

Optimal Dynamic Distributed MIS

Keren Censor-Hillel^{*}
Department of Computer
Science, Technion
ckeren@cs.technion.ac.il

Elad Haramaty[†]
College of Computer and
Information Science,
Northeastern University
eladh@cs.technion.ac.il

Zohar Karnin
Yahoo Labs
zkarnin@ymail.com

ABSTRACT

Finding a maximal independent set (MIS) in a graph is a cornerstone task in distributed computing. The local nature of an MIS allows for fast solutions in a static distributed setting, which are logarithmic in the number of nodes or in their degrees. The result trivially applies for the dynamic distributed model, in which edges or nodes may be inserted or deleted. In this paper, we take a different approach which exploits locality to the extreme, and show how to update an MIS in a dynamic distributed setting, either *synchronous* or *asynchronous*, with only a *single adjustment* and in a single round, in expectation. These strong guarantees hold for the *complete fully dynamic* setting: Insertions and deletions, of edges as well as nodes, gracefully and abruptly. This strongly separates the static and dynamic distributed models, as super-constant lower bounds exist for computing an MIS in the former.

Our results are obtained by a novel analysis of the surprisingly simple solution of carefully simulating the greedy *sequential* MIS algorithm with a random ordering of the nodes. As such, our algorithm has a direct application as a 3-approximation algorithm for correlation clustering. This adds to the important toolbox of distributed graph decompositions, which are widely used as crucial building blocks in distributed computing.

Finally, our algorithm enjoys a useful *history-independence* property, meaning the output is independent of the history of topology changes that constructed that graph. This means the output cannot be chosen, or even biased, by the adversary in case its goal is to prevent us from optimizing some objective function.

1. INTRODUCTION

Dynamic environments are very common in distributed

^{*}Supported in part by ISF grant 1696/14.

[†]Supported in part by NSF grant CCF-1319206. Part of this work was done while the author was at Yahoo Labs.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PODC'16, July 25-28, 2016, Chicago, IL, USA

© 2016 ACM. ISBN 978-1-4503-3964-3/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2933057.2933083>

settings, where nodes may occasionally join and leave the network, and communication links may fail and be restored. This makes solving tasks in a dynamic distributed setting of fundamental interest.

Solutions for problems from the static distributed setting translate nicely into the dynamic distributed setting, by running them in response to topology changes, in order to adjust the output [6, 7, 46]. This can be quite efficient especially for *local* problems, such as finding a maximal independent set (MIS) in the network graph. The cornerstone MIS problem admits fast distributed solutions whose complexities are logarithmic in the size of the graph or in the degrees of the nodes [2, 28, 36, 49].

In this paper, we exploit locality to the extreme and present an MIS algorithm for the dynamic distributed setting, which surprisingly requires in expectation only a *single adjustment* and a single round. The adjustment measure of an algorithm is the number of nodes that need to change their output in response to the topology change. These strong guarantees hold for both *synchronous* or *asynchronous* models, and for the *complete fully dynamic* setting, i.e., we handle all cases of *insertions* and *deletions*, of *edges* as well as *nodes*, *gracefully* and *abruptly*.¹ This is a strong separation between the static model and dynamic distributed model under an oblivious adversary, as super-constant lower bounds exist for the static setting [43, 48]. We further prove that for any deterministic algorithm, there is a topology change that requires n adjustments in the worst case, with n being the number of nodes, thus we also strongly separate randomized and deterministic solutions. Below, we overview the challenge we face, our technique for overcoming it, and the applications of our result.

1.1 Is This Not Trivial?

Indeed, the local structure of an MIS allows the nodes to easily detect whether it has suffered from a topology change. At a first glance, it may seem trivial to fix an MIS after a single topology change, as described next².

Suppose a new node is inserted into the graph. The node simply has to check whether it is connected to any MIS node and enter the MIS if and only if the answer is no. In particular, no other node needs to change its output. Similarly, it is an easy exercise to check that if an edge $\{u, v\}$ is deleted, at most one node needs to change its output.

Despite being simple, these two examples are misleading as they do not capture what happens in case of other topol-

¹See definitions of topology changes in Section 2.

²The cautious reader should remain suspicious.

ogy changes. Suppose an MIS node v is deleted from the graph. All of the nodes who were previously neighbors of v , i.e., who were in $N(v)$, now need to check whether they still have MIS neighbors. This can be a large number of nodes who may need to change their output. This alone is still not a problem from the communication perspective, since no other node needs to change its output, and the nodes who were in $N(v)$ detect this immediately along with detecting the topology change. However, the crucial difficulty here is that these nodes need to coordinate their decisions. If two nodes in $N(v)$ are connected by an edge, then only one of them can enter the MIS. In general, the induced subgraph $N(v) \setminus \{v\}$ can have any structure, and deciding which nodes enter the MIS (among those who now do not have a neighbor in the MIS) is equivalent to *constructing* an MIS on the induced subgraph $N(v) \setminus \{v\}$ from scratch. As this graph can be linear in the size of the original graph, this would lead to fixing the MIS being asymptotically as expensive as structuring it from scratch.

1.2 Our Contribution

Our approach is surprisingly simple: We simulate the greedy *sequential* algorithm for solving MIS. The greedy sequential algorithm orders the nodes and then inspects them by increasing order. A node is added to the MIS if and only if it does not have a lower-order neighbor already in the MIS. We consider *random greedy*, the variant of greedy in which the order is chosen uniformly at random. Consider simulating random greedy in a dynamic environment with the following template (ignoring the model of computation/communication for the moment). Each node needs to maintain the invariant that its state depends only on the states of its neighbors with lower order, such that it is in the MIS if and only if none of its lower order neighbors are in the MIS. When a change occurs in the graph, nodes may need to change their output, perhaps more than once, until they form a new MIS. Our key technical contribution is in proving:

Theorem 1 For any *arbitrary* change in the graph, the expectation over all random orders, of the number of nodes that change their output in the above random greedy template is at most 1.

The Challenge: We denote by π the random order of nodes, we denote by v^* the only node (if any) for which the above invariant does not hold after the topology change, and we denote by S the set of nodes that need to be changed in order for the invariant to hold again at all nodes. We look at S' , the set of nodes that would have needed to be changed if the order was as in π , except for pushing v^* to be the first node in that order. The definition of S' does not depend on the real order of v^* in π . Therefore, we can prove that S can either be equal to S' if the order of v^* in π is minimal in S' , and empty otherwise. Now the question is, what is the probability, given S' , that v^* is indeed its minimal order node? The answer is that if S' were deterministic, i.e. independent of π , the probability would be $1/|S'|$. However, S' is a random set and having knowledge of its members restricts π to be non-uniform, which in turn requires a careful analysis of the required probability. To overcome this issue, we prove that the information that S' gives about π is either about the order between nodes not in S' , or about the order within

$S' \setminus \{v^*\}$, which both do not affect the probability that v^* is the minimal in S' .

Distributed Implementation: This powerful statement of $\mathbb{E}[|S|] \leq 1$ directly implies that a *single* adjustment is sufficient for maintaining an MIS in a dynamic setting. A direct distributed implementation of our template implies that in expectation also a single round is sufficient. This applies both to the synchronous and asynchronous models, where the number of rounds in the asynchronous model is defined as the longest path of communication.

