

Fully Dynamic Maximal Independent Set in Expected Poly-Log Update Time

Shiri Chechik
Tel Aviv University
shiri.chechik@gmail.com

Tianyi Zhang
Tsinghua University
tianyi-z16@mails.tsinghua.edu.cn

Abstract—In the fully dynamic maximal independent set (MIS) problem our goal is to maintain an MIS in a given graph G while edges are inserted and deleted from the graph. The first non-trivial algorithm for this problem was presented by Assadi, Onak, Schieber, and Solomon [STOC 2018] who obtained a deterministic fully dynamic MIS with $O(m^{3/4})$ update time. Later, this was independently improved by Du and Zhang and by Gupta and Khan [arXiv 2018] to $\tilde{O}(m^{2/3})$ update time¹. Du and Zhang [arXiv 2018] also presented a randomized algorithm against an oblivious adversary with $\tilde{O}(\sqrt{m})$ update time.

The current state of art is by Assadi, Onak, Schieber, and Solomon [SODA 2019] who obtained randomized algorithms against oblivious adversary with $\tilde{O}(\sqrt{n})$ and $\tilde{O}(m^{1/3})$ update times.

In this paper, we propose a dynamic randomized algorithm against oblivious adversary with expected worst-case update time of $O(\log^4 n)$. As a direct corollary, one can apply the black-box reduction from a recent work by Bernstein, Forster, and Henzinger [SODA 2019] to achieve $O(\log^6 n)$ worst-case update time with high probability. This is the first dynamic MIS algorithm with very fast update time of poly-log.

I. INTRODUCTION

A maximal independent set (MIS) of a given graph $G = (V, E)$ is a subset M of vertices such that M does not contain two neighbor vertices and every vertex in $V \setminus M$ has a neighbor vertex in M . In this paper, we study the maximal independent set (MIS) problem in the dynamic setting, where the graph G is undergoing a sequence of edge insertions and deletions.

MIS is a fundamental problem with both theoretical and practical importance and is used as a fundamental building block in many applications. For instance, MIS has been used for resource scheduling for parallel threads in a multi-core environment, for leader election [7], for resource allocation [13], etc.

The MIS had received a lot of attention in the distributed and parallel settings since the influential works of [1], [10], [11]. It is considered a central

¹As usual n is the number of vertices, m is the number of edges and $\tilde{O}(\cdot)$ suppresses poly-logarithmic factors.

problem in distributed computing and in particular in the symmetry breaking field. Specifically, attaining a better understanding of MIS in the distributed setting is of particular interest not only because it is a fundamental problem but also because other fundamental problems reduce to it.

Censor-Hillel, Haramaty, and Karnin [6] in their pioneering paper studied the problem of maintaining an MIS in the distributed dynamic setting where the graph changes over time. They gave a randomized algorithm for maintaining an MIS against an oblivious adversary in the distributed dynamic setting; as an open question, the authors asked whether it is possible to maintain an MIS in a dynamic graph with update time faster than recomputing everything from scratch, which triggered a recent line of research.

The first non-trivial algorithm was proposed by Assadi, Onak, Schieber and Solomon [2] who presented a deterministic algorithm with $O(m^{3/4})$ amortized update time. This was the first dynamic algorithm that maintains an MIS with sublinear update time in the sequential model. This upper bound was later improved to $\tilde{O}(m^{2/3})$ independently by Du and Zhang [8] and by Gupta and Khan [9]. In the same paper Du and Zhang [8] overcame the $\tilde{O}(m^{2/3})$ barrier by assuming an oblivious adversary and a randomized algorithm with amortized $\tilde{O}(\sqrt{m})$ was proposed. This randomized upper bound was recently improved to $\tilde{O}(\sqrt{n})$ by Assadi *et al.* [3]. For graphs with bounded arboricity α , a deterministic algorithm with amortized update time of $O(\alpha^2 \log^2 n)$ was proposed in [12].

A. Our contribution

In this paper we present the first dynamic MIS algorithm with very fast update time of poly-logarithmic in n . We obtain a randomized dynamic MIS algorithm that works against an oblivious adversary. Moreover, our algorithm can actually maintain a greedy MIS with respect to a random order on the set of vertices; the concept of greedy MIS is defined as follows.

Definition 1. Given any order π on all vertices in V , the greedy MIS M_π with respect to π is uniquely defined by the following procedure that gradually builds an MIS: starting with $M_\pi = \emptyset$, for each vertex in V under order π , if it is not yet dominated by M_π , add it to M_π .

We say that an algorithm has worst-case expected update time α if for every update σ , the expected time to process σ is at most α .

Our main result argues that when π is a uniformly random permutation, the corresponding greedy MIS can be maintained under edge updates against an oblivious adversary, which is formalized in the following statement.

Theorem 2. Let π be a random permutation over V . The greedy MIS on G according to order π can be maintained under edge insertions and deletions in worst-case expected $O(\log^4 n)$ time against an oblivious adversary, where the expectation is taken over the random choice of π .

As a corollary, we can apply a black-box reduction from worst-case time dynamic algorithms to expected worst-case time dynamic algorithms that appeared in a recent paper [5].

Theorem 3 ([5]). Let A be an algorithm that maintains a dynamic data structure D with worst-case expected time α , and let n be a parameter such that the maximum number of items stored in the data structure at any point in time is polynomial in n , and let l be a parameter for the length of the update sequence to be considered. Then there exists an algorithm A' with the following properties.

- 1) For any sequence of updates $\sigma_1, \sigma_2, \dots$, A' processes each update σ_i in $O(\alpha \log^2 n)$ time with high probability.
- 2) A' maintains $\Theta(\log n)$ data structures $D_1, D_2, \dots, D_{\Theta(\log n)}$, as well as a pointer to some D_i that is guaranteed to be correct at the current time. Query operations are answered with D_i .

Corollary 4. There is a dynamic MIS algorithm against an oblivious adversary that handles edge updates in worst-case $O(\log^6 n)$ time with high probability, and answers MIS membership queries in constant time.

Independent work: Independent of our work, Behnezhad *et al.* [4] also presents a fully dynamic algorithm that maintains a greedy MIS with expected poly-logarithmic running time against oblivious adversaries.

B. Technical overview

Our algorithm is a combination of techniques from [6] and [3]. In paper [6], the authors proved a lemma that the expected number of changes made to a greedy MIS by an edge update is bounded by a constant. Unfortunately, they could not achieve an efficient dynamic algorithm since a straightforward implementation of the lemma has a linear dependence on the maximum degree of the graph which could be large.

The issue with the maximum degree was overcome by the algorithm from [3] which relies on what we informally call the *degree reduction* lemma: if we pick a random subset of k vertices and build a greedy MIS on this subset, then the maximum degree of the induced subgraph on all the rest un-dominated vertices is at most $O(\frac{n \log n}{k})$. Therefore we can do the following to achieve an update time with sub-linear dependence on n . First build an MIS on a randomly selected subset of k vertices and then maintain an MIS on the induced subgraph of all the rest vertices in a brute-force manner. If an edge update lies entirely within the induced subgraph, then it takes time proportional to the maximum degree which is $\tilde{O}(n/k)$; if an edge update lies within the random subset, then we rebuild the whole data structure from scratch. The expected running time of this algorithm is a trade-off between two terms. One the one hand, when the edge update occurs within the induced subgraph, the cost would be proportional to the maximum degree which is $\tilde{O}(n/k)$; on the other hand, when the edge update connects two vertices in the random subset, the cost of rebuilding would be $O(m) = O(n^2)$, and under the assumption of obliviousness, the probability that an edge update lies within the random subset is roughly $O(\frac{k^2}{n^2})$, and so the expected time of rebuilding would be $\tilde{O}(n^2 \cdot \frac{k^2}{n^2}) = \tilde{O}(k^2)$. Taking $k = \lfloor n^{1/3} \rfloor$ gives a balance of $\tilde{O}(n^{2/3})$ update time. In their paper [3], the authors further refined the running time to $\tilde{O}(\sqrt{n})$ using a hierarchical approach.

We believe the main bottleneck of the above algorithm is that it takes no effort to utilize the lemma from [6]. As a first attempt one could try to look for expensive parts of [3]'s algorithm and try to plug in [6]'s lemma. For example, instead of directly rebuilding, we could try to apply [6]'s lemma when restoring a greedy MIS among the random subset of k vertices if an edge update occurs between them. However, we would again encounter the large degree issue within the random subset.

