

# Dynamic Networks of Finite State Machines

Yuval Emek<sup>1</sup> and Jara Uitto<sup>2</sup>

<sup>1</sup>Technion, Israel

<sup>2</sup>Comerge AG, Switzerland

## Abstract

Like distributed systems, biological multicellular processes are subject to dynamic changes and a biological system will not pass the survival-of-the-fittest test unless it exhibits certain features that enable fast recovery from these changes. In most cases, the types of dynamic changes a biological process may experience and its desired recovery features differ from those traditionally studied in the distributed computing literature. In particular, a question seldomly asked in the context of distributed digital systems and that is crucial in the context of biological cellular networks, is whether the system can keep the changing components *confined* so that only nodes in their vicinity may be affected by the changes, but nodes sufficiently far away from any changing component remain unaffected.

Based on this notion of confinement, we propose a new metric for measuring the dynamic changes recovery performance in distributed network algorithms operating under the *Stone Age* model (Emek & Wattenhofer, PODC 2013), where the class of dynamic topology changes we consider includes inserting/deleting an edge, deleting a node together with its incident edges, and inserting a new isolated node. Our main technical contribution is a distributed algorithm for maximal independent set (MIS) in synchronous networks subject to these topology changes that performs well in terms of the aforementioned new metric. Specifically, our algorithm guarantees that nodes which do not experience a topology change in their immediate vicinity are not affected and that all surviving nodes (including the affected ones) perform  $\mathcal{O}((C + 1) \log^2 n)$  computationally-meaningful steps, where  $C$  is the number of topology changes; in other words, each surviving node performs  $\mathcal{O}(\log^2 n)$  steps when amortized over the number of topology changes. This is accompanied by a simple example demonstrating that the linear dependency on  $C$  cannot be avoided.

# 1 Introduction

The biological form of close-range (juxtacrine) message passing relies on designated messenger molecules that bind to crossmembrane receptors in neighboring cells; this binding action triggers a signaling cascade that eventually affects gene expression, thus modifying the neighboring cells' states. This mechanism should feel familiar to members of the distributed computing community as it resembles the message passing schemes of distributed digital systems. In contrast to nodes in distributed digital systems, however, biological cells are not believed to be Turing complete, rather each biological cell is pretty limited in computation as well as communication. In attempt to cope with these differences, Emek and Wattenhofer [13] introduced the *Stone Age* model of distributed computing (a.k.a. networked finite state machines), where each node in the network is a very weak computational unit with limited communication capabilities and showed that several fundamental distributed computing problems can be solved efficiently under this model.

An important topic left outside the scope of [13] is that of *dynamic topology changes*. Just like distributed digital systems, biological systems may experience local changes and the ability of the system to recover from these changes is crucial to its survival. However, the desired recovery features in biological cellular networks typically differ from those traditionally studied in the distributed computing literature. In particular, a major issue in the context of biological cellular networks, that is rarely addressed in the study of distributed digital systems, is that of *confining* the topology changes: while nodes in the immediate vicinity of a topology change are doomed to be affected by it (hopefully, to a bounded extent), isolating the nodes sufficiently far away from any topology change so that their operation remains unaffected, is often critical. Indeed, a biological multicellular system with limited energy resources cannot afford every cell division (or death) to have far reaching effects on the cellular network.

In this paper, we make a step towards bringing the models from computer science closer to biology by extending the Stone Age model to accommodate four types of dynamic topology changes: (1) deleting an existing edge; (2) inserting a new edge; (3) deleting an existing node together with its incident edges; and (4) inserting a new isolated node. We also introduce a new method for measuring the performance of a network in recovering from these types of topology changes that takes into account the aforementioned confinement property. This new method measures the number of “computationally-meaningful” steps made by the individual nodes, which are essentially all steps in which the node participates (in the weakest possible sense) in the global computational process. An algorithm is said to be *effectively confining* if (i) the runtime of the nodes that are not adjacent to any topology change is  $\log^{\mathcal{O}(1)} n$ ; and (ii) the global runtime (including all surviving nodes) is  $(C + 1) \log^{\mathcal{O}(1)} n$ , where  $C$  is the number of topology changes throughout the execution. In other words, the global runtime is  $\log^{\mathcal{O}(1)} n$  when amortized over the number of changes.

Following that, we turn our attention to the extensively studied *maximal independent set (MIS)* problem and design a randomized effectively confining algorithm for it under the Stone Age model extended to dynamic topology changes. This is achieved by carefully augmenting the MIS algorithm

introduced in [13] with new components, tailored to ensure fast recovery from topology changes. Being a first step in the study of recovery from dynamic changes under the Stone Age model, our algorithm assumes a *synchronous* environment and it remains an open question whether this assumption can be lifted. Nevertheless, this assumption is justified by the findings of Fisher et al. [14] that model cellular networks as being subject to a *bounded asynchrony* scheduler, which is equivalent to a synchronous environment from an algorithmic perspective.