Obtaining $O(1)$ Broadcasts and Bits: In fact, in the synchronous model, it is possible to obtain an expected number of $O(1)$ *broadcasts* and *bits*. Here the number of broadcasts is the total number of times, over all nodes, that any node sends a $O(\log n)$ -bit broadcast message (to all of its neighbors)³. Moreover, since we only need a node to know the order between itself and its neighbors, using a similar technique to that of [51], we obtain that in expectation, a node only needs to send a constant number of bits in each broadcast. The above holds for edge insertions and deletions, graceful node deletion, and node *unmuting* (which is defined formally in Section 2), while for an abrupt deletion of a node v^* we will need $O(\min\{\log(n), d(v^*)\})$ broadcasts, and for an insertion of a node v^* we will need $O(d(v^*))$ broadcasts, in expectation.

This is done with a careful dynamic distributed implementation (Section 4) which guarantees that each node that changes its output does so at most $O(1)$ times, as opposed to the direct distributed implementation⁴. Hence, obtaining these broadcast and bit complexities comes at a cost of increasing the round complexity, but it remains constant (albeit not 1).

Matching Lower Bounds: We claim that any deterministic algorithm requires n adjustments in the worst case, which can be seen through the following example. Let A be a dynamic deterministic MIS algorithm. Let G_0 be the complete bipartite graph over two sets of nodes of size k . We denote by L the side of G_0 that is chosen to be the MIS by A , and we denote the other side by R . For every $i \in [k]$ let G_i be the graph obtained after deleting i nodes from L , and consider executing A on G_0, G_1, \dots, G_k . For every i , since G_i is a complete bipartite graph, one of the sides has to be the MIS. Since G_k contains only disconnected nodes of R then R is the only MIS of G_k . This implies that after some specific change along the sequence, the side of the MIS changes from L to R . In this topology change, *all* of the nodes change their output.

This gives a strong separation between our result and deterministic algorithms. Moreover, it shows that (1) the expected adjustment complexity of any algorithm must be at least 1, as we have a sequence of k topology changes that lead to at least k adjustments, and (2) it is impossible to achieve high probability bounds that improve upon a simple Markov bound. Specifically, this explains why we obtain our result in expectation, rather than with high probability.

³We emphasize that the term broadcast is used here to indicate the more restricted setting of not being able to send different messages to different neighbors in the same round. It does not refer to a wireless setting of communication.

⁴This bears some similarity to the method in [60], where the number of *moves* is reduced in an MIS self-stabilizing algorithm by adding a possible *wait* state to the standard *in MIS* and *not in MIS* states.

This is because the example can be inserted into any larger graph on n nodes, showing that *for every* value of k , there exists an instance for which at least $\Omega(k)$ adjustments are needed with probability at least $1/k$.

Approximate Correlation Clustering: In addition to the optimal complexity guarantees, the fact that our algorithm simulates the random greedy sequential algorithm has a significant application to correlation clustering. Correlation clustering requires the nodes to be partitioned into clusters in a way that minimizes the sum of the number of edges outside clusters and the number of non neighboring pairs of nodes within clusters (that is, missing edges within clusters). Ailon et al. [1] show that random greedy obtains a 3-approximation for correlation clustering⁵, by having each MIS node inducing a cluster, and each node not in the MIS belonging to the cluster induced by the smallest random ID among its MIS neighbors. This directly translates to our model, by having the nodes know that random ID of their neighbors. Graph decompositions play a vital role in distributed computing (see, e.g., [54]), and hence the importance of obtaining a 3-approximation for correlation clustering.

History Independence: Finally, our algorithm has a useful property, which we call *history independence* (in full version [16]), which means that the structure output by the algorithm (e.g., the MIS) depends only on the current graph, and does not depend on the history of topology changes. This means that the output cannot be chosen, or even biased, by the adversary, in case its goal is to prevent us from optimizing some objective function. Moreover, history independent algorithms compose nicely, which allows us to obtain history independent coloring and matching algorithms, using standard reductions.

1.3 Previous Work

A strongly related previous work is that of König and Wattenhofer [39], where they also show that the above can be obtained in a single round. However, their setting differs crucially from ours, by two major aspects. In their solution, an MIS node that needs to be removed from the MIS (following a topology change) already has a “last will” for it in the local states of its neighbors, which tells each of them whether it should enter the MIS or not. This constitutes the first huge difference: our algorithm updates only a *single state* in expectation. But even more critical is that in order to compute this last will, each node must be aware of the edges between its neighbors. This is an enormous amount of information, which can be quadratic in the size of the network, and is therefore ruled out in our setting as an unacceptable amount of information to be sent in a single round. Our algorithms send only a constant number of bits in each message.

In fact, [39] argues that without a “last will” it is impossible to fix an MIS in constant time. This because a node v that is in the MIS can be deleted, and its neighborhood $N(v)$ can be arbitrarily connected, which reduces the problem of fixing the solution to that of finding an MIS in $N(v)$. Since $N(v)$ can have linear size, and since lower bounds on

⁵In the same paper, they also provide a 2.5 approximation based on rounding a solution of a linear program. We do not elaborate on the details of this algorithm, nor the history of the correlation clustering problem as it is outside the scope of our paper.

finding an MIS are super-constant [43, 48], this implies that some mechanism is required in order to be able to fix an MIS in a constant number of rounds.

However, notice that the above argument only holds under the implicit assumption that the MIS and the topology change can be arbitrary, i.e., controlled by an adaptive adversary. Our contribution is exactly in making sure that such a bad example does not occur. For this, we make the standard assumption of an *oblivious* non-adaptive adversary. This means that the topology changes do not depend on the randomness of the algorithm. This is the standard assumption in dynamic settings (see literature about sequential dynamic computing, e.g., [18] and references therein), where typically an adaptive adversary that can choose the topology changes based on the status of the data structure that the algorithm maintains is too strong and renders most problems unsolvable in any reasonable complexity measure. An oblivious adversary is natural also for our setting, as, for example, an adaptive adversary can always choose to delete MIS nodes and thereby force worst-case behavior in terms of the number of adjustments.

1.4 Additional Related Work

Distributed MIS: Finding an MIS is a central theme in the classical distributed setting. The classic algorithms [2, 36, 49] complete within $O(\log n)$ rounds, with high probability. More recently, a beautiful line of work reduced the round complexity to depend on Δ , the maximal degree in the graph. These include the $O(\Delta + \log^* n)$ -round and $O(\log^2 \Delta + 2^{O(\sqrt{\log \log n})})$ -round algorithms of [10], the $O(\log \Delta \sqrt{\log n})$ -round algorithm of [11], and the very recent $O(\log \Delta) + 2^{O(\sqrt{\log \log n})}$ -round algorithm of [28]. An excellent source for additional background and relations to coloring can be found in [9].⁶

Distributed dynamic MIS: The problem of finding a fast dynamic distributed MIS algorithm appears as an open problem in [24], which studies the problem of maintaining a sparse spanner a dynamic setting. Additional problems in a dynamic setting are also addressed in [12], and in slightly different settings in [5, 17, 37, 41]. One standard approach for maintaining an MIS is running distributed algorithms that are designed for the static setting. This can be done for any distributed algorithm, sometimes using a corresponding *compiler*, e.g., when applied to an asynchronous dynamic setting [6, 7, 46].