Our new algorithm is a direct way of combining [6]'s lemma and the degree reduction lemma. The

algorithm keeps a random ordering $\pi : V \rightarrow [n]$ of all vertices and tries to maintain the random greedy MIS. In order to do so, we explicitly maintain all the induced subgraphs $G_i = (V_i, E_i)$ ($0 \leq i \leq \log n$) on all vertices which are not dominated by MIS vertices from $\pi^{-1}(1), \pi^{-1}(2), \dots, \pi^{-1}(2^i)$. For simplicity assume edge (u, v) is inserted where $2^b < \pi(u) < \pi(v) \leq 2^{b+1}$ for some integer b . Then, on the one hand, this event happens with probability $O(2^{2b}/n^2)$ when π is uniformly random; on the other hand, all changes to the MIS could only take place in G_b whose maximum degree is bounded by $O(\frac{n \log n}{2^b})$.

Let $S \subseteq V_b$ be the set of all influenced vertices (we will formally define what S is later on; basically S contains all vertices that could possibly enter or leave the MIS during this update). Following similar proofs of [6], we could prove the conditional expectation of S is at most $O(n/2^b)$. As the maximum degree of G_b is bounded by $O(\frac{n \log n}{2^b})$, we could go over all neighbors of S in G_b and maintain memberships of vertices from S in subgraphs G_{b+1}, G_{b+2}, \dots , which takes $\tilde{O}(n^2/2^{2b})$ time, perfectly canceling out the probability $O(2^{2b}/n^2)$ we just mentioned. However, this is not the end of the story. Not only could vertices from S change their memberships in subgraphs G_{b+1}, G_{b+2}, \dots , but neighbors of vertices in S as well, which could be as many as $O(n^2/2^{2b})$ in the worst-case. The key to the running time analysis is that π roughly assigns the set S uniform-random positions in $[2^b + 1, n]$ even when S is given as prior knowledge. Therefore, on average, the number of neighbors in G_b of a vertex in S is bounded by $\tilde{O}(1)$.

II. PRELIMINARIES

For any subgraph $H \subseteq G$, let $\Delta(H)$ be its maximum vertex degree. For any $U \subseteq V$, define $\Gamma(U)$ to be the set of all neighbors of U in G , and $G[U]$ the induced subgraph of G on U . For any permutation π on V and vertex $u \in V$, define I_u^π to be the set of neighbor predecessors of u with respect to π . For any two different vertices $u, v \in V$, we say u has a *higher priority* than v if $\pi(u) < \pi(v)$. For any pair of indices i, j , define $\pi[i, j] = \{w \mid i \leq \pi(w) \leq j\}$. The following lemma states the basic characterization of a greedy MIS.

Lemma 5 (folklore). *An MIS M is the greedy MIS with respect to order π if and only if for all $z \in V$, it satisfies the constraint that either $z \in M$ or $I_z^\pi \cap M \neq \emptyset$. For the rest, we will call this constraint the greedy MIS constraint for z .*

The following lemma appeared in [3].

Lemma 6 ([3]). *Let π be a uniformly random permutation on V and let k be an integer in $[n]$. Let U be the set of all vertices not dominated by $M_\pi \cap \pi[1, k]$, then with high probability of $1 - n^{-4}$, $\Delta(G[U]) \leq O(\frac{n \log n}{k})$.*

The next lemma is a slight modification of the previous lemma where we show that even if we fix the position in the permutation of two vertices the lemma still holds.

Lemma 7. *Let $u_1, u_2 \in V$ be two different vertices and $k_1, k_2 \in [n]$ be two different indices, and let $1 \leq k \leq n$ be an integer. Let π be a uniformly random permutation on V under the condition that $\pi(u_i) = k_i, i \in \{1, 2\}$. Let U be the set of all vertices not dominated by $M_\pi \cap \pi[1, k]$, then with high probability $1 - n^{-2}$, $\Delta(G[U]) \leq O(\frac{n \log n}{k})$.*

Proof: Call a permutation π *bad* if $\Delta(G[U]) \geq \Omega(\frac{n \log n}{k})$. Noticing that $\Pr_\pi[\pi(u_i) = k_i, i \in \{1, 2\}] = \frac{1}{n(n-1)/2}$, by Lemma 6 we have:

$$\begin{aligned} n^{-4} &\geq \Pr_\pi[\pi \text{ is bad}] \\ &= \frac{1}{n(n-1)/2} \Pr_\pi[\pi \text{ is bad} \mid \forall i, \pi(u_i) = k_i] \\ &\quad + (1 - \frac{1}{n(n-1)/2}) \Pr_\pi[\pi \text{ is bad} \mid \exists i, \pi(u_i) \neq k_i] \\ &\geq \frac{1}{n(n-1)/2} \Pr_\pi[\pi \text{ is bad} \mid \pi(u_i) = k_i] \end{aligned}$$

Hence, $\Pr_\pi[\pi \text{ is bad} \mid \forall i, \pi(u_i) = k_i]$ is at most n^{-2} as well, which concludes the proof. \blacksquare

For the rest of this section, we review the notion of *influenced set* which was studied in [6]. Given a total order π , an MIS $M = M_\pi$, as well as an edge update between u, v , we turn to define v 's influenced set S_v^π . If v does not violate the greedy MIS constraint after the edge update, then define $S_v^\pi = \emptyset$; notice that v always preserves the greedy MIS constraint if $\pi(v) < \pi(u)$. Otherwise, initially set $S_0 = \{v\}$. For any $i \geq 1$, define S_i to be the set of all non-MIS vertices whose MIS predecessors are all in S_{i-1} , plus the set of every MIS vertex who has at least one predecessor in S_{i-1} , namely:

$$\begin{aligned} S_i &= \{w \mid w \in M, S_{i-1} \cap I_w^\pi \neq \emptyset\} \\ &\cup \{w \mid w \notin M, I_w^\pi \cap M \subseteq \bigcup_{j=0}^{i-1} S_j\} \end{aligned}$$

Note that the set M refers to the greedy MIS in the old graph, not in the new graph. Eventually, define v 's influenced set to be $S_v^\pi = \bigcup_{i=0}^{\infty} S_i$. When $S_v^\pi \neq \emptyset$, there is a simple characterization which will be used later.

Lemma 8. Let M be the greedy MIS in the old graph. When $S_v^\pi \neq \emptyset$, it is equal to the smallest set S that contains v and satisfies the following two conditions. (1) $\forall z \in M, I_z^\pi \cap S \neq \emptyset$ iff $z \in S$. (2) $\forall z \notin M, I_z^\pi \cap M \subseteq S$ iff $z \in S$.

Proof: Since S_v^π satisfies both of (1) and (2), it suffices to prove that any S containing v that satisfies both (1) and (2) would contain S_v^π as a subset. This is done by an easy induction on $i \geq 0$ that S contains every S_i . ■

The proofs of the following two lemmas are given for completeness in the appendix.

Lemma 9 ([6]). Let π, σ be two permutations, $S \subseteq V$ a nonempty set, and $v \in V$ be an arbitrary vertex. Suppose an edge update occurs between u, v . Assume $S_v^\pi = S$, $\pi(u) < \pi(v)$, $\sigma(u) < \sigma(v)$, σ, π have the same induced relative order on both S and $V \setminus S$, namely $\pi_S = \sigma_S, \pi_{V \setminus S} = \sigma_{V \setminus S}$, then $M_\pi = M_\sigma$ in the old graph before the edge update, and $S_v^\sigma = S$.

Lemma 10 ([6]). Let π, σ be two permutations, $S \subseteq V$ a vertex subset, and $v \in V$ be an arbitrary vertex. Suppose an edge update occurs between u, v . If $S_v^\pi = S \neq \emptyset$, and $\pi(u) < \pi(v)$, $\sigma(u) < \sigma(v)$, σ, π have the same induced relative order on both $S \setminus \{v\}$ and $V \setminus S$, namely $\pi_{S \setminus \{v\}} = \sigma_{S \setminus \{v\}}, \pi_{V \setminus S} = \sigma_{V \setminus S}$. If $v \neq \arg \min_{z \in S} \{\sigma(z)\}$ then $S_v^\sigma = \emptyset$.

It was also shown in [6] that for an edge update (u, v) the expected size of S_v^π is constant. In our algorithm we need the following different variants of this claim; the proofs are deferred to the appendix.

Lemma 11. Suppose an edge update occurs between u, v . Let $1 \leq A < B \leq n$ be two integers. Then

$$\mathbb{E}_\pi[|S_v^\pi| \mid \pi(u) = A, \pi(v) \in [A + 1, B]] < \frac{n}{B - A}$$

Lemma 12. Suppose an edge update occurs between u, v . Let $1 \leq A < B \leq n$ be two integers. Then

$$\mathbb{E}_\pi[|S_v^\pi| \mid A < \pi(u) < \pi(v) \leq B] < \frac{2n}{B - A}$$

III. THE MAIN ALGORITHM

In this section we describe our fully dynamic MIS algorithm.