**Paper’s Organization.** An extension of the Stone Age model of [13] to dynamic topology changes is presented in Sec. 2 together with our new method for evaluating the recovery performance of distributed algorithms. In Sec. 3, we describe the details of our MIS algorithm. Then, in Sec. 4.1, we show that each node not affected by a topology change will reach an output state in  $\mathcal{O}(\log^2 n)$  rounds. In Sec. 4.2 and 4.3, we finish the analysis of our MIS algorithm by establishing an  $\mathcal{O}((C + 1)\log^2 n)$  upper bound on the global runtime. The runtime of the new MIS algorithm is shown to be near-optimal in Sec. 5 by proving that the global runtime of any algorithm is  $\Omega(C)$ . In Apx. A, we show that the runtime of any node  $u$  can be further bounded by  $\mathcal{O}((C_u + 1)\log^2 n)$ , where  $C_u$  is the number of topology changes that occurred within  $\mathcal{O}(\log n)$  hops from  $u$ . Some interesting open questions are listed in Sec. 6.

**Related work.** The standard model for a network of communicating devices is the message passing model [28, 34]. There are several variants of this model, where the power of the network has been weakened. Perhaps the best-known variant of the message passing model is the congest model, where the message size is limited to size logarithmic in the size of the input graph [34]. A step to weaken the model further is to consider interference of messages, i.e., a node only hears a message if it receives a single message per round — cf. the *radio network model* [9]. In the *beeping model* [11, 15], the communication capabilities are reduced further by only allowing to send beeps that do not carry information, where a listening node cannot distinguish between a single beep and multiple beeps transmitted by its neighbors.

The models mentioned above focus on limiting the communication but not the computation, i.e., the nodes are assumed to be strong enough to perform unlimited (local) Turing computations in each round. Networks of nodes weaker than Turing machines have been extensively studied in the context of *cellular automata* [16, 39, 41]. While the cellular automata model typically considers the deployment of finite state machines in highly regular graph structures such as the grid, the question of modeling cellular automata in arbitrary graphs was tackled by Marr and Hütt in [31], where a node changes its binary state according to the densities of its neighboring states. Another extensively studied model for distributed computing in a network of finite state machines is *population protocols* [4] (see also [5, 32]), where the nodes communicate through a series of pairwise rendezvous. Refer to [13] for a comprehensive account of the similarities and differences between the Stone Age model and the models of cellular automata and population protocols.

Distributed computing in dynamic networks has been extensively studied [6, 19, 23, 26, 40]. A

classic result by Awerbuch and Sipser states that under the message passing model, any algorithm designed to run in a static network can be transformed into an algorithm that runs in a dynamic network with only a constant multiplicative runtime overhead [7]. However, the transformation of Awerbuch and Sipser requires storing the whole execution history and sending it around the network, which is not possible under the Stone Age model. Some dynamic network papers rely on the assumption that the topology changes are spaced in time so that the system has an opportunity to recover before another change occurs [20,21,30]. The current paper does not make this assumption.

The *maximal independent set (MIS)* problem has a long history in the context of distributed algorithms [3,10,28,29,33,38]. Arguably the most significant breakthrough in the study of message passing MIS algorithms was the  $\mathcal{O}(\log n)$  algorithm of Luby [29] (developed independently also by Alon et al. [3]). Later, Barenboim et al. [8] showed an upper bound of  $2^{\mathcal{O}(\sqrt{\log \log n})}$  in the case of polylogarithmic maximum degree and an  $\mathcal{O}(\log \Delta + \sqrt{\log n})$  bound for general graphs. For growth bounded graphs, it was shown by Schneider et al. [35] that MIS can be computed in  $\mathcal{O}(\log^* n)$  time. In a recent work, Ghaffari studied the *local* complexity of computing an MIS, where the time complexity is measured from the perspective of a single node, instead of the whole network [17]. In a similar spirit, we provide, in addition to a global runtime bound, a runtime analysis from the perspective of a single node.