One important exception is the aforementioned solution shown by König and Wattenhofer [39], which as in our algorithm, requires a constant number of rounds, but as opposed to our algorithm, has the following two very expensive properties: (1) the number of updates following a topology change can be linear in the number of nodes and (2) the size of messages can be almost quadratic in the number of nodes. The only way to overcome item (2) is to pay the cost of that information in terms of round complexity rather

⁶As in all of these static MIS algorithms, we also do not maintain a Lexicographic MIS (LexMIS). One could argue that our algorithm maintains an MIS that corresponds to the random permutation over the nodes. This, however, by itself is not as powerful as maintaining a LexMIS with respect to any *given* order, say, of node IDs. In fact, any MIS can be seen as an MIS that is induced by *some* order of the nodes, so only an MIS corresponding to a specific given order should be considered as LexMIS.

than message complexity, a solution that is very expensive as well. We are unaware of any other work in this setting about maintaining an MIS.

Additional distributed dynamic algorithms: Indeed, because of its extreme importance, dynamic distributed computing is a widely studied area of research. However, different assumptions on the setting have a significant influence on what can be done.

A huge amount of literature is devoted to devising different algorithms in a self-stabilizing setting (see, e.g., [20, 21, 29, 57] and references therein). This setting is inherently different from ours since it is asynchronous, and considers the *time* that it takes an algorithm to stabilize as the number of asynchronous rounds until the set of outputs satisfies the problem specifications and all of the nodes are quiet (it is typically not possible to actually detect this state and terminate). An asynchronous round requires that each node communicates with all of its neighbors, and hence inherently implies a lot of communication (broadcasts). In our work, we do not assume that the topology ever stops changing, and just assume that enough time is given in order to recover from a change. This is why it is crucial to obtain an algorithm that requires only a constant (in fact, even a single) update and round in order to recover.

An MIS-based clustering algorithm for the asynchronous model that appeared in [23] also uses a random node order for recovering after a change. However, their self-stabilizing setting differs from ours in several aspects, such as assuming a bounded degree graph and discussing corrupted states of multiple nodes and topology changes. In addition, our techniques and analysis are completely different. In particular, the clustering obtained there may not be an approximation to correlation clustering. Furthermore, the number of rounds required by [23] is $O(\log(n))$ as opposed to our single round (in expectation) algorithm.

Related, but not identical, notions of error confinement, fault local and fault amendable algorithms have been studied in [8, 44, 45], where the internal memory of a node may change. Another property that self-stabilizing algorithms should aim for is super-stabilization [22], which means that they are self-stabilizing (eventually output the required structure) and also recover extremely fast from a single topology change. Super-stabilization requires also a small adjustment measure, which is the maximum number of nodes that have to change their output. Our MIS algorithm recovers from a single topology change in a constant number of rounds, and has an adjustment measure of exactly 1, in expectation.

Other studied models are of more severe graph changes, whether arbitrary [42] or evolving randomly [4]. In such models, one typically addresses problems of gathering information rather than maintaining graph structures, as the latter is intuitively unattainable. In contrast, in this paper, we consider a dynamic distributed setting which is in the same spirit of the literature on *sequential dynamic algorithms* [18].

Simulating the sequential greedy algorithm: Simulating random greedy has been used before in order to obtain fast solutions for sequential local computation algorithms (LCA). In this setting, the algorithm does not have access to the entire graph, but rather an oracle access using queries about nodes or edges, and needs to provide an approximate solution for various problems, among which are the problems considered in this paper. We emphasize that the models

are inherently different, and hence is our technical analysis. While we bound the size of the set of nodes that may change their output after a topology change, studies in the local computation literature [3, 47, 50, 52, 61] bound the size of the set of nodes that need to be recursively queried in order to answer a *random* node query. In some sense, these sets are opposite: We begin with a single node v changing its state due to a topology change, and look at the set of nodes that change their state *due* to the change of v . Local computation algorithms begin with a single node v and look at the set of nodes whose states *determine* the state of v .

In [15], the DAG of *all* nodes is considered, according to the same random order we consider here. The paper analyzes its depth (rather than its size, which is n) and obtains a bound of $O(\log^2 n)$, w.h.p. This translates to a bound on the performance of the parallel random greedy algorithm for finding an MIS. It also implies that w.h.p., our algorithm finishes to update after at most $O(\log^2 n)$ steps.

2. DYNAMIC DISTRIBUTED COMPUTATIONS

The distributed setup is a standard message passing model. The network is modeled by an undirected graph $G = (V, E)$ where the node set is V , and E consists of the node pairs that have the ability to directly communicate. We assume a broadcast setting where a message sent by a node is heard by all of its neighbors. Also, we assume a synchronous communication model, where time is divided into rounds and in each round any willing node can broadcast a message to its neighbors. We restrict the size of each message to be $O(\log(n))$ bits, with $n = |V|$ being the size of the network⁷. The computational task is to maintain a graph structure, such as a maximal independent set (MIS) or a node clustering. That is, each node has an output, such that the set of outputs defines the required structure.

Our focus is on a dynamic network, where the graph changes over time. As a result, nodes may need to communicate in order to adjust their outputs. The system is *stable* is when the structure defined by the outputs satisfies the problem requirements.

A graph topology change can be with respect to either an edge or a node. In both cases we address both deletions and insertions, both of which are further split into two different types. For deletions we discuss both a *graceful deletion* and an *abrupt deletion*. In the former, the deleted node (edge) may be used for passing messages between its neighbors (endpoints), and retires completely only once the system is stable again. In the latter, the neighbors of the deleted node simply discover that the node (edge) has retired but it cannot be used for communication. For insertions, we distinguish between a *new node insertion* and an *unmuting* of a previously existing node. In the former, a new node is inserted to the graph, possibly with multiple edges. In the latter, a node that was previously invisible to its neighbors

⁷This is the standard assumption in a distributed setting. In our dynamic setting where the size of the graph may change we assume knowledge of some upper bound $N \geq n$, with $N = n^{O(1)}$, and restrict the message length to $O(\log(N)) = O(\log(n))$.

but heard their communication, becomes visible and enters the graph topology⁸.

We assume that the system is always stable before a change occurs. Since our algorithm requires in expectation only an extremely short single time unit to recover, this even allows topology changes that are more frequent compared to other models. Specifically, many previous work in dynamic models, whether sequential [18] or distributed [30, 32, 40, 53], assume that the topology does not change too fast, so that there is only a single topology change at a time with sufficient recovery time in between changes. Since our algorithm runs in expectation in a super-fast constant time which is completely independent on the size of the network, we can in fact tolerate more changes.

We consider the performance of an algorithm according to three complexity measures. The first is the *adjustment-complexity*, measuring the number of nodes that change their output as a result of the recent topology change. The second is the *round-complexity*, which is the number of rounds required for the system to become stable. Finally, the third, more harsh, score is the *broadcast-complexity*, measuring the total number of broadcasts.

Our algorithms are randomized and thus our results apply to the expected values of the above measures, where the expectation is taken over the randomness of the nodes. We emphasize that this is the only randomness discussed; specifically, the result is not for a random node in the graph nor a random sequence of changes, but rather applies to any node and any sequence of changes. It holds for *every* change in the graph, not only amortized over all changes.

