaction, even steps far in the future, we will be able to add enough vertices to the MIS. We achieve this by maintaining a precise *bucketing* of a carefully maintained subset of the in-neighborhood of each vertex, where vertices in larger-indexed buckets are capable of setting off longer and more fruitful chain reactions. At the beginning of the chain reaction, we add vertices in the larger-indexed buckets to the MIS, and gradually allow for vertices in smaller-indexed buckets to be added.

IV: A tentative set of MIS vertices. In contrast to the algorithm of Assadi et al., here we cannot iteratively choose which vertices to add to the MIS. Instead, we need to build a *tentative* set of vertices to add to the MIS, and only later prune this set to choose which vertices to add. If we do not carefully select vertices added to the MIS, we may not increase the size of the MIS sufficiently in order to amortize the cost of this process. To make sure that the size of the tentative set remains sufficiently large after the pruning, we make critical use of the first property of low arboricity graphs mentioned in Section 1.2.1. More details can be found at the beginning of Section 4.1.

1.2.3 A comparison to other previous work

Censor-Hillel, Haramaty, and Karnin [6] consider the problem of maintaining MIS in the dynamic distributed setting. They show that there is a randomized algorithm that requires only a constant number of rounds in expectation to update the maintained MIS, as long as the sequence of graph updates does not depend on the algorithm's randomness. This assumption is often referred to as *oblivious adversary*. As noted by Censor-Hillel et al., it is unclear whether their algorithm can be implemented with low total work in the centralized setting. This shortcoming is addressed by Assadi et al. [2] and by the current paper.

Kowalik and Kurowski [15] employ a dynamic bounded out-degree orientation to answer shortest path queries in (unweighted) planar graphs. Specifically, given a planar graph and a constant k , Kowalik and Kurowski maintain a data structure that can check in constant time whether two vertices are at distance at most k and if so produce a path of such length. This data structure is fully dynamic with polylogarithmic amortized update time. Like our data structure, their data structure maintains information on the 2-hop neighborhoods of the vertices, however, the nature of the 2-hop neighborhood information necessary for the two problems is different. For answering shortest path queries, one needs to maintain the *complete* 2-hop neighborhood. Whereas, to maintain an MIS, we only need to maintain *partial* information about the 2-hop neighborhood, however, the information that we store must be more detailed in the sense that for each 2-hop neighbor that we store, we need to know which 1-hop neighbors it is adjacent to. This is necessary because when we remove a vertex v from the MIS, we need to know which of v 's neighbors have no neighbors in the MIS in order to know which vertices we need to add to the MIS.

1.3 Dynamic MIS vs. dynamic maximal matching

In the maximal matching problem, the goal is to compute a matching that cannot be extended by adding another edge. The problem is equivalent to finding an MIS in the line graph of the input graph. However, despite this very close relationship between the MIS and maximal matching problems, (efficiently) *maintaining* an MIS appears to be inherently harder than maintaining a maximal matching. As a first potential evidence, one may notice that there

is a significant gap in the performance of the naive algorithms for these dynamic problems. For the maximal matching problem, the naive algorithm works in $O(\Delta)$ time. For the MIS problem, the naive algorithm has to inspect not only neighbors of a vertex v that is being removed from MIS, but also their neighbors (i.e. its 2-hop neighborhood) in order to find out which neighbors need to be added to the MIS; as a result the update time is $O(m)$. Furthermore, the worst-case number of MIS changes (by any algorithm) may be as high as $\Omega(\Delta)$ [6, 2], whereas the worst-case number of changes to the maximal matching maintained by the naive algorithm is $O(1)$. Also, the available body of work on the dynamic MIS problem is significantly sparser than for the dynamic maximal matching problem.

It is therefore plausible that it may be hard to obtain, for MIS, a bound better than the best bounds known for the dynamic maximal matching problem. In particular, the state-of-the-art dynamic deterministic algorithm for maintaining a maximal matching has an update time of $O(\sqrt{m})$ [21], even in the amortized sense. Hence, in order to obtain update time bounds sub-polynomial in m , one may have to exploit the structure of the graph, and bounded arboricity graphs are a natural candidate. A maximal matching can be maintained in graphs of arboricity bounded by α with amortized update time $O(\alpha + \sqrt{\alpha \lg n})$ [21, 10]; as long as the arboricity is polylogarithmic in n , the amortized update time for maintaining a maximal matching is polylogarithmic. In this work we show that essentially the same picture applies to the seemingly harder problem of dynamic MIS.