In the radio networks realm, with  $f$  channels, an MIS can be computed in  $\Theta(\log^2 n/f) + \tilde{\mathcal{O}}(\log n)$  time [12], where  $\tilde{\mathcal{O}}$  hides factors polynomial in  $\log \log n$ . The MIS problem was extensively studied also under the beeping model [1,2,36]. Afek et al. [1] proved that if the nodes are provided with an upper bound on  $n$  or if they are given a common sense of time, then the MIS problem can be solved in  $\mathcal{O}(\log^{\mathcal{O}(1)} n)$  time. This was improved to  $\mathcal{O}(\log n)$  by Scott et al. [36] assuming that the nodes can detect sender collision

On the negative side, the seminal work of Linial [28] provides a runtime lower bound of  $\Omega(\log^* n)$  for computing an MIS in an  $n$ -node ring under the message passing model [28]. Kuhn et al. [22] established a stronger lower bound for general graphs stating that it takes  $\Omega\left(\sqrt{\frac{\log n}{\log \log n}}\right) + \Omega\left(\frac{\log \Delta}{\log \log \Delta}\right)$  rounds to compute an MIS, where  $\Delta$  is the maximum degree of the input graph. For uniform algorithms in radio networks (and therefore, also for the beeping model) with asynchronous wake up schedule, there exists a lower bound of  $\Omega(\sqrt{n/\log n})$  communication rounds [1].

Containing faults within a small radius of the faulty node has been studied in the context of *self-stabilization* [18]. An elegant MIS algorithm was developed under the assumption that the activation times of the nodes are controlled by a *central daemon* who activates the nodes one at a time [27,37]. In contrast, we follow the common assumption that all nodes are activated simultaneously in each round. In the self-stabilization realm, the performance of an algorithm is typically measured as a function of some network parameter, such as the size of the network or the maximum degree, whereas in the current paper, the performance depends also on the number of failures.

With respect to the performance evaluation, perhaps the works closest to ours are by Kutten

and Peleg [24,25], where the concepts of *mending algorithms* and *tight fault locality* are introduced. The idea behind a fault local mending algorithm is to be able to recover a legal state of the network after a fault occurs, measuring the performance in terms of the number of faults. The term tight fault locality reflects the property that an algorithm running in time  $\mathcal{O}(T(n))$  without faults is able to mend the network in time  $\mathcal{O}(T(F))$ , where  $F$  denotes the number of faults. The algorithm of Kutten and Peleg recovers an MIS in time  $\mathcal{O}(\log F)$ , but they use techniques that require nodes to count beyond constant numbers, which is not possible in the Stone Age model. Furthermore, they consider transient faults, whereas we consider permanent changes in the network topology.

## 2 Model

Consider some network represented by an undirected graph  $G = (V, E)$ , where the nodes in  $V$  correspond to the network's computational units and the edges represent bidirectional communication channels. Adopting the (static) Stone Age model of [13], each node  $v \in V$  runs an algorithm captured by the 8-tuple  $\Pi = \langle Q, q_0, q_{yes}, q_{no}, \Sigma, \sigma_0, b, \delta \rangle$ , where  $Q$  is a fixed set of *states*;  $q_0 \in Q$  is the *initial state* in which all nodes reside when they wake up for the first time;  $q_{yes}$  and  $q_{no}$  are the *output states*, where the former (resp., latter) represents membership (resp., non-membership) in the output MIS;  $\Sigma$  is a fixed *communication alphabet*;  $\sigma_0 \in \Sigma$  is the *initial letter*;  $b \in \mathbb{Z}_{>0}$  is a *bounding parameter*; and  $\delta : Q \times \{0, 1, \dots, b\}^{|\Sigma|} \rightarrow 2^{Q \times (\Sigma \cup \{\varepsilon\})}$  is the *transition function*. For convenience, we sometimes denote a state transition from  $q$  to  $q'$  by  $(q \rightarrow q')$  and omit the rules associated with this transition from the notation.<sup>1</sup>

Node  $v$  communicates with its neighbors by transmitting messages that consist of a single letter  $\sigma \in \Sigma$  such that the same letter  $\sigma$  is sent to all neighbors. It is assumed that  $v$  holds a port  $\phi_u(v)$  for each neighbor  $u$  of  $v$  in which the last message (a letter in  $\Sigma$ ) received from  $u$  is stored. Transmitting the designated empty symbol  $\varepsilon$  corresponds to the case where  $u$  does not transmit any message. In other words, when node  $u$  transmits the  $\varepsilon$  letter, the letters in ports  $\phi_u(v)$ , for all neighbors  $v$  of  $u$ , remain unchanged. In the beginning of the execution, all ports contain the initial letter  $\sigma_0$ .