In what follows we discuss the problem of computing an MIS. Here, the outputs of the nodes define a set M , where any two nodes in M are not connected by an edge, and any node not in M has a neighbor in M . The second problem we discuss is that of *correlation clustering*. Here, the objective is to find a partitioning \mathcal{C} of the node set V , where we favor partitions with a small number of “contradicting edges”. That is, we aim to minimize the sum $\sum_{C \in \mathcal{C}} \sum_{u, v \in C} \mathbb{1}_{[(u, v) \notin E]} + \sum_{C_1 \neq C_2 \in \mathcal{C}} \sum_{u \in C_1, v \in C_2} \mathbb{1}_{[(u, v) \in E]}$.

3. A TEMPLATE FOR MAINTAINING A MAXIMAL INDEPENDENT SET

In this section we describe a template for maintaining a maximal independent set (MIS). Initially, we are given a graph $G = (V, E)$ along with an MIS that satisfies certain properties, and after a topology change occurs in the graph, applying the template results in an MIS that satisfies the same properties. That is, the template describes what we do after a single topology change, and if one considers a long-lived process of topology changes, then this would correspond to having initially an empty graph and maintaining an MIS as it evolves. We emphasize that the template describes a process that is not in any particular model of computation, and later in Section 4 we show how to implement it efficiently in our dynamic distributed setting. This also means that there are only four topology changes we need to consider: edge-insertion, edge-deletion, node-insertion and node-deletion. For example, the notions of abrupt and graceful node deletions are defined with re-

⁸The distinction is only relevant for nodes insertions, as there is no knowledge associated with an edge.

spect to the dynamic distributed setting because they affect communication, and therefore the implementation of the template will have to address this distinction, but the template itself is only concerned with a single type of node deletion, not in any particular computation model.

Throughout, we assume a uniformly random permutation π on the nodes $v \in V$. We define two *states* in which each node can be: M for an MIS node, and \bar{M} for a non-MIS node. We abuse notations and also denote by M and \bar{M} the sets of all MIS and non-MIS nodes, respectively. Our goal is to maintain the following *MIS invariant*: A node v is in M if and only if all of its neighbors $u \in N(v)$ which are ordered before it according to π , i.e., for which $\pi(u) < \pi(v)$, are not in M . It is easy to verify that whenever the MIS invariant is satisfied, it holds that the set M is a maximal independent set in G . Furthermore, it is easy to verify that this invariant simulates the greedy sequential algorithm, as defined in the introduction.

When any of the four topology changes occurs, there is at most a single node for which the MIS invariant no longer holds, as we explain next and in Footnote 9. We denote this node by $v^* = v^*(G_{\text{old}}, G_{\text{new}}, \pi)$, where G_{old} and G_{new} are the graphs before and after the topology change. For an edge insertion or deletion, v^* is the endpoint with the larger order according to π . For a node insertion or deletion, v^* is that node⁹. In case the topology change is an edge change, we will need also to take into consideration its other endpoint. We denote it by $v^{**} = v^{**}(G_{\text{old}}, G_{\text{new}}, \pi)$, and notice that by our notation, it must be the case that $\pi(v^{**}) < \pi(v^*)$. In order to unify our proofs for all of the four possible topology changes, we talk about a node v^{**} also for node changes. In this case we define v^{**} to be v^* itself, and we have that $\pi(v^{**}) = \pi(v^*)$. Therefore, for any topology change, it holds that $\pi(v^{**}) \leq \pi(v^*)$.

To describe our template, consider the case where a new edge is inserted and it connects two nodes $\pi(v^{**}) < \pi(v^*)$, where both nodes are in M . As a result, v^* must now be deleted from the MIS and hence we need to change its state. Notice that as a result of the change in the state of v^* , additional nodes may need their state to be changed, causing multiple state changes in the graph. An important observation is that it is possible that during this process of propagating local corrections of the MIS invariant, we change the state of a node more than once. As a simple example, consider the case in which v^* has two neighbors, u_1 and u_2 , for which $\pi(v^*) < \pi(u_1), \pi(u_2)$, and that u_1 and u_2 are connected by a path (u_1, w_1, w_2, u_2) , with $\pi(u_1) < \pi(w_1) < \pi(w_2) < \pi(u_2)$. Now, when we change the state of v^* to \bar{M} , both u_1 and u_2 need to be changed to M , for the MIS invariant to hold. This implies that w_1 needs to be changed to \bar{M} and w_2 needs to be changed to M . In this case, since $\pi(w_2) < \pi(u_2)$, the node u_2 needs to be changed back to state \bar{M} .

The above observation leads us to define a set of *influenced* nodes, denoted by $S = S(G_{\text{old}}, G_{\text{new}}, \pi)$, containing

⁹For a node deletion, we slightly abuse the definition of v^* in order to facilitate the presentation, and consider it to be the deleted node. This means that here we consider an intermediate stage of having v^* still belong to the graph w.r.t. the MIS invariant of all the other nodes, but for v^* the MIS invariant no longer holds. This is in order to unify the four cases, otherwise we would have to consider all of the neighbors of a deleted node as nodes for which the MIS invariant no longer holds after the topology change.

v^* in the scenario where we need to change its state, and all other nodes whose state we must subsequently change as a result of the state change of v^* . To formally define the set S we introduce some notations. The notations rely on the graph structure of G_{new} unless the change is a node deletion in which case they rely on G_{old} . For each node u , we define $I_\pi(u) = \{v \in N(u) \mid \pi(v) < \pi(u)\}$, the set of neighbors of u that are ordered before it according to π . These are the nodes that can potentially *influence* the state of u according to the MIS invariant. The definition of S is recursive, according to the ordering induced by π . If immediately after the topology change, in the new graph G with the order π it holds that the MIS invariant still holds for v^* , then we define $S = \emptyset$. (This is motivated by the fact that no node is influenced by this change.) Otherwise, we denote $S_0 = \{v^*\}$, and inductively define

$$S_i = \{u \mid u \in M, \text{ and } S_{i-1} \cap I_\pi(u) \neq \emptyset\} \cup \{u \mid u \in \bar{M}, \text{ and every } v \in I_\pi(u) \cap M \text{ is in } \cup_{j=0}^{i-1} S_j\}. \quad (1)$$

The set S is then defined as $S = \bigcup_i S_i$.¹⁰ Notice that a node u can be in more than one set S_i , as is the case for u_2 in the example above, which is in both S_1 and S_4 . The impact of a node u being in more than one S_i is that in order to maintain the MIS invariant, we need to make sure that we update the state of u after we update that of w , for any w such that $w \in I_\pi(u)$. Instead of updating the state of u twice, we can simply wait and update it only after the state of every such w is updated. For this, we denote by $i_u = \max\{i \mid u \in S_i\}$ the maximal index i for which u is in S_i .

Algorithm 1 A Template for MIS.

Initially, $G = (V, E)$ satisfies the MIS invariant.

On topology change at node v^* do:

1. Update state of v^* if required for MIS to hold
 2. For $i \leftarrow 1$, until $S_i = \emptyset$, do:
 3. For every $u \in S_i$ such that $i = i_u$:
 4. Update state of u
 5. $i \leftarrow i + 1$
-

We formally describe our template in Algorithm 1. By construction, the updated states after executing Algorithm 1 satisfy the MIS invariant. In addition, the crucial property that is satisfied by the above template is that in expectation, the size of the set S is 1. The remainder of this section is devoted to proving the following, which is our main technical result.