A. Data structure

When π is a fixed permutation over V , our algorithm is entirely deterministic. Let $M \subseteq V$ be the greedy MIS with respect to π , and for any $1 \leq k \leq n$, define $M_k = M \cap \pi[1, k]$. Since M is defined in a greedy manner, M_k dominates the entire set $\pi[1, k]$.

The algorithm explicitly maintains the induced subgraph $G_i = (V_i, E_i), \forall 0 \leq i \leq \log n - 1$, where $V_i = V \setminus (M_{2^i} \cup \Gamma(M_{2^i}))$; by definition $G_0 \supseteq G_1 \supseteq G_2 \supseteq \dots \supseteq G_{\log n - 1}$. More precisely, given a permutation π , our algorithm maintains at any given point of time the graphs G_i for $0 \leq i \leq \log n - 1$ and the greedy MIS M_π . In the following subsection we describe our update algorithm to maintain both the graphs G_i and the MIS M_π .

B. Update algorithm

Suppose an edge is updated, either inserted or deleted, between $u, v \in V$ with $\pi(u) < \pi(v)$. Suppose $2^a < \pi(u) \leq 2^{a+1}$ and $2^b < \pi(v) \leq 2^{b+1}$ for integers a and b . There are several *easy cases*, where $S_v^\pi = \emptyset$ and thus we do not need to make changes to M as M stays the greedy MIS with respect to π , and we only need to maintain the subgraphs $G_0, G_1, \dots, G_{\log n - 1}$.

- (i) $u \notin M$. In this case, we simply add or remove, depending whether the edge update is an insertion or deletion, the edge (u, v) to/from E_0, E_1, \dots, E_i , where i is the largest index such that $u, v \in V_i$.
- (ii) $u \in M, v \notin M$, the update is a deletion and $I_v^\pi \cap M \neq \{u\}$. This case can be handled in the same way as in (1): remove the edge (u, v) in E_0, E_1, \dots, E_i , where i is the largest index such that $u, v \in V_i$, and recompute v 's position in the subgraphs $G_a, G_{a+1}, \dots, G_{\log n - 1}$.
- (iii) $u \in M, v \notin M$ and the update is an insertion. In this case, if $v \in V_a$, then since now v is dominated by $u \in V_a$ we should remove v from all subgraphs $G_k, \forall k > a$. After that, add (u, v) to E_0, E_1, \dots, E_i , where i is the largest index such that $u, v \in V_i$.

For the rest of this section we consider the case where an edge is inserted between $u, v \in M$, or deleted between $u \in M, v \notin M$ with $I_v^\pi \cap M = \{u\}$. In both of these cases, $S_v^\pi \neq \emptyset$ and thus we need to change v 's status in the MIS, and then we must try to fix the greedy MIS M within G_b . We start by computing the nonempty influenced set S_v^π with respect to edge update between u, v .

- (1) Initialize an output set $S = \emptyset$ that is promised to be equal to S_v^π by the end of the algorithm, as well as a set $T = \{v\}$ that contains a set of candidate vertices that might be included in S during the process.
- (2) In each iteration, extract $z = \arg \min_{z \in T} \{\pi(z)\}$ from T . If $z \in M$, then suppose $2^k < \pi(z) \leq 2^{k+1}$; by definition it must be $z \in V_k$. First we add z to S , and scan all neighbors w of z in V_k such that $\pi(w) > \pi(z)$ and add w to T .

If $z \notin M$, first scan its adjacency list in G_b ; if all its MIS neighbors with higher priority are in S , then add z to S and add all of its MIS neighbors $w \in V_b$ with $\pi(w) > \pi(z)$ to T .

(3) When T becomes empty, output S as S_v^π .

For convenience we summarize the above procedure as pseudo-code 1.

Algorithm 1: FindInfluencedSet(u, v, b)

```

1  $S \leftarrow \emptyset$ , in easy cases (i)(ii)(iii)  $T \leftarrow \emptyset$ , and
   otherwise  $T \leftarrow \{v\}$ ;
2 while  $T \neq \emptyset$  do
3    $z \leftarrow \arg \min_{z \in T} \{\pi(z)\}$ ,  $T \leftarrow T \setminus \{z\}$ ;
4   if  $z \in M$  then
5      $S \leftarrow S \cup \{z\}$ ;
6     suppose  $2^k < \pi(z) \leq 2^{k+1}$ , and assert
        $z \in V_k$ ;
7     for neighbor  $w \in V_k$  of  $z$  such that
        $\pi(w) > \pi(z)$  do
8        $T \leftarrow T \cup \{w\}$ ;
9   else
10    flag  $\leftarrow$  true;
11    for neighbor  $w \in V_b \cap M$  of  $z$  such that
        $\pi(w) < \pi(z)$  do
12      if  $w \notin S$  then
13         $\text{flag} \leftarrow$  false and break;
14      if flag then
15         $S \leftarrow S \cup \{z\}$ ;
16        for neighbor  $w \in V_b \cap M$  of  $z$  with
            $\pi(w) > \pi(z)$  do
17           $T \leftarrow T \cup \{w\}$ ;
18 return  $S$ ;
```

It will be proved that the output S of Algorithm 1 is equal to S_v^π . Once we have found $S = S_v^\pi$, we can try to fix the greedy MIS by only looking at $G[S]$; note that it might be the case that not every vertex in S needs to change its status in the MIS (for example if $G[S]$ is a triangle and v is removed from M due to an insertion, we would not add both vertices in S to M). If the edge update is an insertion, we first remove v from all $V_k, k > a$, and then compute the greedy MIS on $G[S \setminus \{v\}]$ with respect to π ; if the edge update is a deletion, we add v to all $V_k, \forall a < k \leq b$, and then compute the greedy MIS on $G[S]$ with respect to π .

Last but not least, we also need to update $G_k, k \geq b + 1$. This is done in the straightforward manner:

go over every vertex z that has changed its status in MIS in the increasing order with respect to $\pi(z)$. Assuming $2^k < \pi(z) \leq 2^{k+1}$, directly go over all of its neighbors in G_k and recompute their memberships in $G_{b+1}, \dots, G_{\log n - 1}$. More specifically, consider the following two cases.

- (1) If z has been added to M , then for every neighbor $w \in \Gamma(z) \cap V_k$, we remove w from all $G_l, l > k$.
- (2) If z has been removed from M , then z belonged to V_k before the update. Instead of enumerating every neighbor from the current version of $\Gamma(z) \cap V_k$, we go over all of its old neighbors $w \in V_k$ before the update, and compute their memberships in $G_{b+1}, \dots, G_{\log n - 1}$.

We also summarize this procedure as pseudo-code 2. After that we can summarize the main update algorithm as pseudo-code 3.

Algorithm 2: FixSubgraphs(S, b)

```

1 for  $z \in S$  that has changed its status, in the
   increasing order in terms of  $\pi$  do
2   assume  $2^k < \pi(z) \leq 2^{k+1}$ ;
3   if  $z$  has joined  $M$  then
4     for  $w \in V_k \cap \Gamma(z)$  do
5       remove  $w$  from all  $G_l, l > k$ ;
6   else if  $z$  has left  $M$  then
7     for neighbor  $w$  of  $z$  in the old version of  $V_k$ 
       before the edge update do
8       compute  $w$ 's memberships in
          $G_k, G_{k+1}, \dots, G_{\log n - 1}$ ;
```

C. Correctness

In this section we prove the correctness of our algorithm. We start by proving that the algorithm correctly computed the set S_v^π .