2 Algorithm overview

Using a bounded out-degree orientation of the edges, a very simple algorithm suffices to handle edge updates that fall into certain cases. The nontrivial case occurs when we remove a vertex v from the MIS and need to determine which vertices in v 's in-neighborhood have no neighbors in the MIS, and thus need to be added to the MIS. The in-neighborhood of v could be very large and it would be costly to spend even constant time per in-neighbor of v . Furthermore, it would be costly to maintain a data structure that stores for each vertex in the MIS which of its in-neighbors have no other neighbors in the MIS. Suppose we stored such a data structure for a vertex v . Then the removal of a vertex u from the MIS could cause the entirety of the common neighborhood of u and v to change their status in v 's data structure. If this common neighborhood is large, then this operation is costly.

To address this issue, our algorithm does not even attempt to determine the exact set of neighbors of v that need to be added to the MIS. Instead, we maintain *partial* information about which vertices will need to be added to the MIS. Then, when we are unsure about whether we need to add a specific vertex to the MIS, we simply add it to the MIS and remove its conflicting neighbors from the MIS, which triggers a chain reaction of changes to the MIS. Despite the fact that this chain reaction may take a long time to resolve, we obtain an amortized time bound by using a potential function: the number of vertices not in the MIS. That is, we ensure that if we spend a lot of time processing an edge update, then the size of the MIS increases substantially as a result.

The core of our algorithm is to carefully choose which vertices to add to the MIS at each step of the chain reaction to ensure that the size of the MIS increases sufficiently. To accomplish this, we store an intricate data structure for each vertex which includes a partition of a subset of its in-neighborhood into numbered buckets. The key idea is to ensure that whenever we remove a vertex from the MIS, it has at least one full bucket of vertices, which we add to the MIS.

When we remove a vertex from the MIS whose top bucket is full of vertices, we begin the chain reaction by adding these vertices to the MIS (and removing the conflicting vertices from the MIS). In each subsequent step of the chain reaction, we process more vertices, and for each processed vertex, we add to the MIS the set of vertices in its *topmost full bucket*. To guarantee that every vertex that we process has at least one full bucket, we utilize an invariant (the “Main Invariant”) which says that for all i , when we process a vertex whose bucket i is full, then in the next iteration of the chain reaction we will only process vertices whose bucket $i - 1$ is full. This implies that if the number of iterations in the chain reaction is at most the number of buckets, then we only process vertices with at least one full bucket. To bound the number of iterations of the chain reaction, we prove that the number of processed vertices *doubles* at every iteration. This way, there cannot be more than a logarithmic number of iterations. Thus, by choosing the number of buckets to be logarithmic, we only process vertices with at least one full bucket, which results in the desired increase in the size of the MIS.

3 Algorithm setup

Let \mathcal{M} be the MIS that we maintain. Our algorithm uses a dynamic edge orientation algorithm as a black box.

For each vertex v , let $N(v)$ denote the neighborhood of v , let $N^+(v)$ denote the out-neighborhood of v , and let $N^-(v)$ denote the in-neighborhood of v .

3.1 The trivial cases

For certain cases of edge updates, there is a simple algorithm to update the \mathcal{M} . Here, we introduce this simple algorithm and then describe the case that this algorithm does not cover.

► **Definition 3.1.** We say that a vertex v is *resolved* if either v is in \mathcal{M} or a vertex in $N^-(v)$ is in \mathcal{M} . Otherwise we say that v is *unresolved*.

The data structure is simply that each vertex v stores a partition of its in-neighborhood into resolved vertices and unresolved vertices. To maintain this data structure, whenever a vertex v enters or exits \mathcal{M} , v notifies its 2-hop out-neighborhood.

DELETE(u, v):

- It cannot be the case that both u and v are in \mathcal{M} since \mathcal{M} is an independent set.
- If neither u nor v is in \mathcal{M} then both must already have neighbors in \mathcal{M} and we do nothing.
- If $u \in \mathcal{M}$ and $v \notin \mathcal{M}$, then we may need to add v to \mathcal{M} . If v is resolved, we do not add v to \mathcal{M} . Otherwise, we scan $N^+(v)$ and if no vertex in $N^+(v)$ is in \mathcal{M} , we add v to \mathcal{M} .