THEOREM 1. $\mathbb{E}_\pi [|S(G_{\text{old}}, G_{\text{new}}, \pi)|] \leq 1$.

Outline of the proof.

In order to prove that $\mathbb{E}[|S|] \leq 1$, instead of analyzing the set S directly, we analyze the set $S' = S'(G_{\text{old}}, G_{\text{new}}, \pi, v^*)$, which is defined via recursion similarly to S with three modifications: (1) It is always the case that $S'_0 = \{v^*\}$ (2) The graph according to which S' is defined is G_{old} in the case of a node deletion or an edge insertion, and G_{new} otherwise. (3) For any permutation π' that is identical to π on

all pairs not containing v^* , $S'(G_{\text{old}}, G_{\text{new}}, \pi, v^*)$ remains the same set. In other words, having knowledge of S' does not provide information regarding the location of v^* in π .

In Lemma 2, we prove that if $\pi(v^*) \neq \min\{\pi(u) \mid u \in S'\}$ then $S = \emptyset$, and otherwise $S = S'$ (in fact, it would be enough that $S \subseteq S'$). Then, in Lemma 3, we prove that for any set $P \subseteq V$, given the event that $S' = P$, the probability, over the random choice of π , that $\pi(v^*) = \min\{\pi(u) \mid u \in P\}$ is $1/|P|$. This leads to the required result of Theorem 1. Lemma 3 would be trivial if there was no dependency between π and S' . However, the trap we must avoid here is that S' is defined according to π , and therefore when analyzing its size we cannot treat π as a *uniformly* random permutation. To see why, suppose we know that inside $S' \setminus \{v^*\}$ we have nodes with large order in π . Then the probability that the order of v^* in π is smaller than all nodes in $S' \setminus \{v^*\}$, is much larger than $1/|S'|$, and can in fact be as large as $1 - o(1)$. In other words, S' gives some information over π . Nevertheless, we show that for our particular definition of S' this information is either about the order between nodes outside of S' , or about the order between nodes within $S' \setminus \{v^*\}$. Both types of restrictions on π do not affect the probability that v^* is the minimal of S' .

We now formally prove our result as outlined above. Throughout we use the notation $u \in M$ or $u \in \bar{M}$. This applies only to nodes u for which we are guaranteed that their states remain the same despite the topology change.

LEMMA 2. *If $\pi(v^*) \neq \min\{\pi(u) \mid u \in S'\}$ then $S = \emptyset$. Otherwise, $S \subseteq S'$.*

PROOF. First, assume that $\pi(v^*) \neq \min\{\pi(u) \mid u \in S'\}$. We show that the MIS invariant still holds after the topology change, and so $S = \emptyset$. Consider the node w , for which $\pi(w) = \min\{\pi(u) \mid u \in S'\}$. Notice that $w \notin S$, because $\pi(w) < \pi(v^*)$ and by the definition of S it can contain only nodes whose order in π is larger than that of v^* . We claim that $w \in M$. Assume, towards a contradiction, that $w \in \bar{M}$. This implies that w has a neighbor $u \in M$ such that $\pi(u) < \pi(w)$. For this node u we must have $u \notin S'$ due to the minimality of $\pi(w)$. It follows, according to the construction of S' that w cannot be an element of S' , leading to a contradiction.

We have that $w \in M$ and due to the minimality of $\pi(w)$, it must be that $w \in S'_1$, which implies that w is a neighbor of v^* . But then, when considering S , v^* has a neighbor other than v^* (since v^* and v^* cannot be neighbors in the graph defining S') which is ordered before it according to π which is in M . In the case of an edge insertion or deletion, this means that v^* remains in \bar{M} despite the topology change meaning that $S = \emptyset$. In the case of a node deletion, v^* was not in M prior to the change hence $S = \emptyset$. In the case of a node insertion, v^* does not enter M hence again, $S = \emptyset$.

Next, assume that $\pi(v^*) = \min\{\pi(u) \mid u \in S'\}$. We show that either $S = \emptyset$ or $S = S'$. If there is no need to change the state of v^* as a result of the topological change then $S_0 = \emptyset$, and so $S = \emptyset$ and the claim holds. It remains to analyze the case where $S_0 = S'_0 = \{v^*\}$. If $u \in S'_1$ then $\pi(v^*) < \pi(u)$ hence according to its definition $u \in S_1$. If $u \notin S'_1$ then u must have a neighbor $w \in M$ with $\pi(w) < \pi(u)$ meaning that $u \notin S_1$. We have that $S_1 = S'_1$ and similarly $S_i = S'_i$ for all $i > 1$. We conclude that $S' = S$ as required. \square

The following lemma shows that the probability of having

¹⁰All states above refer to the time at which the topology change occurs. The times of actual updates will in fact be captured by the indexes i of the sets S_i to which a node belongs.

$S = S'$ is $1/|S'|$, which immediately lead to Theorem 1 as the only other alternative is $S = \emptyset$.

LEMMA 3. *For any set of nodes $P \subseteq V$, it holds that $\Pr[\pi(v^*) = \min\{\pi(u) \mid u \in P\} \mid S' = P \text{ and } \pi(v^{**}) \leq \pi(v^*)] = \frac{1}{|P|}$.*

To prove this lemma we focus on S' . Notice that the events we considered in the previous lemma depend only on the ordering implied by π and hold for any configuration of states for the nodes that satisfy the MIS invariant. Roughly speaking, the lemma will follow from the fact that the event $S' = P$ does not give any information about the order implied by π between nodes in P and nodes in $V \setminus P$. To this end, for every permutation τ on V , we define $S'(\tau) = S'(G_{\text{old}}, G_{\text{new}}, \tau, v^*)$ as the set corresponding to S' under the ordering induced by τ . We denote by Π_P the set of all permutations τ for which it holds that $S'(\tau) = P$. We first need to establish the following about permutations in Π_P : If π and σ are two permutations on V such that $\pi|_P = \sigma|_P$ and $\pi|_{V \setminus P} = \sigma|_{V \setminus P}$, then $\sigma \in \Pi_P$ if and only if $\pi \in \Pi_P$.

CLAIM 4. *Let $P \subseteq V$ be a set of nodes, and let π and σ be two permutations such that $\pi|_P = \sigma|_P$ and $\pi|_{V \setminus P} = \sigma|_{V \setminus P}$. Assume $\pi \in \Pi_P$. We have that $V \setminus P \subseteq V \setminus S'(\sigma)$ and every $u \in V \setminus P$ has the same state according to π and σ .*

PROOF. Let $u \in V \setminus P$. We prove that $u \in V \setminus S'(\sigma)$ and that its state under σ is the same as it is under π by induction on the order of nodes in $V \setminus P$ according to π (which is equal to their order according to σ).