Lemma 13. *Algorithm 1 correctly outputs the influenced set with respect to v , namely $S = S_v^\pi$ when it terminates.*

Proof: Let $v = z_1, z_2, \dots, z_l$ be the sequence of vertices that are added to S sorted in the increasing order with respect to π . We prove inductively that for any $1 \leq i \leq l$, $S_v^\pi \cap \pi[\pi(z_1), \pi(z_i)] = \{z_1, z_2, \dots, z_i\}$. When $i = 1$, the equality holds trivially as $S_v^\pi \cap \pi[\pi(z_1), \pi(z_i)] = \{v\} = \{z_1\}$. For the inductive step, suppose we have $S_v^\pi \cap \pi[\pi(z_1), \pi(z_i)] =$

Algorithm 3: Update(u, v)

```
1 suppose  $\pi(u) < \pi(v)$ , and  $2^a < \pi(u) \leq 2^{a+1}$ ,  
    $2^b < \pi(v) \leq 2^{b+1}$ ;  
2  $S \leftarrow \text{FindInfluencedSet}(u, v, b)$ ;  
3 if  $S = \emptyset$  then  
4   recompute  $v$ 's memberships among  
    $G_a, G_{a+1}, \dots, G_{\log n-1}$ ;  
5 else  
6   if the update is insertion then  
7     remove  $v$  from  $V_k, k > a$ ;  
8     run the greedy MIS algorithm on  
      $G[S \setminus \{v\}]$  with respect to order  $\pi$ ;  
9   else  
10    add  $v$  to all  $V_k, a < k \leq b$ ;  
11    run the greedy MIS algorithm on  $G[S]$   
    with respect to order  $\pi$ ;  
12 FixSubgraphs( $S, b$ );
```

$\{z_1, z_2, \dots, z_i\}$ for some $i \geq 1$. Next we prove $S_v^\pi \cap \pi[\pi(z_i) + 1, \pi(z_{i+1})] = \{z_{i+1}\}$ in two steps.

- $z_{i+1} \in S_v^\pi$.
This can be verified according to the specification of Algorithm 1 and definition of S_v^π in the following way. If z_{i+1} were added to S on line-5, namely $z_{i+1} \in M$, then it must have been introduced to T on line-17 by a neighboring by one of its neighbor that appears before in π . Note that this predecessor cannot in M , and so it was added to S on line-15, and thus z_{i+1} was added to T on line-17. Then according to the definition of S_v^π , $z_{i+1} \in S_v^\pi$. If otherwise z_{i+1} was added to S as a non-MIS vertex, then on the one hand z_{i+1} does not have MIS predecessor neighbors not in V_b as $z_{i+1} \in V_b$; on the other hand, z_{i+1} can be added to S only when all of MIS its neighboring predecessors belong to $\{z_1, z_2, \dots, z_i\} \subseteq S_v^\pi$. Therefore, according to the definition of S_v^π , it should be $z_{i+1} \in S_v^\pi$.
- For any $w \in \pi[\pi(z_i) + 1, \pi(z_{i+1}) - 1]$, $w \notin S_v^\pi$.
Suppose we choose $w \in S_v^\pi \cap \pi[\pi(z_i) + 1, \pi(z_{i+1}) - 1]$ with the smallest order in π . We first rule out the case where $w \in M$. If this should be the case, the w must be adjacent to a vertex $z \in \{z_1, z_2, \dots, z_i\}$; this is not possible because w would have been added to T , when z was added to S on line-15, and then later it would be added to S .
Now we suppose $w \notin M$. By definition of S_v^π

and the inductive hypothesis, all preceding MIS neighbors of w belong to $\{z_1, z_2, \dots, z_i\}$. Let $z \in \{z_1, z_2, \dots, z_i\}$ be the one with the smallest order among its MIS neighbors, and suppose $2^k < \pi(z) < 2^{k+1}$. Since z is the MIS vertex that dominates w with the smallest order, it must be $w \in V_k$, and therefore when z was added to S on line-5, w would be added to T on line-8, and later to S on line-15, which is a contradiction. ■

Lemma 14. *Algorithm 3 correctly restores the greedy MIS with respect to π .*

Proof: We only need to consider the case when $S = S_v^\pi \neq \emptyset$ since otherwise no changes are made to the greedy MIS. We first claim that none of the vertices outside S need to change their status in the greedy MIS. This is because, on the one hand, for any $z \in M \setminus S$, by Lemma 8 we know $I_z^\pi \cap S = \emptyset$, and so any changes within S cannot affect z ; on the other hand, for any $z \in V \setminus (M \cup S)$, by Lemma 8, there exists a vertex from $M \setminus S$ that dominates z as a predecessor, and therefore z stays a non-MIS vertex, irrespective of changes in S .

Secondly, we claim that recomputing the greedy MIS on $G[S \setminus \{v\}]$ or $G[S]$, depending on whether the update is an insertion or a deletion, has no conflicts with MIS vertices in $M \setminus S$. This is because, again by Lemma 8, for any $z \in S$ that was originally a non-MIS vertex, z is not adjacent to any MIS vertex from $M \setminus S$, and so adding z to M has no conflicts with vertices in $M \setminus S$. ■

Lemma 15. *In each iteration of the outermost loop of Algorithm 2, by the time when line-2 is executed, V_k is already equal to $V \setminus (M_{2^k} \cup \Gamma(M_{2^k}))$.*

Proof: We prove the claim by an induction on the value of $\pi(z)$. For the base case where $z = v, k = b$, note that the only possible change to V_b is v : if the edge update is an insertion, then v would leave V_b ; if the edge update is a deletion, then v might join V_b . In both cases, we have already fixed it right before recomputing the greedy MIS on $G[S \setminus \{v\}]$ or $G[S]$. Since we turn to fix subgraphs $G_b, G_{b+1}, \dots, G_{\log n-1}$ after we have finished restoring the greedy MIS, it should be $V_b = V \setminus (M_{2^b} \cup \Gamma(M_{2^b}))$ at the beginning of Algorithm 2.

Next we turn to look at the inductive step. We first argue that any vertex w that leaves $V \setminus (M_{2^k} \cup \Gamma(M_{2^k}))$ due to changes in S has already been removed from V_k in previous iterations. This is because we iterate over S in the increasing order with respect to π , and we must have already visited another vertex $z' \in S \cap M$

with $2^l < \pi(z') \leq 2^{l+1} \leq 2^k$ who is the earliest neighbor of w . By the inductive hypothesis, when z' was enumerated in the for-loop, $V_l = V \setminus (M_{2^l} \cup \Gamma(M_{2^l}))$, and thus w is removed from V_k by then.

Secondly we argue that any vertex w that joins $V \setminus (M_{2^k} \cup \Gamma(M_{2^k}))$ due to changes in S has already been added to V_k in previous iterations. For w to join $V \setminus (M_{2^k} \cup \Gamma(M_{2^k}))$, it must have lost all of its MIS neighbors whose order is less or equal to 2^k . Let $z' \in S \setminus M$ be the one with smallest order and assume $2^l < \pi(z') \leq 2^{l+1} \leq 2^k$, and so z' must have been removed from M by Algorithm 3. By the inductive hypothesis, by the time when z' was enumerated by Algorithm 2, we fix all old neighbors of z' in V_l , which include w , and hence w 's memberships in $G_l, G_{l+1}, \dots, G_{\log n-1}$ were already recomputed from scratch by then. ■

Corollary 16. *The update algorithm correctly maintains subgraphs $G_0, G_1, \dots, G_{\log n-1}$ by the end of its execution.*

D. Running time analysis

Define \mathcal{B} to be the set of all permutations π such that there exists an index $0 \leq k \leq \log n - 1$ for which $\Delta(G_k) \geq \Omega(n \log n / 2^k)$ either before or after the edge update; we need to emphasize it here that the constant hidden in the $\Omega(\cdot)$ notation is larger than the constant hidden in the notation $O(\cdot)$ in the statement of Lemma 7.

Lemma 17. *Let a, b be fixed integers. Denote $\mathcal{E} = \{\pi(u) < \pi(v), \pi(u) \in [2^a + 1, 2^{a+1}], \pi(v) \in [2^b + 1, 2^{b+1}]\}$. Let T_0 be the set of all vertices that have once belonged to T , and let T_1 be the set of all vertices that need to change their memberships among subgraphs $G_{b+1}, \dots, G_{\log n-1}$ during the execution of Algorithm 2. Note that in the easy cases where $S_v^\pi = \emptyset$, we have $T = T_0 = T_1 = \emptyset$. Then we have the following bound on the conditional expectation:*

$$\mathbb{E}_\pi[|T_0 \cup T_1| \mid \mathcal{E}] = O(n \log^2 n / 2^b + n^2 \cdot \Pr[\pi \in \mathcal{B} \mid \mathcal{E}])$$

We break the proof of the above lemma into several steps.

Lemma 18. $\mathbb{E}_\pi[|S_v^\pi| \mid \mathcal{E}] = O(n/2^b)$.