INSERT(u, v):

- If it is not the case that both u and v are in \mathcal{M} , then we do nothing and \mathcal{M} remains maximal.
- If both u and v are in \mathcal{M} we remove v from \mathcal{M} . Now, some of v 's neighbors may need to be added to \mathcal{M} , specifically, those with no neighbors in \mathcal{M} . For each unresolved vertex $w \in N^+(v)$, we scan $N^+(w)$ and if $N^+(w) \cap \mathcal{M} = \emptyset$, then we add w to \mathcal{M} . For each resolved vertex $w \in N^-(v)$, we know not to add w to \mathcal{M} . On the other hand, for each unresolved vertex $w \in N^-(v)$, we do not know whether to add w to \mathcal{M} and it could be costly to scan $N^+(w)$ for all such w . This simple algorithm does not handle the case where v has many unresolved in-neighbors.

In summary, the nontrivial case occurs when we delete a vertex v from \mathcal{M} and v has many unresolved in-neighbors.

3.2 Data structure

As in the trivial cases, each vertex v maintains a partition of $N^-(v)$ into resolved vertices and unresolved vertices. In addition, we further refine the set of unresolved vertices. One important subset of the unresolved vertices in $N^-(v)$ is the *active set of v* , denoted A_v . As motivated in the algorithm overview, A_v is partitioned into b buckets $A_v(1), \dots, A_v(b)$ each of size at most s . We will set b and s so that $b = \Theta(\log n)$ and $s = \Theta(\alpha)$.

The purpose of maintaining A_v is to handle the event that v is removed from \mathcal{M} . When v is removed from \mathcal{M} , we use the partition of A_v into buckets to carefully choose which neighbors of v to add to \mathcal{M} to begin a chain reaction of changes to \mathcal{M} . For the rest of the vertices in A_v , we scan through them and update the data structure to reflect the fact that $v \notin \mathcal{M}$. This scan of A_v is why it is important that each active set is small (size $O(\alpha \log n)$).

One important property of active sets is that each vertex is in the active set of at most one vertex. For each vertex v , let $a(v)$ denote the vertex whose active set contains v . Let $B(v)$ denote the bucket of $A_{a(v)}$ that contains v .

For each vertex v the data structure maintains the following partition of $N^-(v)$:

- Z_v is the set of resolved vertices in $N^-(v)$.
- A_v (*the active set*) is a subset of the unresolved vertices in $N^-(v)$ partitioned into $b = \Theta(\log n)$ buckets $A_v(1), \dots, A_v(b)$ each of size at most $s = \Theta(\alpha)$. A_v is empty if $v \notin \mathcal{M}$.
- P_v (*the passive set*) is the set of unresolved vertices in $N^-(v)$ in the active set of some vertex other than v . P_v is partitioned into b buckets $P_v(1), \dots, P_v(b)$ such that each vertex $u \in P_v$ is in the set $P_v(i)$ if and only if $B(u) = i$.
- R_v (*the residual set*) is the set of unresolved vertices in $N^-(v)$ not in the active set of any vertex.

We note that while Z_v depends only on \mathcal{M} and the orientation of the edges, the other three sets depend on internal choices made by the algorithm. In particular, for each vertex v , the algorithm picks at most one vertex $a(v)$ for which $v \in A_{a(v)}$ and this choice uniquely determines for every vertex $u \in N^+(v)$, which set (A_u , P_u , or R_u) v belongs to.

We now outline the purpose of the passive set and the residual set. Suppose a vertex v is removed from \mathcal{M} . We do not need to worry about the vertices in P_v because we know that all of these vertices are in the active set of a vertex in \mathcal{M} and thus none of them need to be added to \mathcal{M} . On the other hand, we do not know whether the vertices in R_v need to be added to \mathcal{M} . We cannot afford to scan through them all and we cannot risk not adding them since this might cause \mathcal{M} to not be maximal. Thus, we add them all to \mathcal{M} and set off a chain reaction of changes to \mathcal{M} . That is, even though our analysis requires that we carefully choose which vertices of A_v to add to \mathcal{M} during the chain reaction, it suffices to simply add every vertex in R_v to \mathcal{M} (except for those with edges to other vertices we are adding to \mathcal{M}).

3.3 Invariants

We maintain several invariants of the data structure. The invariant most central to the overall argument is the Main Invariant (Invariant 6), whose purpose is outlined in the algorithm overview. The first four invariants follow from the definitions.

- ▶ **Invariant 1.** (Resolved Invariant). *For all resolved vertices v , for all $u \in N^+(v)$, v is in Z_u .*
- ▶ **Invariant 2.** (Orientation Invariant). *For all v , $Z_v \cup A_v \cup P_v \cup R_v = N^-(v)$.*
- ▶ **Invariant 3.** (Empty Active Set Invariant). *For all $v \notin \mathcal{M}$, A_v is empty.*
- ▶ **Invariant 4.** (Consistency Invariant).
 - *If v is resolved, then for all vertices $u \in N^+(v)$, $v \notin A_u$.*
 - *If v is unresolved then v is in A_u for at most one vertex $u \in N^+(v)$.*
 - *If v is in the active set of some vertex u , then for all $w \in N^+(v) \setminus \{u\}$, v is in $P_w(i)$ where i is such that $B(v) = i$.*
 - *If v is in the residual set for some vertex u , then for all $w \in N^+(v)$, v is in R_w .*

The next invariant says that the active set of a vertex is filled from lowest bucket to highest bucket and only then is the residual set filled.