For the base case, assume that u has the minimal order in $V \setminus P$. We claim that u cannot have a neighbor in P . Assume, towards a contradiction, that u has a neighbor $w \in P$. Since $w \in P$ then it is possible that after the updates, w will be in M . In particular, w cannot have a neighbor in M , otherwise w would not be in P , since two nodes in M cannot be neighbors and $u \notin P$. Hence, u must be in \bar{M} according to π . In this case there is a node $z \in I_\pi(u) \cap V \setminus P$ that is in M according to π . But this cannot occur due to the minimality of $\pi(u)$ in $V \setminus P$. Therefore, u has no neighbors in P as required.

We have that all of the neighbors of u are in $V \setminus P$ and that u is the minimal among its neighbors according to π . Since $\pi|_{V \setminus P} = \sigma|_{V \setminus P}$ we have that u has the minimal order among its neighbors according to σ . This translates into u having a state of M under σ and in particular, u is not an element of $S'(\sigma)$, thus proving our base case.

For the induction step, consider a node $u \in V \setminus P$, and assume the claim holds for every $w \in V \setminus P \cap I_\pi(u)$. We consider two cases, depending on whether u has a neighbor in P or not.

Case 1: u does not have any neighbor in P . If $u \in \bar{M}$, then there is a node $z \in I_\pi(u) \cap V \setminus P$ that is in M according to π . By the induction hypothesis, $z \in V \setminus S'(\sigma)$ and $z \in M$ also according to σ . Since $\pi|_{V \setminus P} = \sigma|_{V \setminus P}$, we have that u is in \bar{M} according to σ too. Otherwise, if $u \in M$, then every $w \in I_\pi(u)$ (which is also $V \setminus P$) is in \bar{M} according to π . Any node $w \in I_\sigma(u)$ is also in $w \in I_\pi(u)$, since it is not in P and $\pi|_{V \setminus P} = \sigma|_{V \setminus P}$. The induction hypothesis on w gives that it is also in $V \setminus S'(\sigma)$ (otherwise it would be in $S'_\pi = P$ in contradiction to the assumption of case 1), and its state according to σ is \bar{M} . Hence, u must be in $V \setminus S'(\sigma)$ as well, and in state M according to σ .

Case 2: Assume u has a neighbor $w \in P$. Since $w \in P$, it is possible that after the updates, w will be in M . Since two nodes in M cannot be neighbors and $u \notin P$, then u must be in \bar{M} according to π . In this case there is a node $z \in I_\pi(u) \setminus P$ that is in M according to π . By the induction hypothesis, $z \in V \setminus S'(\sigma)$ and $z \in M$ also according to σ . Since $\pi|_{V \setminus P} = \sigma|_{V \setminus P}$, then u is in \bar{M} and in $V \setminus S'(\sigma)$ according to σ too. \square

CLAIM 5. *Let $P \subseteq V$ be a set of nodes, and let π and σ be two permutations such that $\pi|_P = \sigma|_P$ and $\pi|_{V \setminus P} = \sigma|_{V \setminus P}$. Assume $\pi \in \Pi_P$. We have that $P \subseteq S'(\sigma)$.*

The proof of Claim 5 appears in the full version [16]. Claims 4 and 5 combined imply that if $\pi|_P = \sigma|_P$ and $\pi|_{V \setminus P} = \sigma|_{V \setminus P}$ then $\sigma \in \Pi_P$ if and only if $\pi \in \Pi_P$. This is used in the full version [16] to prove Lemma 3. Lemmas 2 and 3 immediately lead to Theorem 1. As an immediate corollary of Theorem 1 we get:

COROLLARY 6. *A direct distributed implementation of Algorithm 1 has, in expectation, both a single adjustment and round, in both the synchronous and asynchronous models.*

4. A CONSTANT BROADCAST IMPLEMENTATION

Theorem 1 promises that the expected number of nodes that need to change their output according to our template algorithm is 1. However, a direct implementation of the template in Algorithm 1 in a dynamic distributed setting may require a much larger broadcast complexity because it may be the case that a node needs to change its state several times until the MIS invariant holds at all nodes. This is because a node can be in more than a single set S_i , as discussed in the previous section. In such a case, despite the fact that the expected number of nodes in S is a constant, it may be that the expected number of state changes is much larger. Specifically, in a naive implementation, the number of broadcasts may be as large as $|S|^2$. Hence, although $\mathbb{E}[|S|] = 1$, the expected number of broadcasts may be as large as n .

We thus take a different approach for implementing the template in Algorithm 1, in the *synchronous* setting, where each node waits until it knows the maximal i for which it belongs to S_i , and changes its state only once. This allows to obtain, for almost all of the possible topology changes a constant broadcast complexity at the cost of a constant, rather than single, round complexity.

In order to implement the random permutation π we assume each node $v \in V$ has a uniformly random and independent ID $\ell_v \in [0, 1]$. We will maintain the property that each node has knowledge of its ℓ value and those of its neighbors. We describe our algorithm in Algorithm 2. This directly applies to the following topology changes: edge-insertion, graceful-edge-deletion, abrupt-edge-deletion, graceful-node-deletion and node-unmuting. An extension of the analysis is provided in the full version [16] for the case of an abrupt node deletion, and a slight modification is provided there for the case of node-insertion. The following summarizes the guarantees of our implementation, and is proven in Lemma 9 and in the full version [16].

THEOREM 7. *There is a complete fully dynamic distributed MIS algorithm which requires in expectation a single adjustment and $O(1)$ rounds for all topology changes. For edge insertions and deletions, graceful node deletion, and node unmuting, the algorithm requires $O(1)$ broadcasts, for an abrupt deletion of a node v^* it requires $O(\min\{\log(n), d(v^*)\})$ broadcasts, and for an insertion of a node v^* it requires $O(d(v^*))$ broadcasts, in expectation.*

In the algorithm a node may be in one of four states: M for an MIS node, \bar{M} for a non-MIS node, C for a node that may need to change from M to \bar{M} or vice-versa, and R for a node that is ready to change. We will sometimes abuse notation and consider a state as the set of nodes which are in that state. Our goal is to maintain the MIS invariant.

Algorithm 2 MIS Algorithm for node v

- 1: $v \in M$: If some $u \in I_\pi(v)$ changes to state C , change state to C .
 - 2: $v \in \bar{M}$: If some $u \in I_\pi(v)$ changes to state C and all other $w \in I_\pi(v)$ are not in M , change state to C .
 - 3: $v \in C$: If (1) all neighbors u with $\pi(v) < \pi(u)$ are not in state C and (2) v changed to state C at least 2 rounds ago, change state to R .
 - 4: $v \in R$: If all $u \in I_\pi(v)$ are in states \bar{M} or M , change state to M if all $u \in I_\pi(v)$ are in \bar{M} , and change state to \bar{M} otherwise.
-

Any change of state of a node is followed by a broadcast of the new state to all of its neighbors. We now define our implementation as a sequence of state changes. When a topology change occurs at node v^* , if the MIS invariant still holds then v^* does not change its state and algorithm consists of doing nothing. Otherwise, v^* changes its state to C .

From states M or \bar{M} , a node changes to state C when it discovers it is in the set S of influenced nodes, as defined in Equation (1). From state C , a node v changes to state R when (1) none of its neighbors u for which $\pi(v) < \pi(u)$ are in state C and (2) v changed its state to C at least two rounds ago. Finally, from state R a node v returns to states M or \bar{M} when all of its neighbors u for which $\pi(u) < \pi(v)$ are in states M or \bar{M} . In order to bound the complexity of the algorithm we first show that every node can change from state R to either M or \bar{M} at most once.