Proof: If $a < b$, then u belongs to $\pi[1, 2^b]$. Directly apply Lemma 11 by fixing an arbitrary position $\pi(u) \in [2^a + 1, 2^{a+1}]$ and setting $A = \pi(u), B = 2^{b+1}$, and then we would have $\mathbb{E}_\pi[|S_v^\pi| \mid \mathcal{E}] \leq n/(2^{b+1} - A) \leq n/2^b$. If $a = b$, then $u, v \in \pi[2^b + 1, 2^{b+1}]$. Apply

Lemma 12 with $A = 2^b, B = 2^{b+1}$, and then we also have $\mathbb{E}_\pi[|S_v^\pi| \mid \mathcal{E}] \leq n/2^{b-1}$. ■

Fix any set S such that $v \in S \subseteq V$, as well as any relative order σ_+ on S and any relative order σ_- on $V \setminus S$, such that there exists a permutation π with $S_v^\pi = S, \pi_S = \sigma_+, \pi_{V \setminus S} = \sigma_-$. Therefore, we can further decompose the conditional expectations as follows:

$$\begin{aligned} & \mathbb{E}_\pi[|T_0 \cup T_1| \mid \mathcal{E}] \\ &= \sum_{S, \sigma_+, \sigma_-} \Pr[S_v^\pi = S, \pi_S = \sigma_+, \pi_{V \setminus S} = \sigma_- \mid \mathcal{E}] \\ & \cdot \mathbb{E}_\pi[|T_0 \cup T_1| \mid \mathcal{E}, S_v^\pi = S, \pi_S = \sigma_+, \pi_{V \setminus S} = \sigma_-] \end{aligned}$$

Therefore, it suffices to study the upper bound on $\mathbb{E}_\pi[|T_0 \cup T_1| \mid \mathcal{E}, S_v^\pi = S, \pi_S = \sigma_+, \pi_{V \setminus S} = \sigma_-]$. For notational convenience, define $\Omega = \{\pi \mid \mathcal{E}, \pi_S = \sigma_+, \pi_{V \setminus S} = \sigma_-\}$. By Lemma 9, if there exists one $\pi \in \Omega$ such that $S_v^\pi = S$, then all $\pi \in \Omega$ would satisfy $S_v^\pi = S$; plus $\forall \pi \in \Omega$, all M_π 's are the same in the old graph before the edge update, which we can safely denote as a common MIS M .

First we study the conditional expectation $\mathbb{E}_\pi[|T_0| \mid \pi \in \Omega]$. As can be seen from Algorithm 1, any vertex, which belonged to M before the edge update, that has once been added to T must have eventually joined S . So we only need to bound the total number of vertices in $T_0 \setminus M$. Again by Algorithm 1, any $w \in T_0 \setminus M$ was added to T by an MIS predecessor $z \in S$ on line-8. Therefore, $|T_0 \setminus M|$ is bounded by the sum of (lower priority) neighbors of all $z \in S \cap M$. So it suffices to study individual contribution of all $z \in S \cap M$ to $T_0 \setminus M$. Formally, $\forall z \in S \setminus M, w \in T_0 \setminus M$, we say z contributes w to T_0 if w was added to T on line-8 when z is being processed. First we notice a basic property of T_0 .

Lemma 19. $v = \arg \min_{z \in T_0} \{\pi(z)\}$, for any $\pi \in \Omega$.

Proof: This property is directly guaranteed by Algorithm 1: on line-8 or line-17, it only adds vertices w to T whose order is strictly larger than z who has just entered S . Since v is the first vertex that has been added to S , all vertices that join T should have larger order than v . ■

Lemma 20. For any $k > b$, $\mathbb{E}_{\pi \in \Omega}[|(S \setminus \{v\}) \cap \pi[2^b + 1, 2^k]|] < \frac{2^k |S|}{n}$.

Proof: Decompose the expectation as following:

$$\begin{aligned} \mathbb{E}_{\pi \in \Omega}[|(S \setminus \{v\}) \cap \pi[2^b + 1, 2^k]|] &= \sum_{j=2^b+1}^{2^{b+1}} \Pr_{\pi \in \Omega}[\pi(v) = j] \\ & \cdot \mathbb{E}_{\pi \in \Omega}[|(S \setminus \{v\}) \cap \pi[2^b + 1, 2^k]| \mid \pi(v) = j] \end{aligned}$$

When $\pi(v) = j$, the rest of $S \setminus \{v\}$ are free to choose positions on $[j+1, n]$, as v always takes the smallest order among S , which is guaranteed by Lemma 19 as $S \subseteq T_0$. Hence, for any $l \in [1, \min\{2^k - j, |S| - 1\}]$, conditioned on $\pi(v) = j$, the probability that $|(S \setminus \{v\}) \cap \pi[2^b + 1, 2^k]| = l$ is equal to $\frac{\binom{2^k - j}{l} \cdot \binom{n - 2^k}{|S| - 1 - l}}{\binom{n - j}{|S| - 1}}$. Therefore, the expectation is computed as follows:

$$\begin{aligned}
& \mathbb{E}_{\pi \in \Omega} [|(S \setminus \{v\}) \cap \pi[2^b + 1, 2^k]| \mid \pi(v) = j] \\
&= \sum_{l=1}^{\min\{2^k - j, |S| - 1\}} l \cdot \Pr_{\pi \in \Omega} [|(S \setminus \{v\}) \cap \pi[2^b + 1, 2^k]| = l \mid \pi(v) = j] \\
&= \sum_{l=1}^{\min\{2^k - j, |S| - 1\}} l \cdot \frac{\binom{2^k - j}{l} \cdot \binom{n - 2^k}{|S| - 1 - l}}{\binom{n - j}{|S| - 1}} \\
&= \sum_{l=1}^{\min\{2^k - j, |S| - 1\}} (2^k - j) \cdot \binom{2^k - j - 1}{l - 1} \\
&\quad \cdot \binom{n - 2^k}{|S| - 1 - l} / \binom{n - j}{|S| - 1} \\
&= (2^k - j) \cdot \binom{n - j - 1}{|S| - 2} / \binom{n - j}{|S| - 1} \\
&= \frac{(2^k - j)(|S| - 1)}{n - j} < \frac{2^k |S|}{n}
\end{aligned}$$

Finally, we have:

$$\begin{aligned}
& \mathbb{E}_{\pi \in \Omega} [|(S \setminus \{v\}) \cap \pi[2^b + 1, 2^k]|] \\
&= \sum_{j=2^{b+1}}^{2^{b+1}} \Pr_{\pi \in \Omega} [\pi(v) = j] \\
&\quad \cdot \mathbb{E}_{\pi \in \Omega} [|(S \setminus \{v\}) \cap \pi[2^b + 1, 2^k]| \mid \pi(v) = j] \\
&< \sum_{j=2^{b+1}}^{2^{b+1}} \Pr_{\pi \in \Omega} [\pi(v) = j] \cdot \frac{2^k |S|}{n} = \frac{2^k |S|}{n}
\end{aligned}$$

Lemma 21. *The expected contribution of all $z \in S \cap M \setminus \{v\}$ to T_0 is at most $O(|S| \log^2 n + |S|n \cdot \Pr_{\pi \in \Omega} [\pi \in \mathcal{B}])$.*

Proof: Consider any index $b \leq k \leq \log n - 1$. When $2^k < \pi(z) \leq 2^{k+1}$, the total number of neighbors of z in V_k is at most $O(n \log n / 2^k + n \cdot \mathbb{1}[\pi \in \mathcal{B}])$, by definition of \mathcal{B} . Therefore, by Lemma 20 the expected total contribution of $z \in S \cap M \setminus \{v\}$ to T_0 that lies in $[2^k + 1, 2^{k+1}]$ is bounded by $O(|S| \log n + 2^k |S| \cdot \mathbb{1}[\pi \in \mathcal{B}])$. Taking a summation over all k we can finalize the proof. ■

By Lemma 21, we have an upper bound on conditional expectation:

$$\begin{aligned}
& \mathbb{E}_{\pi} [|T_0| \mid \pi \in \Omega] \\
&\leq O(|S| \log^2 n + n \log n / 2^b + |S|n \cdot \Pr_{\pi \in \Omega} [\pi \in \mathcal{B}])
\end{aligned}$$

Here we have an extra additive term as an upper bound on the contribution of v to T_0 .

Next we try to study $\mathbb{E}_{\pi} [|T_1| \mid \pi \in \Omega]$. By Algorithm 2, for any $z \in S$ that has changed its status in M , we go over some of the neighbors of z and update their memberships in $G_{k+1}, \dots, G_{\log n - 1}$ using brute force, and by definition these neighbors would belong to T_1 . Similar to what we did with T_0 , we say z *contributes* these neighbors to T_1 . Next we need to carefully analyze the total number of these neighbors.

Lemma 22. *The expected contribution of all $z \in S \setminus \{v\}$ to T_1 is at most $O(|S| \log^2 n + |S|n \cdot \Pr_{\pi \in \Omega} [\pi \in \mathcal{B}])$.*

Proof: Let $k \in [b, \log n - 1]$ be any index. Assume $2^k < \pi(z) \leq 2^{k+1}$. Consider the following two possibilities.