▶ **Definition 3.2.** We say that a bucket $A_v(i)$ is *full* if $|A_v(i)| = s$. We say A_v is *full* if all b of its buckets are full.

▶ **Invariant 5.** (Full Invariant). *For all vertices v and all $i < b$, if $A_v(i)$ is not full then $A_v(i + 1)$ is empty. Also, if A_v is not full then R_v is empty.*

The next invariant, the Main Invariant, says that if we were to move v from $A_{a(v)}$ to the active set of a different vertex u by placing v in the lowest non-full bucket of A_u , then $B(v)$ would not decrease.

▶ **Invariant 6.** (Main Invariant). *For all v , if $B(v) = i > 1$ then for all $u \in N^+(v) \cap \mathcal{M}$, $A_u(i - 1)$ is full.*

4 Algorithm

The algorithm works in four phases.

1. Update \mathcal{M} .
2. Update the data structure.
3. Run a black-box edge orientation algorithm.
4. Update the data structure.

The data structure is completely static during phases 1 and 3.

4.1 Updating \mathcal{M}

When an edge (u, v) is deleted, we run the procedure $\text{DELETE}(u, v)$ specified in the trivial cases section. When an edge (u, v) is inserted and it is not the case that both u and v are in \mathcal{M} , then we do nothing and \mathcal{M} remains maximal. In the case that both u and v are in \mathcal{M} , we need to remove either u or v from \mathcal{M} which may trigger many changes to \mathcal{M} .

The procedure of updating \mathcal{M} happens in two stages. In the first stage, we iteratively build two sets of vertices, S^+ and S^- . Intuitively, S^+ is a subset of vertices that we intend to add to \mathcal{M} and S^- is the set of vertices that we intend to delete from \mathcal{M} . The aforementioned chain reaction of changes to \mathcal{M} is captured in the construction of S^+ and S^- . In the second stage we make changes to \mathcal{M} according to S^+ and S^- . In particular, the set of vertices that we add to \mathcal{M} contains a large subset of S^+ as well as some vertices not in S^+ and the set of vertices that we remove from \mathcal{M} is a subset of S^- . In accordance with our goal of increasing

the size of \mathcal{M} substantially, we ensure that S^+ is much larger than S^- . The most interesting part of updating \mathcal{M} is captured in the first stage. We defer the description of the second stage to the full version [23].

Why is it important to build S^+ before choosing which vertices to add to \mathcal{M} ? The answer is that it is important that we add a *large* subset of S^+ to \mathcal{M} since our goal is to increase the size of \mathcal{M} substantially. We find this large subset of S^+ by finding a large MIS in the graph induced by S^+ , which exists (and can be found in linear time) because the graph has bounded arboricity (see the full version for details [23]). Suppose that instead of iteratively building S^+ , we tried to iteratively add vertices directly to \mathcal{M} in a greedy fashion. This could result in only very few vertices successfully being added to \mathcal{M} . For example, if we begin by adding the center of a star graph to \mathcal{M} and subsequently try to add the leaves of the star, we will not succeed in adding any of the leaves to \mathcal{M} . On the other hand, if we first add the vertices of the star to S^+ then we can find a large MIS in the star (the leaves) to add it to \mathcal{M} .

For the rest of this section we consider an edge insertion (u, v) .

4.1.1 Stage 1: Constructing S^+ and S^-

A key property of the construction is that S^+ is considerably larger than S^- :

► **Lemma 4.1.** *If $|S^-| > 1$ then $|S^+| \geq 4\alpha|S^-|$.*

After constructing S^+ and S^- we will add at least $\frac{|S^+|}{2\alpha}$ vertices to \mathcal{M} and remove at most $|S^-|$ vertices from \mathcal{M} . Thus, Lemma 4.1 implies that \mathcal{M} increases by $\Omega(\frac{|S^+|}{\alpha})$.

To construct S^+ and S^- , we define a recursive procedure $\text{PROCESS}(w)$ which adds at least one full bucket of A_w to S^+ . A key idea in the analysis is to guarantee that for every call to $\text{PROCESS}(w)$, A_w indeed has at least one full bucket.