LEMMA 8. *In Algorithm 2, a node u changes its state from R to another state at most once.*

PROOF. First, note that every $u \notin S$ never changes its state. Consider a node u changing its state from R to either M or \bar{M} . Since u changes from state R , if $u \neq v^*$ then it must have a neighbor $w \in I_\pi(u)$ that was in state C , changed to state R and then changed to M or \bar{M} . It follows that v^* must be the first node to change its state from R to M or \bar{M} . This event occurs only when all neighbors u of v^* are not in C , which in turn can happen only when all neighbors of each such u with higher π value have changed from C to R at least once. But, since no node could have changed its state from R to another state before v^* has done so, we have that when v^* changes its state from R to another, all $u \in S$ are in state R .

In particular, we have that at the round of the first change of a node from R to another state, there are no nodes in state

C . Since a node can only change to state C due to a neighbor at state C we have that any node changing its state from R to M or \bar{M} will not change its state again, thus proving our claim. \square

LEMMA 9. *For edge-insertion, graceful-edge-deletion, abrupt-edge-deletion, graceful-node-deletion and node-unmuting, Algorithm 2 requires in expectation a single adjustment, $O(1)$ rounds, and $O(1)$ broadcasts.*

PROOF. Since only nodes in S can change their outputs, the number of adjustments is bounded by $|S|$, and hence is 1 in expectation, by Theorem 1. According to Lemma 8, if a node changes its state then it does so exactly three times. First it changes from either M or \bar{M} to C , then it changes to R , and finally it changes to either M or \bar{M} again. Since only nodes in S change their states and since the round and broadcast complexities are clearly bounded by the number of state changes plus 1 (due to the forced waiting round before changing from C to R), the claim follows. \square

The cases of node insertion and abrupt node deletion are addressed in the full version [16].

5. DISCUSSION

This paper studies computing an MIS in a distributed dynamic setting. The strength of our analysis lies in obtaining that for an algorithm that simulates the sequential random greedy algorithm, the size of the set of nodes that need to change their output is in expectation 1. This brings the locality of the fundamental MIS problem to its most powerful setting.

We are able to extend our analysis to show that our algorithm can cope with any number k of topology changes that overlap. The required time for updating the MIS naturally depends on k , but remains constant if k is constant. For clarity of presentation, we defer this extension to a future extended version of this paper.

The expected update time required for our algorithm is with respect to a random order of nodes. One might argue that once we have a network in a fixed state then this order is given, and so an interesting question is whether our algorithm can perform well for this particular given order, even with possibly relaxing the topology changes to be random. However, a simple counter-example shows that this is not the case: consider a path with the nodes ordered from left to right (v_1, \dots, v_n) . The greedy MIS consists of all v_i for odd i . Any topology change of removing v_i causes $n - i$ nodes to change their state. This means that in expectation we will suffer a linear number of changes.

Our work sets the ground for more research in this crucial setting. First, there are many additional problems that can be addressed in the dynamic distributed setting, especially in the synchronous case. We believe that our contribution can find applications in solving many additional dynamic distributed tasks.

A major open question is whether our techniques can be adapted to sequential dynamic graph algorithms, which constitutes a major area of research in the sequential setting [13, 14, 19, 25–27, 31, 33–35, 38, 55, 56, 58, 59]. A formal definition and description of typical problems can be found in, e.g., [18]. Notice that our algorithms are *fully dynamic*, which means that they handle both insertions and deletions (of edges and nodes). Although our template for finding

an MIS can be easily implemented in a sequential dynamic setting, it would come with a cost of at least $O(\Delta)$ for the update complexity in a direct implementation. This is because we would have to access neighbors of the set of nodes analyzed in Theorem 1. Our distributed implementation avoids this by having them simply not respond since they do not need to change their output, and hence they do not contribute to the communication. Nevertheless, we believe that our approach may be useful for designing an MIS algorithm for the dynamic sequential setting, and leave this for future research.

Acknowledgements.

We thank Reuven Bar-Yehuda, Mohsen Ghaffari, and Boaz Patt-Shamir for interesting discussions.

6. REFERENCES

- [1] N. Ailon, M. Charikar, and A. Newman. Aggregating inconsistent information: Ranking and clustering. *J. ACM*, 55(5), 2008.
- [2] N. Alon, L. Babai, and A. Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms*, 7(4):567–583, 1986.
- [3] N. Alon, R. Rubinfeld, S. Vardi, and N. Xie. Space-efficient local computation algorithms. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1132–1139, 2012.
- [4] C. Avin, M. Koucký, and Z. Lotker. How to explore a fast-changing world (cover time of a simple random walk on evolving graphs). In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming, ICALP*, pages 121–132, 2008.
- [5] B. Awerbuch, I. Cidon, and S. Kutten. Communication-optimal maintenance of replicated information. In *31st Annual Symposium on Foundations of Computer Science, FOCS*, pages 492–502, 1990.
- [6] B. Awerbuch and M. Sipser. Dynamic networks are as fast as static networks (preliminary version). In *29th Annual Symposium on Foundations of Computer Science, FOCS*, pages 206–220, 1988.
- [7] B. Awerbuch and G. Varghese. Distributed program checking: a paradigm for building self-stabilizing distributed protocols (extended abstract). In *32nd Annual Symposium on Foundations of Computer Science, FOCS*, pages 258–267, 1991.
- [8] Y. Azar, S. Kutten, and B. Patt-Shamir. Distributed error confinement. *ACM Trans. Algorithms*, 6(3):48:1–48:23, July 2010.
- [9] L. Barenboim and M. Elkin. *Distributed Graph Coloring: Fundamentals and Recent Developments*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2013.
- [10] L. Barenboim, M. Elkin, and F. Kuhn. Distributed $(\Delta+1)$ -coloring in linear (in Δ) time. *SIAM J. Comput.*, 43(1):72–95, 2014.
- [11] L. Barenboim, M. Elkin, S. Pettie, and J. Schneider. The locality of distributed symmetry breaking. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 321–330, 2012.
- [12] S. Baswana, S. Khurana, and S. Sarkar. Fully dynamic randomized algorithms for graph spanners. *ACM Transactions on Algorithms*, 8(4):35, 2012.
- [13] A. Bernstein and C. Stein. Fully dynamic matching in bipartite graphs. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming ICALP*, pages 167–179, 2015.
- [14] S. Bhattacharya, M. Henzinger, and G. F. Italiano. Deterministic fully dynamic data structures for vertex cover and matching. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 785–804, 2015.
- [15] G. E. Blelloch, J. T. Fineman, and J. Shun. Greedy sequential maximal independent set and matching are parallel on average. In *24th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA*, pages 308–317, 2012.
- [16] K. Censor-Hillel, E. Harnamty, and Z. S. Karnin. Optimal dynamic distributed MIS. *CoRR*, abs/1507.04330, 2015.
- [17] S. Cicerone, G. D. Stefano, D. Frigioni, and U. Nanni. A fully dynamic algorithm for distributed shortest paths. *Theor. Comput. Sci.*, 297(1-3):83–102, 2003.
- [18] C. Demetrescu, D. Eppstein, Z. Galil, and G. F. Italiano. Dynamic graph algorithms. In M. J. Atallah and M. Blanton, editors, *Algorithms and Theory of Computation Handbook*, chapter 9. Chapman & Hall/CRC, 2010.
- [19] C. Demetrescu and G. F. Italiano. Fully dynamic transitive closure: Breaking through the $o(n^2)$ barrier. In *41st Annual Symposium on Foundations of Computer Science, FOCS*, pages 381–389, 2000.
- [20] E. W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Commun. ACM*, 17(11):643–644, 1974.
- [21] S. Dolev. *Self-Stabilization*. MIT Press, 2000.
- [22] S. Dolev and T. Herman. Superstabilizing protocols for dynamic distributed systems. *Chicago J. Theor. Comput. Sci.*, 1997, 1997.
- [23] S. Dolev and N. Tzachar. Empire of colonies: Self-stabilizing and self-organizing distributed algorithm. *Theor. Comput. Sci.*, 410(6-7):514–532, 2009.
- [24] M. Elkin. A near-optimal distributed fully dynamic algorithm for maintaining sparse spanners. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Principles of Distributed Computing, PODC*, pages 185–194, 2007.
- [25] D. Eppstein, Z. Galil, G. F. Italiano, and A. Nissenzweig. Sparsification - a technique for speeding up dynamic graph algorithms. *J. ACM*, 44(5):669–696, 1997.
- [26] S. Even and Y. Shiloach. An on-line edge-deletion problem. *J. ACM*, 28(1):1–4, 1981.
- [27] G. N. Frederickson. Data structures for on-line updating of minimum spanning trees, with applications. *SIAM J. Comput.*, 14(4):781–798, 1985.
- [28] M. Ghaffari. An improved distributed algorithm for maximal independent set. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 270–277, 2016.