- z has joined M during the update algorithm. In this case, z must belong to V_k and thus the total number of its neighbors in $V \setminus (M_{2^k} \cup \Gamma(M_{2^k}))$ is at most $O(n \log n / 2^k + n \cdot \mathbb{1}[\pi \in \mathcal{B}])$, and by Lemma 15 we know $V_k = V \setminus (M_{2^k} \cup \Gamma(M_{2^k}))$ by the time z is processed by Algorithm 2, and thus the total number of its neighbors in V_k is at most $O(n \log n / 2^k)$.
- z has just left M during the update algorithm. In this case, z was selected by M and thus belonged to V_k before the update. As Algorithm 2 only iterates over z 's old neighbors in V_k , the total number of these neighbors is also bounded by $O(n \log n / 2^k + n \cdot \mathbb{1}[\pi \in \mathcal{B}])$.

In any case, the contribution of z to T_1 is at most $O(n \log n / 2^k + n \cdot \mathbb{1}[\pi \in \mathcal{B}])$. Therefore, by Lemma 20 the expected total contribution of $z \in S \cap M \setminus \{v\}$ to T_1 that lies in $[2^k + 1, 2^{k+1}]$ is bounded by $O(|S| \log n + 2^k |S| \cdot \mathbb{1}[\pi \in \mathcal{B}])$. Taking a summation over all k we can finalize the proof. ■

Taking a summation over all $z \in S \setminus \{v\}$ that has changed its status in the MIS we have:

$$\begin{aligned}
& \mathbb{E}_{\pi} [|T_1| \mid \pi \in \Omega] \\
&\leq O(|S| \log^2 n + n \log n / 2^b + |S|n \cdot \Pr_{\pi \in \Omega} [\pi \in \mathcal{B}])
\end{aligned}$$

Here the extra additive term also stands for v 's contribution to T_1 .

Proof of Lemma 17: To summarize, by Lemma 21 and Lemma 22, we have proved:

$$\begin{aligned} & \mathbb{E}_\pi[|T_0 \cup T_1| \mid \pi \in \Omega] \\ & \leq O(|S| \log^2 n + n \log n / 2^b + |S|n \cdot \Pr_{\pi \in \Omega}[\pi \in \mathcal{B}]) \end{aligned}$$

Recall a previous decomposition we would then have:

$$\begin{aligned} & \mathbb{E}_\pi[|T_0 \cup T_1| \mid \mathcal{E}] \\ & = \sum_{S, \sigma_+, \sigma_-} \Pr_\pi[S_v^\pi = S, \pi_S = \sigma_+, \pi_{V \setminus S} = \sigma_- \mid \mathcal{E}] \\ & \cdot \mathbb{E}_\pi[|T_0 \cup T_1| \mid \mathcal{E}, S_v^\pi = S, \pi_S = \sigma_+, \pi_{V \setminus S} = \sigma_-] \\ & \leq \sum_{S, \sigma_+, \sigma_-} O(|S| \log^2 n + n \log n / 2^b + |S|n \cdot \Pr_{\pi \in \Omega}[\pi \in \mathcal{B}]) \\ & \cdot \Pr_\pi[S_v^\pi = S, \pi_S = \sigma_+, \pi_{V \setminus S} = \sigma_- \mid \mathcal{E}] \\ & = \sum_S O(|S| \log^2 n \cdot \Pr_\pi[S_v^\pi = S \mid \mathcal{E}]) + O(n \log n / 2^b) \\ & + \sum_{S, \sigma_+, \sigma_-} |S|n \cdot \Pr_\pi[S_v^\pi = S, \pi_S = \sigma_+, \pi_{V \setminus S} = \sigma_- \mid \mathcal{E}] \\ & \cdot \Pr_\pi[\pi \in \mathcal{B} \mid \mathcal{E}, S_v^\pi = S, \pi_S = \sigma_+, \pi_{V \setminus S} = \sigma_-] \\ & \leq O(\mathbb{E}_\pi[|S_v^\pi| \log^2 n \mid \mathcal{E}] + n \log n / 2^b + n^2 \Pr_\pi[\pi \in \mathcal{B} \mid \mathcal{E}]) \\ & \leq O(n \log^2 n / 2^b + n^2 \Pr_\pi[\pi \in \mathcal{B} \mid \mathcal{E}]) \end{aligned}$$

The last inequality holds by Lemma 18. \blacksquare

To remove the extra term $\Pr_\pi[\pi \in \mathcal{B} \mid \mathcal{E}]$, apply Lemma 7 by fixing values of $\pi(u), \pi(v)$ and taking union bound over all k equal to powers of 2, we would know that $\pi \notin \mathcal{B}$ with high probability, namely $\Pr_\pi[\pi \in \mathcal{B} \mid \mathcal{E}] = n^{-2} \log n$, and thus $\mathbb{E}_\pi[|T_0 \cup T_1| \mid \mathcal{E}_0] \leq O(n \log^2 n / 2^b + \log n) = O(n \log^2 n / 2^b)$.

By definition of T_0 and T_1 , the total update time is proportional to $\Delta(G_b) \cdot (|T_0| + |T_1|)$ whose expectation is then bounded by $O(n^2 \log^3 n / 2^{2b})$. Since fixing the memberships of v takes time at most $O(n \log^2 n / 2^a)$, it immediately says that the expected update time is $O(n^2 \log^3 n / 2^{2b} + n \log^2 n / 2^a)$. Since the adversary is oblivious to the randomness used in the algorithm, the probability of inserting an edge between V_a and V_b with respect to π is $O(2^{a+b} / n^2)$. Hence, the expected running time would be $O(2^{a+b} / n^2 \cdot (n^2 \log^3 n / 2^{2b} + n \log^2 n / 2^a)) = O(2^{a-b} \log^3 n + \log^2 n)$. Summing over all different indices a, b , the total time would be $O(\log^4 n)$.

ACKNOWLEDGEMENT

The second author would like to thank Hengjie Zhang for helpful discussions.

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, December 1986.
- [2] Sepehr Assadi, Krzysztof Onak, Baruch Schieber, and Shay Solomon. Fully dynamic maximal independent set with sublinear update time. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 815–826, 2018.
- [3] Sepehr Assadi, Krzysztof Onak, Baruch Schieber, and Shay Solomon. Fully dynamic maximal independent set with sublinear in n update time. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1919–1936. SIAM, 2019.
- [4] S. Behnezhad, M. Derakhshan, M. Hajiaghayi, C. Stein, and M. Sudan. Fully dynamic maximal independent set with polylogarithmic update time. *To appear in FOCS*, 2019.
- [5] Aaron Bernstein, Sebastian Forster, and Monika Henzinger. A deamortization approach for dynamic spanner and dynamic maximal matching. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1899–1918. SIAM, 2019.
- [6] Keren Censor-Hillel, Elad Haramaty, and Zohar Karnin. Optimal dynamic distributed mis. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, pages 217–226. ACM, 2016.
- [7] Sebastian Daum, Seth Gilbert, Fabian Kuhn, and Calvin Newport. Leader election in shared spectrum radio networks. In *Proceedings of the 2012 ACM Symposium on Principles of Distributed Computing, PODC '12*, pages 215–224. ACM, 2012.
- [8] Yuhao Du and Hengjie Zhang. Improved algorithms for fully dynamic maximal independent set. *arXiv preprint arXiv:1804.08908*, 2018.
- [9] Manoj Gupta and Shahbaz Khan. Simple dynamic algorithms for maximal independent set and other problems. *arXiv preprint arXiv:1804.01823*, 2018.
- [10] Nathan Linial. Distributive graph algorithms global solutions from local data. In *Proceedings of the 28th Annual Symposium on Foundations of Computer Science, SFCS '87*, pages 331–335, Washington, DC, USA, 1987. IEEE Computer Society.
- [11] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15(4):1036–1055, November 1986.
- [12] Krzysztof Onak, Baruch Schieber, Shay Solomon, and Nicole Wein. Fully dynamic mis in uniformly sparse graphs. *arXiv preprint arXiv:1808.10316*, 2018.

- [13] Dongxiao Yu, Yuexuan Wang, Qiang-Sheng Hua, and Francis C.M. Lau. Distributed $(\delta + 1)$ -coloring in the physical model. *Theor. Comput. Sci.*, 553(C):37–56, December 2014.

APPENDIX

A. Proof of Lemma 9

Proof: Here we present an conceptually simpler proof than the one presented in [6]. Notice that it suffices to consider the case where $\sigma(z) = \pi(z), \forall z \notin \{x, y\}$ and $\sigma(x) = \pi(y), \sigma(y) = \pi(x)$, where x, y is an arbitrary pair of consecutive vertices in π such that $x \in S$ and $y \notin S$. As $\sigma(u) < \sigma(v), \pi(u) < \pi(v)$, it can never be $x = v$ and $y = u$. Denote $M = M_\pi$ be the greedy MIS in the old graph. The proof follows from the two statements below.