Algorithm description

We say that a vertex $w \in S^-$ has been *processed* if $\text{PROCESS}(w)$ has been called and otherwise we say that w is *unprocessed*. We maintain a partition of S^- into the processed set and the unprocessed set and we maintain a partition of the set of unprocessed vertices w into two sets based on whether A_w is full or not. We also maintain a queue \mathcal{Q} of vertices to process, which is initially empty. Recall that (u, v) is the inserted edge and both u and v are in \mathcal{M} . The algorithm is as follows.

First, we add v to S^- . Then, if A_v is not full, we terminate the construction of S^+ and S^- . Otherwise, we call $\text{PROCESS}(v)$.

$\text{PROCESS}(w)$:

1. If A_w is full, then add all vertices in $A_w(b) \cup R_w$ to S^+ . If A_w is not full, then let i be the largest full bucket of A_w and add all vertices in $A_w(i)$ and $A_w(i+1)$ to S^+ . We will claim that such an i exists (Lemma 4.2).
2. For all vertices x added to S^+ in this call to PROCESS , we add $N^+(x) \cap \mathcal{M}$ to S^- .
3. If S^- contains an unprocessed vertex x with full A_x , we call $\text{PROCESS}(x)$.

When a call to PROCESS terminates, including the recursive calls, we check whether Lemma 4.1 is satisfied (that is, whether $|S^+| \geq 4\alpha|S^-|$), and if so, we terminate. Otherwise, if \mathcal{Q} is not empty, we let w be the next vertex in \mathcal{Q} and call $\text{PROCESS}(w)$. If \mathcal{Q} is empty we enqueue a new batch of vertices to \mathcal{Q} . This batch consists of the set of all unprocessed vertices in S^- . We will claim that such vertices exist (Lemma 4.2).

► **Remark.** The reason we terminate without calling $\text{PROCESS}(v)$ if A_v is not full (i.e. R_v is empty) is because R_v is the only set for which we cannot afford to determine whether or not each vertex has another neighbor in \mathcal{M} (besides v): We know that each vertex $w \in Z_v \cup P_v$ has another neighbor in \mathcal{M} , and the set A_v is small enough to scan. For the same reason, step 3 of PROCESS is necessary because it ensures that for every vertex w in S^- , all vertices in R_w are in S^+ . If this weren't the case and we removed a vertex w in S^- from \mathcal{M} , we might be left in the “hard case” of needing to deal with R_w .

Lemma 4.1 follows from the algorithm specification: either the algorithm terminates immediately with $S^- = \{v\}$ or the algorithm terminates according to the termination condition, which is that Lemma 4.1 is satisfied.

Several steps in the algorithm (Step 2 of $\text{PROCESS}(w)$ and the last sentence of the algorithm specification) rely on Lemma 4.2:

► **Lemma 4.2.**

1. If we call $\text{PROCESS}(w)$, then A_w has at least one full bucket.
2. Every batch of vertices that we enqueue to \mathcal{Q} is nonempty.

Proof of Lemma 4.2

Let epoch 1 denote the period of time until the first batch of vertices has been enqueued to \mathcal{Q} . For all $i > 1$, let epoch i denote the period of time from the end of epoch $i - 1$ to when the i^{th} batch of vertices has been enqueued to \mathcal{Q} .

To prove Lemma 4.2, we prove a collection of lemmas that together show that (i) Lemma 4.2 holds for all calls to PROCESS before epoch b ends (recall that b is the number of buckets) and (ii) the algorithm terminates before the end of epoch b .

For all i , let p_i and u_i be the number of processed and unprocessed vertices in S^- respectively, when epoch i ends. Let S_i^+ and S_i^- be the sets S^+ and S^- respectively when epoch i ends. Recall that s is the size of a full bucket. Let $s = 8\alpha$ and let $b = \log_2 n + 1$.

► **Lemma 4.3.** For all $1 \leq j \leq b$, every time we call $\text{PROCESS}(w)$ during epoch j , $A_w(b-j+1)$ is full.

Proof. We proceed by induction on j .

Base case. If $j = 1$ then the algorithm only calls $\text{PROCESS}(w)$ on vertices w with full A_w and thus full $A_w(b)$.

Inductive hypothesis. Suppose that during epoch j , all of the processed vertices have full $A_w(b - j + 1)$.

Inductive step. We will show that during epoch $j + 1$, all of the processed vertices have full $A_w(b - j)$. We first note that during $\text{PROCESS}(w)$, the algorithm only adds the vertices in the topmost full bucket of A_w to S^+ . Thus, the inductive hypothesis implies that at the end of epoch j for all vertices $x \in S^+$, $B(x) \geq b - j + 1$.