The execution of the algorithm proceeds in discrete synchronous rounds indexed by the positive integers. In each round  $r \in \mathbb{Z}_{>0}$ , node  $v$  is in some state  $q \in Q$ . Let  $\sharp(\sigma)$  be the number of appearances of the letter  $\sigma \in \Sigma$  in  $v$ 's ports in round  $r$  and let  $\langle \min\{\sharp(\sigma), b\} \rangle_{\sigma \in \Sigma}$  be a  $\Sigma$ -indexed vector whose  $\sigma$ -entry is set to the minimum between  $\sharp(\sigma)$  and the bounding parameter  $b$ . Then the state  $q'$  in which  $v$  resides in round  $r + 1$  and the message  $\sigma'$  that  $v$  sends in round  $r$  (appears in the corresponding ports of  $v$ 's neighbors in round  $r + 1$  unless  $\sigma' = \varepsilon$ ) are chosen uniformly at

---

<sup>1</sup>For simplicity, the definition presented here is already tailored to the MIS problem. For a general problem  $P$ , the initial state  $q_0$  can be replaced with a set of initial states, representing the possible inputs of  $P$ , and we can have as many output states as needed to represent the possible outputs of  $P$  (captured in the MIS case by the two states  $q_{yes}$  and  $q_{no}$ ).

random among the pairs in

$$\delta(q, \langle \min\{\#\!(\sigma), b\} \rangle_{\sigma \in \Sigma}) \subseteq Q \times (\Sigma \cup \{\varepsilon\}).$$

This means that  $v$  tosses an unbiased die (with  $|\delta(q, \langle \min\{\#\!(\sigma), b\} \rangle_{\sigma \in \Sigma})|$  faces) when deciding on  $q'$  and  $\sigma'$ .

To ensure that our state transitions are well defined, we require that  $|\delta(q, x)| \geq 1$  for all  $q \in Q$  and  $x \in \{0, 1, \dots, b\}^{|\Sigma|}$  and say that a state transition is *deterministic* if  $|\delta(q, x)| = 1$ .

**Topology changes.** In contrast to the model presented in [13], our network model supports dynamic *topology changes* that belong to the following four classes:

1. *Edge deletion:* Remove a selected edge  $e = \{u, v\}$  from the current graph. The corresponding ports  $\phi_u(v)$  and  $\phi_v(u)$  are removed and the messages stored in them are erased.
2. *Edge insertion:* Add an edge connecting nodes  $u, v \in V$  to the current graph. New ports  $\phi_u(v)$  and  $\phi_v(u)$  are introduced storing the initial letter  $\sigma_0 \in \Sigma$ .
3. *Node deletion:* Remove a selected node  $v \in V$  from the current graph with all its incident edges. The corresponding ports  $\phi_v(u)$  of all  $v$ 's neighbors  $u$  are eliminated and the messages stored in them are erased.
4. *Node insertion:* Add a new isolated (i.e., with no neighbors) node to the current graph. Initially, the node resides in the initial state  $q_0$ .

We assume that the schedule of these topology changes is controlled by an *oblivious adversary*. Formally, the strategy of the adversary associates a (possibly empty) set of topology changes with each round  $r \in \mathbb{Z}_{>0}$  of the execution so that the total number of changes throughout the execution, denoted by  $C$ , is finite. This strategy may depend on the algorithm  $\Pi$ , but not on the random choices made during the execution.

To be precise, each round is divided into 4 successive steps as follows: (i) messages arrive at their destination ports; (ii) topology changes occur; (iii) the transition function is applied; and (iv) messages are transmitted. In particular, a message transmitted by node  $v$  in round  $r$  will not be read by node  $u$  in round  $r + 1$  if edge  $\{u, v\}$  is inserted or deleted in round  $r + 1$ . It is convenient to define the *adversarial graph sequence*  $\mathcal{G} = G_1, G_2, \dots$  so that  $G_1$  is the initial graph and  $G_{r+1}$  is the graph obtained from  $G_r$  by applying to it the topology changes scheduled for round  $r$ . By definition,  $\mathcal{G}$  is fully determined by the initial graph and the adversarial policy (and vice versa). The requirement that  $C$  is finite implies, in particular, that  $\mathcal{G}$  admits an infinite suffix of identical graphs. Let  $n$  be the largest number of nodes that co-existed in the same round, i.e.,  $n = \max_r |V(G_r)|$ , where  $V(G_r)$  is the node set of  $G_r$ .

**Correctness.** We say that node  $u$  resides in state  $q$  in round  $r$  if the state of  $u$  is  $q$  at the beginning of round  $r$ . The algorithm is said to be in an *output configuration* in round  $r$  if every node  $u \in G_r$  resides in an output state ( $q_{yes}$  or  $q_{no}$ ). The output configuration is said to be *correct* if

the states of the nodes (treated as their output) correspond to a valid MIS of the current graph  $G_r$ . An algorithm  $\Pi$  is said to be *correct* if the following conditions are satisfied for every adversarial graph sequence: (C1) If  $\Pi$  is in a non-output configuration in round  $r$ , then it will move to an output configuration in some (finite) round  $r' > r$  w.p. 1.<sup>2</sup> (C2) If  $\Pi$  is in an output configuration in round  $r$ , then this output configuration is correct (with respect to  $G_r$ ). (C3) If  $\Pi$  is in an output configuration in round  $r$  and  $G_{r+1} = G_r$ , then  $\Pi$  remains in the same output configuration in round  $r + 1$ .