Claim 23. M was also the greedy MIS on the old version of G with respect to σ .

Proof: Recall from Lemma 5, M is the greedy MIS with respect to σ in the old graph if M is an MIS and for all $z \in V \setminus M, I_z^\sigma \cap M \neq \emptyset$. The first half is easy: M was the greedy MIS in the old version of G with respect to π , so M is certainly an MIS in the old graph. Now we turn to verify the greedy MIS constraints.

Since σ agrees with π on every vertex except for $\{x, y\}$, we only need to verify $\forall z \in \{x, y\} \setminus M, I_z^\sigma \cap M \neq \emptyset$. We can assume x, y are neighbors in the updated graph G ; otherwise switching the orders between x, y in π does not affect the greedy MIS constraint. Consider several cases.

- $x \in M, y \notin M$. In this case, if $\pi(y) > \pi(x)$, then $\sigma(x) > \sigma(y)$, and thus $x \in I_y^\sigma \cap M \neq \emptyset$. If $\pi(y) < \pi(x)$, then since $y \notin S, I_y^\pi \cap M \setminus S$ must be nonempty, and so there exists $z \in I_y^\pi \cap M \setminus S$ that dominates y . As $\sigma(z) = \pi(z) < \pi(x) = \sigma(y)$, $I_y^\sigma \cap M$ is also nonempty.
- $x, y \notin M$. Since x, y are consecutive in π , switching their positions in σ does not affect the invariant that $I_z^\sigma \cap M \neq \emptyset, \forall z \in \{x, y\}$.
- $x \notin M, y \in M$. By definition of $S, \pi(y) < \pi(x)$ as otherwise y would belong to S , and so $\sigma(y) > \sigma(x)$. If $x \neq v$, then x cannot belong to S by definition since x is dominated by some MIS vertices outside of S . If $x = v$, then $y \neq u$ as $\sigma(v) > \sigma(u)$. Right after the edge update x is still dominated by a vertex in M , namely y , which is also a predecessor in π , so $S = \emptyset$ which is a contradiction. ■

Claim 24. $S_v^\sigma = S$.

Proof: By the previous claim, M was also the greedy MIS on G with respect to order σ . We first argue that $S_v^\sigma \supseteq S$. To do this, we prove by an induction that for every $i \geq 0, S_i \subseteq S_v^\sigma$; we refer readers to the definition of influenced sets for the meaning of S_i , where S_i 's are defined with respect to permutation π , not σ .

- Basis. For $i = 0$, to argue $v \in S_v^\sigma$ we only need to prove $S_v^\sigma \neq \emptyset$. As $S_v^\pi \neq \emptyset$, the edge update can only be an insertion (u, v) and $u, v \in M$, or an edge deletion (u, v) and $u \in M, v \notin M$ plus that u is the only MIS predecessor that dominates v . Since σ and π agree on all vertices whose orders are $\leq \pi(v)$, v would also violate its greedy MIS constraint with respect to σ , and so $S_v^\sigma \neq \emptyset$.
- Induction. Suppose we already have $S_{i-1} \subseteq S_v^\sigma$. Then, by Lemma 8, any $z \in M$ such that $S_{i-1} \cap I_z^\sigma \neq \emptyset$ should belong to S_v^σ . Since π and σ have the same relative order on $S, S_{i-1} \cap I_z^\sigma$ would be the same as $S_{i-1} \cap I_z^\pi$ for any $z \in S_i \cap M$. On the other hand, for any $z \in S_i \setminus (M \cup \{v\})$, we claim $I_z^\sigma \cap M$ is also equal to $I_z^\pi \cap M$. The only possible violation comes from the case that $z = x$ and $y \in M$. However this is also not possible: if $\pi(y) > \pi(x)$, then as $y \notin S$, by definition when $x \neq u$, it would have been excluded from S , and otherwise if $x = v$ we would have $S_v^\pi = \emptyset$; if $\pi(y) < \pi(x)$, then y would have been added to S ; both lead to contradictions. Therefore, by definition of S_i , we also have $S_i \subseteq S_v^\sigma$.

To prove $S_v^\sigma \subseteq S$, by Lemma 8 it suffices to verify that (1) $\forall z \in M, I_z^\sigma \cap S \neq \emptyset$ iff $z \in S$; (2) $\forall z \notin M, I_z^\sigma \cap M \subseteq S$ iff $z \in S$. As σ is equal to π except for x, y , we only need to consider $z \in \{x, y\}$ in (1)(2). We can assume x, y are adjacent; otherwise switching the orders between x, y in π does not affect the invariant. Then it can never be the case where $x \notin M, y \in M$ as it would contradict the definition of S . So it is either $x \in M, y \notin M$ or $x, y \notin M$. Consider two cases.

- $x \in M, y \notin M$. In this case, $I_x^\sigma \cap S = \emptyset$ always holds as switching the positions between x, y does not affect the equality $I_x^\sigma \cap S = I_x^\pi \cap S \neq \emptyset$. If $\pi(y) < \pi(x)$, then since $y \notin M$, it must be $I_y^\pi \cap M \neq \emptyset$, and because $y \notin S$, there exists $z \in I_y^\pi \cap M \setminus S$. So $\sigma(z) = \pi(z) < \pi(y) = \sigma(x)$. By the previous claim we already know $M_\sigma = M$, and so $I_z^\sigma \cap M \not\subseteq S$. If $\pi(y) > \pi(x)$, then $I_z^\sigma \cap S \subseteq I_z^\pi \cap S = \emptyset$.
- $x, y \notin M$. Since x, y are consecutive in π , switching their positions in σ does not change

$I_z^\sigma \cap M, \forall z \in \{x, y\}$.

- $x \notin M, y \in M$. By definition of S , $\pi(y) < \pi(x)$ as otherwise y would belong to S , and so $\sigma(y) > \sigma(x)$. If $x \neq u$, then x cannot belong to S by definition since x is dominated by some MIS vertices outside of S . If $x = v$, then $y \neq u$ as $\sigma(v) > \sigma(u)$. Right after the edge update x is still dominated by a vertex in M , namely y , which is also a predecessor in π , so $S = \emptyset$ which is a contradiction. ■

B. Proof of Lemma 10

Proof: As the lemma is stated in a slightly different way from [6], for completeness we also present a proof here. Define an intermediate permutation τ by this operation: remove v from order σ and reinsert it back right after u . Then $\tau(u) < \tau(v), \tau_S = \pi_S, \tau_{V \setminus S} = \pi_{V \setminus S}$, and thus by Lemma 9 we have $S_v^\tau = S$. Namely, τ and π satisfy the same condition in the statement of the lemma.

Let $w = \arg \min_{x \in S \setminus \{v\}} \{\tau(x)\}$. First we argue that w and v are neighbors. If w was in M_τ , then by the inductive definition of S_v^τ , there exists $z \in S \setminus M_\tau$ such that z is a predecessor neighbor of w . By minimality of w , z must be equal to v , and hence w and v are adjacent. If w was not in M_τ , then it has an MIS predecessor $z \in S \cap M_\tau$, similarly by minimality of w , z must be equal to v , and hence w and v are adjacent.

Recalling the relation between τ and σ , we can view σ as a permutation derived from τ by first removing v from τ and then reinsert v back to τ at a certain position somewhere behind w . We claim that right after we remove v from τ before reinsertion, w belongs to the greedy MIS M_τ with respect to the current τ (which is without v). Consider the only two cases where S_v^τ could be nonempty.

- The edge update is an insertion and both of u, v were in M_τ . After the removal of v , w is no longer dominated by any MIS predecessor in M_τ , hence w must join M_τ .
- The edge update is a deletion, and u was in M_τ while v was not in M_τ , plus that u is the only MIS predecessor that dominates v . Since v was not in M_τ , then by minimality of $\tau(w)$ among $S \setminus \{v\}$, the only predecessor of w in S was v , and thus $w \in M_\tau$ before and after v 's removal.