- [29] N. Guellati and H. Kheddouci. A survey on self-stabilizing algorithms for independence, domination, coloring, and matching in graphs. *J. Parallel Distrib. Comput.*, 70(4):406–415, 2010.
- [30] T. P. Hayes, J. Saia, and A. Trehan. The forgiving graph: a distributed data structure for low stretch under adversarial attack. *Distributed Computing*, 25(4):261–278, 2012.
- [31] M. Henzinger, S. Krumm, and D. Nanongkai. Dynamic approximate all-pairs shortest paths: Breaking the $O(mn)$ barrier and derandomization. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 538–547, 2013.
- [32] M. Henzinger, S. Krumm, and D. Nanongkai. Sublinear-time maintenance of breadth-first spanning tree in partially dynamic networks. In *Proceedings of the 40th International Colloquium on Automata, Languages, and Programming ICALP*, pages 607–619, 2013.
- [33] M. R. Henzinger and V. King. Fully dynamic biconnectivity and transitive closure. In *36th Annual Symposium on Foundations of Computer Science FOCS*, pages 664–672, 1995.
- [34] M. R. Henzinger and V. King. Randomized fully dynamic graph algorithms with polylogarithmic time per operation. *J. ACM*, 46(4):502–516, 1999.
- [35] J. Holm, K. de Lichtenberg, and M. Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 79–89, 1998.
- [36] A. Israeli and A. Itai. A fast and simple randomized parallel algorithm for maximal matching. *Inf. Process. Lett.*, 22(2):77–80, 1986.
- [37] G. F. Italiano. Distributed algorithms for updating shortest paths (extended abstract). In *Proceedings of the 5th International Workshop on Distributed Algorithms, WDAG*, pages 200–211, 1991.
- [38] B. M. Kapron, V. King, and B. Mountjoy. Dynamic graph connectivity in polylogarithmic worst case time. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1131–1142, 2013.
- [39] M. König and R. Wattenhofer. On local fixing. In *Proceedings of the 17th International Conference on Principles of Distributed Systems, OPODIS*, pages 191–205, 2013.
- [40] A. Korman. Improved compact routing schemes for dynamic trees. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on Principles of Distributed Computing, PODC*, pages 185–194, 2008.
- [41] A. Korman and D. Peleg. Labeling schemes for weighted dynamic trees. In *Proceedings of the 30th International Colloquium on Automata, Languages and Programming ICALP*, pages 369–383, 2003.
- [42] F. Kuhn, N. A. Lynch, and R. Oshman. Distributed computation in dynamic networks. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC*, pages 513–522, 2010.
- [43] F. Kuhn, T. Moscibroda, and R. Wattenhofer. What cannot be computed locally! In *Proceedings of the Twenty-Third Annual ACM Symposium on Principles of Distributed Computing, PODC*, pages 300–309, 2004.
- [44] S. Kutten and D. Peleg. Fault-local distributed mending. *J. Algorithms*, 30(1):144–165, 1999.
- [45] S. Kutten and D. Peleg. Tight fault locality. *SIAM J. Comput.*, 30(1):247–268, 2000.
- [46] C. Lenzen, J. Suomela, and R. Wattenhofer. Local algorithms: Self-stabilization on speed. In *Proceedings of the 11th International Symposium on Stabilization, Safety, and Security of Distributed Systems SSS*, pages 17–34, 2009.
- [47] R. Levi, R. Rubinfeld, and A. Yodanis. Local computation algorithms for graphs of non-constant degrees. *CoRR*, abs/1502.04022, 2015.
- [48] N. Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
- [49] M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15(4):1036–1053, 1986.
- [50] Y. Mansour, A. Rubinfeld, S. Vardi, and N. Xie. Converting online algorithms to local computation algorithms. In *Proceedings of the 39th International Colloquium on Automata, Languages, and Programming ICALP*, pages 653–664, 2012.
- [51] Y. Métivier, J. M. Robson, N. Saheb-Djahromi, and A. Zemmari. An optimal bit complexity randomized distributed MIS algorithm. *Distributed Computing*, 23(5-6):331–340, 2011.
- [52] H. N. Nguyen and K. Onak. Constant-time approximation algorithms via local improvements. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 327–336, 2008.
- [53] G. Pandurangan and A. Trehan. Xheal: a localized self-healing algorithm using expanders. *Distributed Computing*, 27(1):39–54, 2014.
- [54] D. Peleg. *Distributed Computing: A Locality-sensitive Approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [55] L. Roditty and U. Zwick. On dynamic shortest paths problems. *Algorithmica*, 61(2):389–401, 2011.
- [56] L. Roditty and U. Zwick. Dynamic approximate all-pairs shortest paths in undirected graphs. *SIAM J. Comput.*, 41(3):670–683, 2012.
- [57] M. Schneider. Self-stabilization. *ACM Comput. Surv.*, 25(1):45–67, 1993.
- [58] R. E. Tarjan. Efficiency of a good but not linear set union algorithm. *J. ACM*, 22(2):215–225, 1975.
- [59] M. Thorup. Near-optimal fully-dynamic graph connectivity. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing STOC*, pages 343–350, 2000.
- [60] V. Turau. Linear self-stabilizing algorithms for the independent and dominating set problems using an unfair distributed scheduler. *Inf. Process. Lett.*, 103(3):88–93, 2007.
- [61] Y. Yoshida, M. Yamamoto, and H. Ito. Improved constant-time approximation algorithms for maximum matchings and other optimization problems. *SIAM J. Comput.*, 41(4):1074–1093, 2012.