**Restrictions on the Output States.** Under the model of [13], the nodes are not allowed to change their output once they have entered an output state. On the other hand, this model allows for multiple output states that correspond to the same problem output (“yes” and “no” in the MIS case). In other words, it is required that the output states that correspond to each problem output form a sink of the transition function. Since our model accommodates dynamic topology changes which might turn a correct output configuration into an incorrect one, we lift this restriction: our model allows transitions from an output state to a non-output state, thus providing the algorithm designer with the possibility to escape output configurations that become incorrect. Nevertheless, to prevent nodes in an output state from taking any *meaningful* part in the computation process, we introduce the following new (with respect to [13]) restrictions: (1) each possible output is represented by a unique output state; (2) a transition from an output state to itself is never accompanied by a letter transmission (i.e., it transmits  $\varepsilon$ ); (3) all transitions originating from an output state must be deterministic; and (4) a transition from an output state must lead either to itself or to a non-output state.

**Runtime.** Fix some adversarial graph sequence  $\mathcal{G} = G_1, G_2, \dots$  and let  $\eta$  be an execution of a correct algorithm  $\Pi$  on  $\mathcal{G}$  (determined by the random choices of  $\Pi$ ). Round  $r$  in  $\eta$  is said to be *silent* if  $\Pi$  is in a correct output configuration and no topology change occurs in round  $r$ . The *global runtime* of  $\eta$  is the number of non-silent rounds. Node  $v$  is said to be *active* in round  $r$  of  $\eta$  if it resides in a non-output state, i.e., some state in  $Q - \{q_{yes}, q_{no}\}$ . The *(local) runtime* of  $v$  in  $\eta$  is the number of rounds in which  $v$  is active.

Let  $N_r(v)$  be the (inclusive) neighborhood of  $v$  and let  $E_r(v)$  be the edges incident to  $v$  in  $G_r$ . Node  $v$  is said to be *affected* under  $\mathcal{G}$  if either  $v$  or one of its neighbors experienced an edge insertion/deletion, that is, there exists some round  $r$  and some node  $u \in N_r(v)$  such that  $E_r(u)$  is not identical to  $E_{r+1}(u)$ . Algorithm  $\Pi$  is said to be *effectively confining* if the following conditions are satisfied for every adversarial graph sequence  $\mathcal{G}$ : (1) the expected maximum local runtime of the non-affected nodes is  $\log^{\mathcal{O}(1)} n$ ; and (2) the expected global runtime is  $(C + 1) \cdot \log^{\mathcal{O}(1)} n$ , namely,  $\log^{\mathcal{O}(1)} n$  when amortized over the number of topology changes. Notice that the bound on the global runtime directly implies the same bound on the local runtime of any affected node.

---

<sup>2</sup>Throughout, we use w.p. to abbreviate “with probability” and w.h.p. to abbreviate “with high probability”, i.e., with probability  $n^{-c}$  for any constant  $c$ .

### 3 An MIS Algorithm

Our main goal is to design an algorithm under the Stone Age model for the maximal independent set (MIS) problem that is able to tolerate topology changes. For a graph  $G = (V, E)$ , a set of nodes  $I \subseteq V$  is independent if for all  $u, v \in I$ ,  $\{u, v\} \notin E$ . An independent set  $I$  is maximal if there is no other set  $I' \subseteq V$  such that  $I \subset I'$  and  $I'$  is independent.

Following the terminology introduced in the model section, we show that our algorithm is effectively confining. In Sec. 5, we provide a straightforward lower bound example that shows that under our model, our solution is within a polylogarithmic factor from optimal. In other words, the linear dependency on the number of changes is inevitable. Throughout, the bounding parameter of our algorithm is 1. This indicates that a node is able to distinguish between the cases of having a single appearance of a letter  $\sigma$  in some of its ports and not having any appearances of  $\sigma$  in any of its ports.

The basic idea behind our algorithm is that we first use techniques from [13] to come up with an MIS quickly and then we fix any errors that the topology changes might induce. In other words, our goal is to first come up with a *proportional* MIS, where the likelihood of a node to join the MIS is inversely proportional to the number of neighbors the node has that have not yet decided their output. We partition the state set of our algorithm into two components. One of the components contains the input state and is responsible for computing the proportional MIS. Once a node has reached an output state or detects that it has been affected by a topology change, it transitions to the second component, which is responsible for fixing the errors, and never enters an active state of the first component. One of the reasons behind dividing our algorithm into two seemingly similar components is that the first component has stronger runtime guarantees, but requires that the nodes start in a “nice” configuration, whereas the second component does not have this requirement while providing weaker runtime bounds.