Then, by the Main Invariant, at the end of epoch j , for all $x \in S^+$ and all $y \in N^+(x) \cap \mathcal{M}$, $A_y(b - j)$ is full. By construction, the only vertices in S^- other than v are those in $N^+(x) \cap \mathcal{M}$ for some $x \in S^+$. Thus, at the end of epoch j , for all vertices $y \in S^-$, $A_y(b - j)$ is full. During epoch $j + 1$, the set of vertices that we process consists only of vertices w that are either in S^- at the end of epoch j or have full A_w . We have shown that all of these vertices w have full $A_w(b - j)$. ◀

► **Lemma 4.4.** For all $1 \leq j \leq b$, $|S_j^+| \geq p_j s$.

Proof. By Lemma 4.3, for all calls to $\text{PROCESS}(w)$ until the end of epoch j , A_w has at least one full bucket. During each call to $\text{PROCESS}(w)$, the algorithm adds at least one full bucket (of size s) of A_w to S^+ . By the Consistency Invariant, (i) every vertex is in the active set of at most one vertex and (ii) if a vertex w appears in the active set of some vertex, then w is not in the residual set of any vertex. The only vertices added to S^+ are those in some active set or some residual set, so every vertex in some active set that is added to S^+ , is added at most once. Thus, for each processed vertex, there are at least s distinct vertices in S^+ . ◀

► **Lemma 4.5.** *For all $1 \leq j \leq b$, $p_j < u_j$. That is, there are more unprocessed vertices than processed vertices.*

Proof. At the end of epoch j , Lemma 4.1 is not satisfied because if it were then the algorithm would have terminated. That is, $|S_j^+| < 4\alpha|S_j^-|$. Combining this with Lemma 4.4 and the fact that $p_j + u_j = |S_j^-|$, we have $p_j s < 4\alpha(p_j + u_j)$. Choosing $s = 8\alpha$ completes the proof. ◀

► **Lemma 4.6.** *For all $1 \leq j \leq b$, $p_j > 2p_{j-1}$. That is, the number of processed vertices more than doubles during each epoch.*

Proof. At the end of epoch $j-1$, we add all unprocessed vertices to \mathcal{Q} . As a result of calling PROCESS on each vertex in \mathcal{Q} , the number of processed vertices increases by u_{j-1} by the end of epoch j . That is, $p_j \geq p_{j-1} + u_{j-1}$. By Lemma 4.5, $p_{j-1} < u_{j-1}$, so $p_j > 2p_{j-1}$. ◀

We apply these lemmas to complete the proof of Lemma 4.2:

1. In epoch 1 we process at least one vertex, so $p_1 \geq 1$. By Lemma 4.6, $p_j > 2p_{j-1}$. Thus, $p_j \geq 2^{j-1}$. If $j = b = \log_2 n + 1$, then $p_j > n$, a contradiction. Thus, the algorithm never reaches the end of epoch b . Then, by Lemma 4.3, every time we call $\text{PROCESS}(w)$, A_w has at least one full bucket.
2. Suppose by way of contradiction that we enqueue no vertices to \mathcal{Q} at the end of some epoch $1 \leq j \leq b$. Then, $u_j = 0$. By Lemma 4.5, $p_j < u_j$, so $p_j < 0$, a contradiction.

5 Analysis

In this section we present the most interesting part of the analysis and defer the rest of the analysis to the full version [23]. When a vertex v is added to \mathcal{M} , A_v is empty and needs to be populated in order to satisfy the Main Invariant. This process is the bottleneck of the runtime and we analyze it here.

We begin by analyzing the runtime of two basic processes that happen while updating the data structure: adding a vertex to some active set (Lemma 5.1) and removing a vertex from some active set (Lemma 5.2).

Recall that our algorithm uses a dynamic edge orientation algorithm as a black box. Let T be the amortized update time of this algorithm and let D be the out-degree of the orientation. Ultimately, we will apply the algorithm of Brodal and Fagerberg [5] to get $D = O(\alpha)$ and $T = O(\alpha + \log n)$.