When we insert v back to τ at some position after w , which produces permutation σ , since w is now an MIS predecessor of v , v does not belong to M_σ . If the edge

update is insertion then no changes would be made to M_σ and thus $S_v^\sigma = \emptyset$; if the edge update is deletion, then since v has a neighboring MIS predecessor other than u , which is w , M_σ would also stay unchanged, and thus $S_v^\sigma = \emptyset$. ■

C. Proof of Lemma 11

Proof: For notational convenience, define $\mathcal{E} = \{\pi(u) = A, \pi(v) \in [A + 1, B]\}$. For any vertex set $S \subseteq V \setminus \{u_j\}_{1 \leq j \leq a}$ containing v , and partial orders σ_+, σ_- on $S \setminus \{v\}$ and $V \setminus S$, with the property that there exists at least one permutation π that satisfies event \mathcal{E} , as well as $S_v^\pi = S, \pi_{S \setminus \{v\}} = \sigma_+, \pi_{V \setminus S} = \sigma_-$, define a set of permutations

$$\Omega_{S, \sigma_+, \sigma_-} = \{\pi \mid \mathcal{E}, \pi_{S \setminus \{v\}} = \sigma_+, \pi_{V \setminus S} = \sigma_-\}$$

By Lemma 9 and Lemma 10, for any $\pi \in \Omega_{S, \sigma_+, \sigma_-}$, $S_v^\pi = S$ when $\pi(v) = \min_{z \in S} \{\pi(z)\}$, and $S_v^\pi = \emptyset$ otherwise. Here is a basic property of $\Omega_{S, \sigma_+, \sigma_-}$.

Claim 25. For any two different $\Omega_{S, \sigma_+, \sigma_-} = \{\pi \mid \mathcal{E}, \pi_{S \setminus \{v\}} = \sigma_+, \pi_{V \setminus S} = \sigma_-\}$ and $\Omega_{S', \sigma'_+, \sigma'_-} = \{\pi \mid \mathcal{E}, \pi_{S' \setminus \{v\}} = \sigma'_+, \pi_{V \setminus S'} = \sigma'_-\}$, $\Omega_{S, \sigma_+, \sigma_-}$ and $\Omega_{S', \sigma'_+, \sigma'_-}$ are disjoint.

Proof: Suppose otherwise there exists $\tau \in \Omega_{S, \sigma_+, \sigma_-} \cap \Omega_{S', \sigma'_+, \sigma'_-}$. By definition, there exists $\pi \in \Omega_{S, \sigma_+, \sigma_-}$ that satisfies event \mathcal{E} , as well as $S_v^\pi = S, \pi_{S \setminus \{v\}} = \sigma_+, \pi_{V \setminus S} = \sigma_-$. By Lemma 10, v takes the minimum in π among S .

Remove v from τ and reinsert v back to τ right at position $A + 1$. We claim τ stays in $\Omega_{S, \sigma_+, \sigma_-} \cap \Omega_{S', \sigma'_+, \sigma'_-}$; this is because removal and reinsertion of v preserves τ 's induced order on $S \setminus \{v\}, V \setminus S$ and $S' \setminus \{v\}, V \setminus S'$. Now, since v takes the minimum among S in τ , we have $\tau_S = \pi_S, \tau_{V \setminus S} = \pi_{V \setminus S}$. By Lemma 9, $S_v^\tau = S_v^\pi = S$. Similarly we can also have $S_v^\tau = S'$. Therefore, $S = S'$. As $\tau \in \Omega_{S', \sigma'_+, \sigma'_-}$, we know immediately $\sigma_+ = \tau_S = \tau_{S'} = \sigma'_+, \sigma_- = \tau_{V \setminus S} = \tau_{V \setminus S'} = \sigma'_-$, which is a contradiction that Ω and Ω' are different. ■

By this claim, we can decompose the expectation as a sum of conditional ones:

$$\begin{aligned} & \mathbb{E}_\pi[|S_v^\pi| \mid \mathcal{E}] \\ &= \sum_{S, \sigma_+, \sigma_-} \Pr_\pi[\pi \in \Omega_{S, \sigma_+, \sigma_-} \mid \mathcal{E}] \cdot \mathbb{E}_\pi[|S_v^\pi| \mid \mathcal{E}, \pi \in \Omega_{S, \sigma_+, \sigma_-}] \end{aligned}$$

So it suffices to compute each term in the summation. Fix any S, σ_+, σ_- and $\Omega = \Omega_{S, \sigma_+, \sigma_-}$. Notice that by Lemma 9 and Lemma 10 we have:

$$\mathbb{E}_\pi[|S_v^\pi| \mid \mathcal{E}, \pi \in \Omega] = |S| \cdot \Pr_\pi[\pi(v) = \min_{z \in S} \{\pi(z)\} \mid \mathcal{E}, \pi \in \Omega]$$

To bound the probability $\Pr_{\pi}[\pi(v) = \min_{z \in S} \{\pi(z)\} \mid \mathcal{E}, \pi \in \Omega]$, on the one hand, any permutation $\pi \in \Omega$ can be constructed by picking an arbitrary position for v among $[A + 1, B]$, and then assign arbitrary positions for $S \setminus \{v\}$, so $|\Omega| = (B - A) \cdot \binom{n-A-1}{|S|-1}$. On the other hand, the total number of permutations such that v takes the minimum among S is $\binom{n-A}{|S|} - \binom{n-B}{|S|}$. Therefore, as π is uniformly drawn from Ω , we have:

$$\begin{aligned} & \Pr_{\pi}[\pi(v) = \min_{z \in S} \{\pi(z)\} \mid \mathcal{E}, \pi \in \Omega] \\ &= \frac{\binom{n-A}{|S|} - \binom{n-B}{|S|}}{(B-A) \cdot \binom{n-A-1}{|S|-1}} \\ &= \frac{\binom{n-A}{|S|} - \binom{n-B}{|S|}}{(B-A) \cdot \binom{n-A}{|S|} \cdot \frac{|S|}{n-A}} < \frac{n-A}{(B-A)|S|} \end{aligned}$$

Hence, $\mathbb{E}_{\pi}[|S_v^{\pi}| \mid \mathcal{E}, \pi \in \Omega] = |S| \cdot \Pr_{\pi}[\pi(v) = \min_{z \in S} \{\pi(z)\} \mid \mathcal{E}, \pi \in \Omega] < \frac{n-A}{B-A}$. Since all Ω are disjoint, ranging over all different choices for S, σ_+, σ_- , we have

$$\begin{aligned} & \mathbb{E}_{\pi}[|S_v^{\pi}| \mid \pi(u) = A, \pi(v) \in [A + 1, B]] \\ & < \frac{n-A}{B-A} < \frac{n}{B-A} \end{aligned}$$

■

D. Proof of Lemma 12

Proof: For u, v to both lie in $[A + 1, B]$, B must be larger than $A + 1$. For notational convenience, define $\mathcal{E} = \{A < \pi(u) < \pi(v) \leq B\}$. We decompose the expectation as:

$$\begin{aligned} & \mathbb{E}_{\pi}[|S_v^{\pi}| \mid \mathcal{E}] \\ &= \sum_{k=A+1}^{B-1} \Pr_{\pi}[\pi(u) = k \mid \mathcal{E}] \cdot \mathbb{E}_{\pi}[|S_v^{\pi}| \mid \mathcal{E}, \pi(u) = k] \\ &= \sum_{k=A+1}^{B-1} \frac{B-k}{\binom{B-A}{2}} \cdot \mathbb{E}_{\pi}[|S_v^{\pi}| \mid \mathcal{E}, \pi(u) = k] \end{aligned}$$

The second equality holds as $\Pr_{\pi}[\pi(u) = k \mid \mathcal{E}] = \frac{B-k}{\binom{B-A}{2}}$; this is because, conditioned on $\pi(u) = k$ as well as event \mathcal{E} , there are $(B-k) \cdot (n-A-2)!$ permutations π , while there are $\binom{B-A}{2} \cdot (n-A-2)!$ permutations π that satisfy event \mathcal{E} . Since π is drawn uniformly at random from the set of all permutations that satisfy event \mathcal{E} , we have $\Pr_{\pi}[\pi(u) = k \mid \mathcal{E}] = \frac{B-k}{\binom{B-A}{2}}$.

Using Lemma 11, we have:

$$\mathbb{E}_{\pi}[|S_v^{\pi}| \mid \mathcal{E}, \pi(u) = k] \leq \frac{n}{B-k}$$

Therefore,

$$\begin{aligned} \mathbb{E}_{\pi}[|S_v^{\pi}| \mid \mathcal{E}] &= \sum_{k=A+1}^{B-1} \frac{B-k}{\binom{B-A}{2}} \cdot \mathbb{E}_{\pi}[|S_v^{\pi}| \mid \mathcal{E}, \pi(u) = k] \\ &< \sum_{k=A+1}^{B-1} \frac{B-k}{(B-A-1)(B-A)/2} \cdot \frac{n}{B-k} < \frac{2n}{B-A} \end{aligned}$$

■