#### 3.1 The Proportional Component

The component used to compute the proportional MIS, referred to as **Proportional**, follows closely the design from [13]. The goal of the rest of the section is to introduce a slightly modified version of their algorithm that allows the non-affected nodes to ignore affected nodes while constructing the initial MIS.

**Proportional** consists of states  $Q = \{S, D_1, D_2, U_0, U_1, U_2, W, L\}$ , where  $Q_a = Q - \{W, L\}$  are referred to as *active* states and  $q_{yes} = W$  and  $q_{no} = L$  as *passive* states. We set  $S$  as the initial state. The communication alphabet is identical to the set of states; the algorithm is designed so that node  $v$  transmits letter  $q$  in round  $r$  whenever it resides in state  $q$  in round  $r + 1$ , i.e., state  $q$  was returned by the application of the transition function in round  $r$ .

***Intuition spotlight:** The main idea of **Proportional** is that each node  $v$  iteratively competes*





**Figure 1:** The transition function of **Proportional**. Each edge describes a possible transition and is labeled by the condition that is associated with the corresponding state transition. The lower case letters correspond to the number (counted up to the bounding parameter 1) of appearances of the corresponding letter, e.g., the condition  $u_1 + u_2 \geq 1$  is satisfied for node  $v$  if at least one port of  $v$  contains either letter  $U_1$  or  $U_2$ . Each transition  $q \rightarrow q'$  is delayed by state  $q''$  if there exists a transition  $q'' \rightarrow q$  (rules associated with delays are omitted from the figure for clarity).


*against its neighbors and the winner of a competition enters the MIS. In every competition, node  $v$  tosses fair coins in a synchronized way with its active neighbors. If a coin toss shows tails, node  $v$  loses the current competition. Conversely, node  $v$  wins the competition if it is the last node that remains in the competition. In the case of a tie, a new competition is started.*

Let us denote a state transition between states  $q$  and  $q'$  by  $q \rightarrow q'$  omitting the rules associated with this transition. To ease our notation and to make our illustrations more readable, we say that each state transition  $q \rightarrow q'$  is *delayed* by a set  $\mathcal{D} = \mathcal{D}(q \rightarrow q') \subseteq Q$  of delaying states. For  $q \rightarrow q'$  the set of delaying states corresponds to the states in  $Q - \{q\}$  from which there is a state transition to  $q$ . Transition  $q \rightarrow q'$  being delayed by state  $q''$  indicates that node  $v$  does not execute transition  $q \rightarrow q'$  as long as there is at least one letter in its ports that corresponds to the state  $q''$ . We say that node  $v$  in state  $q$  is delayed if there is a neighbor  $w$  of  $v$  that resides in a state that delays at least one of the transitions from  $q$ , i.e., node  $v$  cannot execute some transition because of node  $w$ . The transition function of **Proportional** is depicted in Figure 1.

Every competition, for an active node  $v$  and its active neighbors, begins by the nodes entering state  $U_0$ . During each round every node in state  $U_j, j \in \{0, 1, 2\}$ , assuming that it is not delayed, tosses a fair coin and proceeds to  $U_{j+1 \bmod 3}$  if the coin shows heads and to  $D_2$  otherwise. Notice that a competition does not necessarily start in the same round for two neighboring nodes. The circular delaying logic of the  $U$ -states ensure that node  $v$  is always at most one coin toss ahead of its neighbors. If node  $v$  observes that it is the only node in its neighborhood in a  $U$ -state, it enters state  $W$ , which corresponds to joining the MIS. In the case where, due to unfortunate coin tosses,  $v$  moves to state  $D_2$  along with all of its neighbors, node  $v$  restarts the process by entering state  $D_1$ .

We call a maximal contiguous sequence of rounds  $v$  spends in state  $q \in Q_a$  a  $q$ -turn and a maximal contiguous sequence of turns starting from a  $D_1$ -turn and not including another  $D_1$ -turn a *tournament*. We index the tournaments and the turns within the tournament by positive integers. The correctness of **Proportional** in the case of no topology changes follows the same line of arguments as in [13].





**Figure 3:** The transition function of **Greedy**. The conditions for the transitions are denoted by the labels on the edges and a lower case letter corresponds to the number of appearances of the corresponding letter. Unlike in **Proportional**, a transition from  $q'$  to  $q$  does not imply that  $q'$  delays any transition from  $q$ .

### 3.2 The Greedy Component

Now we extend **Proportional** to fix the MIS in the case that a topology change leaves the network in an illegal state. This extension of **Proportional** is referred to as the *greedy component* and denoted by **Greedy**.