► **Lemma 5.1.** *Suppose vertex v is not in any active set. Adding v to some active set and updating the data structure accordingly takes time $O(D)$.*

Proof. When we add a vertex v to some $A_w(i)$, for all $w \in N^+(v)$ this could cause a violation to the Consistency Invariant. To remedy this, it suffices to remove v from whichever set it was previously in with respect to w (which is not A_w) and add it to $P_w(i)$. ◀

► **Lemma 5.2.** *Removing a vertex v from some active set A_u and updating the data structure accordingly takes time $O(D \log n)$.*

Proof. When we remove a vertex v from some $A_u(i)$, this leaves bucket $A_u(i)$ not full so the Full Invariant might be violated. To remedy this, we move a vertex, the *replacement vertex*, from a higher bucket or the residual set to $A_u(i)$. That is, the replacement vertex w is chosen to be any arbitrary vertex from $R_u \cup A_u(i+1) \cup \dots \cup A_u(b) \cup P_u(i+1) \cup \dots \cup P_u(b)$. We can choose a vertex from this set in constant time by maintaining a list of all non-empty $P_x(i)$ and $A_x(i)$ for each vertex x . If w is chosen from P_u , we remove w from $A_{a(w)}$ before adding w to $A_u(i)$.

The removal of w from its previous bucket in its previous active set may leave this bucket not full, so again the Full Invariant might be violated and again we remedy this as described above, which sets off a chain reaction. The chain reaction terminates when either there does not exist a viable replacement vertex or until the replacement vertex comes from the residual set. Since the number of the bucket that we choose the replacement vertex from increases at every step of this process, the length of this chain reaction is at most b .³

For each vertex v that we add to an active set, we have already removed v from its previous active set, so Lemma 5.1 applies. Overall, we move at most b vertices to a new bucket and by Lemma 5.1, for each of these b vertices we spend time $O(D)$. Thus, the runtime is $O(bD) = O(D \log n)$. ◀

► **Lemma 5.3.** *The time to update the data structure in response to a violation of the Main Invariant triggered the addition of a single vertex to \mathcal{M} is $O(D\alpha \log^2 n)$.*

Proof. To satisfy the Main Invariant, we need to populate A_v . We fill A_v in order from bucket 1 to bucket b . First, we add the vertices in R_v until either A_v is full or R_v becomes empty. If R_v becomes empty, then we start adding the vertices of $P_v(i)$ in order from $i = b$ to $i = 1$; however, we only add vertex u to A_v if this causes $B(u)$ to decrease. Once we reach a vertex u in P_v where moving u to the lowest numbered non-full bucket of A_v does not cause $B(u)$ to decrease, then we stop populating A_v . Each time we add a vertex u to A_v from P_v , we first remove u from $A_{a(u)}$ and apply Lemma 5.2. Also, each time we add a vertex to A_v , we apply Lemma 5.1. We note that this method of populating A_v is consistent with the Main Invariant.

We add at most $sb = O(\alpha \log n)$ vertices to A_v and for each one we could apply Lemmas 5.2 and 5.1 in succession. Thus, the total time is $O(\alpha D \log^2 n)$. ◀

References

- 1 Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms*, 7(4):567–583, 1986. doi:10.1016/0196-6774(86)90019-2.
- 2 Sepehr Assadi, Krzysztof Onak, Baruch Schieber, and Shay Solomon. Fully dynamic maximal independent set with sublinear update time. In *Proc. 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC*, 2018.

³ We note that the length of the chain reaction can be shortened by choosing the replacement vertex from the highest possible bucket. However, we could be forced to choose from bucket $i+1$ (if all buckets higher than $A_u(i+1)$ or $P_u(i+1)$ are empty).