***Intuition spotlight:** The basic idea is that nodes in states  $L$  and  $W$  verify locally that the configuration corresponds to a legal state of the network. If a node detects that the local configuration does not correspond to an MIS anymore, it revokes its output and tries again to join the MIS according to a slightly modified competition logic. The crucial difference is that we design **Greedy** without the circular delay logic. The reason behind the design is twofold: First, we cannot afford long chains of nodes being delayed. Second, since a dynamic change is only allowed to affect its 1-hop neighborhood, we cannot maintain the local synchrony of adjacent nodes similarly to **Proportional**.*

More precisely, the state set of the algorithm is extended by  $Q_2 = \{U', D'\}$ , where states  $U'$  or  $D'$  are referred to as active. The state transition function of the algorithm is illustrated in Figure 3. To detect an invalid configuration, we add the following state transitions from the output states  $W$  and  $L$ : a transition from state  $L$  to state  $D'$  in case that a node  $v$  resides in state  $L$  and does not have any letters  $W$  in its ports and a transition from state  $W$  to state  $D'$  in case  $v$  resides in state  $W$  and reads a letter  $W$  in its ports. Finally, to prevent nodes in the active states of **Proportional** and **Greedy** from entering state  $W$  at the same time and thus, inducing an incorrect output configuration, we add a transition from  $W$  and  $U'$  to  $D'$  in case that node  $v$  reads the initial letter  $S$ .

The logic of the new states  $U'$  and  $D'$  is the following: node  $v$  in state  $D'$  goes into state  $U'$  if it does not read the letter  $U'$  in its ports. Then  $u$  and its neighbors that transitioned from  $D'$  to  $U'$  during the same round compete against each other: In every round,  $v$  tosses a fair coin and goes back to state  $D'$  if the coin shows tails. We emphasize that nodes in state  $U'$  are not delayed by nodes in state  $D'$ . Then, if  $v$  remains the only node in its neighborhood in state  $U'$ , it declares itself as a winner and moves to state  $W$ . Conversely, if a neighbor of  $v$  wins,  $v$  goes into state  $L$ .

**Observation 1.** *A non-affected node is never in an active state of **Greedy**.*

*Proof.* Consider a non-affected node  $u$ . Since  $u$  is non-affected, it will never read the initial letter

$S$  in its ports after the first round. The design of **Proportional** guarantees that either  $u$  enters state  $W$  and no neighbor of  $u$  enters state  $W$  or  $u$  enters state  $L$  and at least one neighbor  $v$  of  $u$  enters state  $W$ . Due to the design of our algorithm, a node will never enter state  $W$  if it has a neighbor in state  $W$  and therefore, if  $u$  once enters state  $W$ , it will never reside in any other state. By the same argument, if  $u$  is in state  $L$ , the neighbor  $v$  in state  $W$  will guarantee that  $u$  never leaves state  $L$ .  $\square$

## 4 Analysis

**An Auxiliary Execution of Proportional.** The authors of [13] analyze an algorithm very similar to the one induced by **Proportional**, showing that its runtime (on a static graph) is  $O(\log^2 n)$  in expectation and w.h.p. In their analysis, they prove that with a positive constant probability, the number of edges decreases exponentially from one tournament to the next.

***Intuition spotlight:** The dynamic nature of the graph  $\mathcal{G}$  prevents us from analyzing the execution of **Proportional** in the same manner. To overcome this obstacle, we introduce an auxiliary execution  $\mathcal{E}$  of **Proportional** that turns out to be much easier to analyze, yet serves as a good approximation for the real execution.*

Consider some execution  $\eta$  of **Proportional** on  $\mathcal{G}$  and recall that  $\eta$  is determined by the adversarial graph sequence  $\mathcal{G}$  and by the coin tosses of the nodes. We associate with  $\eta$  the *auxiliary execution*  $\mathcal{E} = \mathcal{E}(\eta)$  obtained from  $\eta$  by copying the same coin tosses (for each node) and modifying the adversarial policy by applying to it the following two rules for  $i = 1, 2, \dots$ : (i) If node  $u$  is *excluded* during tournament  $i$  under  $\eta$ , then, instead, we *delete* node  $u$  immediately following its last  $U$ -turn under  $\mathcal{E}(\eta)$ . (ii) If node  $u$  becomes passive during tournament  $i$  under  $\eta$  and it would have remained active throughout tournament  $i$  under  $\mathcal{E}(\eta)$ , then we also *delete* node  $u$  immediately following its last  $U$ -turn under  $\mathcal{E}(\eta)$ .