- 3 Edvin Berglin and Gerth Stølting Brodal. A simple greedy algorithm for dynamic graph orientation. In *Proc. 28th International Symposium on Algorithms and Computation, ISAAC 2017, December 9-12, 2017, Phuket, Thailand*, pages 12:1–12:12, 2017. doi:10.4230/LIPIcs.ISAAC.2017.12.
- 4 Guy E. Blelloch, Jeremy T. Fineman, and Julian Shun. Greedy sequential maximal independent set and matching are parallel on average. In *Proc. 24th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2012, Pittsburgh, PA, USA, June 25-27, 2012*, pages 308–317, 2012. doi:10.1145/2312005.2312058.
- 5 Gerth Stølting Brodal and Rolf Fagerberg. Dynamic representations of sparse graphs. In *Proc. 6th International Workshop on Algorithms and Data Structures WADS*, pages 342–351. Springer-Verlag, 1999.
- 6 Keren Censor-Hillel, Elad Haramaty, and Zohar S. Karnin. Optimal dynamic distributed MIS. In *Proc. ACM Symposium on Principles of Distributed Computing, PODC 2016, Chicago, IL, USA, July 25-28, 2016*, pages 217–226, 2016. doi:10.1145/2933057.2933083.
- 7 Sebastian Daum, Seth Gilbert, Fabian Kuhn, and Calvin C. Newport. Leader election in shared spectrum radio networks. In *Proc. ACM Symposium on Principles of Distributed Computing, PODC 2012, Funchal, Madeira, Portugal, July 16-18, 2012*, pages 215–224, 2012. doi:10.1145/2332432.2332470.
- 8 Manuela Fischer and Andreas Noever. Tight analysis of parallel randomized greedy MIS. In *Proc. 29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2152–2160, 2018. doi:10.1137/1.9781611975031.140.
- 9 Gaurav Goel and Jens Gustedt. Bounded arboricity to determine the local structure of sparse graphs. In *Proc. Graph-Theoretic Concepts in Computer Science, 32nd International Workshop, WG*, pages 159–167, 2006. doi:10.1007/11917496_15.
- 10 Meng He, Ganggui Tang, and Norbert Zeh. Orienting dynamic graphs, with applications to maximal matchings and adjacency queries. In *Proc. Algorithms and Computation - 25th International Symposium, ISAAC 2014, Jeonju, Korea, December 15-17, 2014*, pages 128–140, 2014. doi:10.1007/978-3-319-13075-0_11.
- 11 John E. Hopcroft and Richard M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM J. Comput.*, 2(4):225–231, 1973. doi:10.1137/0202019.
- 12 Tomasz Jurdzinski and Dariusz R. Kowalski. Distributed backbone structure for algorithms in the SINR model of wireless networks. In *Proc. Distributed Computing - 26th International Symposium, DISC 2012, Salvador, Brazil, October 16-18, 2012*, pages 106–120, 2012. doi:10.1007/978-3-642-33651-5_8.
- 13 Tsvi Kopelowitz, Robert Krauthgamer, Ely Porat, and Shay Solomon. Orienting fully dynamic graphs with worst-case time bounds. In *Proc. Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Part II*, pages 532–543, 2014. doi:10.1007/978-3-662-43951-7_45.
- 14 Lukasz Kowalik. Adjacency queries in dynamic sparse graphs. *Inf. Process. Lett.*, 102(5):191–195, 2007. doi:10.1016/j.ip1.2006.12.006.
- 15 Lukasz Kowalik and Maciej Kurowski. Short path queries in planar graphs in constant time. In *Proc. 35th Annual ACM Symposium on Theory of Computing, STOC*, pages 143–148, 2003.
- 16 Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Initializing newly deployed ad hoc and sensor networks. In *Proc. 10th Annual International Conference on Mobile Computing and Networking, MOBICOM 2004, Philadelphia, PA, USA, September 26 - October 1, 2004*, pages 260–274, 2004. doi:10.1145/1023720.1023746.

- 17 Nathan Linial. Distributive graph algorithms-global solutions from local data. In *Proc. 28th Annual Symposium on Foundations of Computer Science, FOCS 1987, Los Angeles, California, USA, 27-29 October 1987*, pages 331–335, 1987. doi:10.1109/SFCS.1987.20.
- 18 Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15(4):1036–1053, 1986. doi:10.1137/0215074.
- 19 Crispin St.J. Nash-Williams. Edge-disjoint spanning trees of finite graphs. *J. London Math. Soc.*, 36(1):445–450, 1961.
- 20 Crispin St.J. Nash-Williams. Decomposition of finite graphs into forests. *J. London Math. Soc.*, 39(1):12, 1964.
- 21 Ofer Neiman and Shay Solomon. Simple deterministic algorithms for fully dynamic maximal matching. In *Proc. Symposium on Theory of Computing Conference, STOC 2013, Palo Alto, CA, USA, June 1-4, 2013*, pages 745–754, 2013. doi:10.1145/2488608.2488703.
- 22 Huy N. Nguyen and Krzysztof Onak. Constant-time approximation algorithms via local improvements. In *Proc. 49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 327–336, 2008. doi:10.1109/FOCS.2008.81.
- 23 Krzysztof Onak, Baruch Schieber, Shay Solomon, and Nicole Wein. Fully dynamic mis in uniformly sparse graphs. *arXiv CoRR*, 2018.
- 24 William T. Tutte. On the problem of decomposing a graph into n connected factors. *J. London Math. Soc.*, 36(1):221–230, 1961.
- 25 Dongxiao Yu, Yuexuan Wang, Qiang-Sheng Hua, and Francis C. M. Lau. Distributed $(\Delta+1)$ -coloring in the physical model. *Theor. Comput. Sci.*, 553:37–56, 2014. doi:10.1016/j.tcs.2014.05.016.