Let  $V_\eta(i)$  and  $V(i)$  be the sets of nodes for which tournament  $i$  exists under  $\eta$  and  $\mathcal{E}(\eta)$ , respectively. Let  $X^v(i)$  and  $Y^v(i)$  denote the number of  $U$ -turns of node  $v$  in tournament  $i$  under  $\eta$  and  $\mathcal{E}(\eta)$ , respectively; to ensure that  $X^v(i)$  and  $Y^v(i)$  are well defined, we set  $X^v(i) = 0$  if  $v \notin V_\eta(i)$  and  $Y^v(i) = 0$  if  $v \notin V(i)$ , i.e., if tournament  $i$  does not exist for  $v$  under  $\eta$  and  $\mathcal{E}(\eta)$ , respectively. Let  $\Gamma_\eta(v, i)$  and  $\Gamma(v, i)$  denote the (exclusive) neighborhood of  $v$  at the beginning of tournament  $i$  under  $\eta$  and  $\mathcal{E}(\eta)$ , respectively. Observe that by definition,  $\Gamma_\eta(v, i) \subseteq V_\eta(i)$  and  $\Gamma(v, i) \subseteq V(i)$ . We say that node  $v$  *wins* in tournament  $i$  under  $\eta$  if  $X^v(i) > X^w(i)$  for all  $w \in \Gamma_\eta(v, i)$  and if  $v$  is not excluded in tournament  $i$ ; likewise, we say that node  $v$  *wins* in tournament  $i$  under  $\mathcal{E}(\eta)$  if  $Y^v(i) > Y^w(i)$  for all  $w \in \Gamma(v, i)$  and if  $v$  is not deleted in tournament  $i$ . The following lemma plays a key role in showing that the number of tournaments in  $\mathcal{E}(\eta)$  is at least as large as that in  $\eta$ .

**Lemma 1.** *If  $v \in V(i)$  wins in tournament  $i$  under  $\mathcal{E}(\eta)$ , then  $v$  wins in tournament  $i$  under  $\eta$ .*

*Proof.* First, we show that  $X^w(i) \leq Y^w(i)$  for any tournament  $i$  that exists in both  $\eta$  and  $\mathcal{E}(\eta)$

for any node  $w$ . We observe that if  $w$  is not excluded in tournament  $i$ , the number of  $U$ -turns is uniquely determined by the coin tosses in both executions, i.e.,  $X^w(i) = Y^w(i)$ . Then, if  $w$  is excluded in tournament  $i$ , there are at least as many  $U$ -turns in  $\mathcal{E}(\eta)$  as in  $\eta$ . Thus,  $X^w(i) \leq Y^w(i)$ .

Assume then that some node  $v$  wins in tournament  $i$  according to  $\mathcal{E}(\eta)$ . Since  $v$  can only win in tournament  $i$  if it is not excluded in this tournament, we get that  $X^v(i) = Y^v(i)$ . Node  $v$  winning in tournament  $i$  implies  $X^v(i) = Y^v(i) > Y^w(i) \geq X^w(i)$  for any  $w \in \Gamma(v, i)$  and thus,  $v$  also wins in  $\eta$ .

To conclude the proof, we observe that if tournament  $i$  does not exist for node  $v$  in  $\mathcal{E}(\eta)$ , then either  $v$  was deleted before entering state  $D_1$  of tournament  $i$  or either  $v$  or some of its neighbors won in tournament  $j < i$ . In all aforementioned cases, tournament  $i$  also does not exist for  $v$  in  $\eta$ . Finally, by construction, if tournament  $i$  does not exist for  $v$  in  $\eta$ , it also does not exist for  $v$  in  $\mathcal{E}(\eta)$  and the claim follows.  $\square$

#### 4.1 Runtime of Proportional

Next, we analyze the number of tournaments in  $\mathcal{E}(\eta)$ . Once we have a bound on the number of tournaments executed in  $\mathcal{E}(\eta)$ , it is fairly easy to obtain the same bound for  $\eta$ . Similarly as before, the set of nodes for which tournament  $i$  exists according to  $\mathcal{E}(\eta)$  is denoted by  $V(i)$ . Furthermore, let  $\Gamma(v, i)$  be the set of neighbors of  $v$  for which tournament  $i$  exists in  $\mathcal{E}(\eta)$  and let  $E(i) = \bigcup_{v \in V(i)} \{\{v, w\} \mid w \in \Gamma(v, i)\}$  and  $G(i) = (V(i), E(i))$ . Notice that  $G(i)$  is defined for the sake of the analysis and the topology of the underlying graph does not necessarily correspond to  $G_i$  in any round  $i$ , where  $G_i$  is the  $i$ th element in the adversarial graph sequence.

According to the design of **Proportional**, no node that once enters a passive state can enter an active state of **Proportional**. Furthermore, if any edge is added adjacent to node  $v$  after  $v$  has transitioned out of state  $S$ , node  $v$  will not participate in any tournament after the addition. Therefore, we get that  $V(i+1) \subseteq V(i)$  and  $E(i+1) \subseteq E(i)$  for any  $i$ . The following lemma plays a crucial role in the runtime analysis of **Proportional**.

**Intuition spotlight:** *In the case of no topology changes, the proof is (almost) the same as the one in [13]. By a delicate adjustment in the details of the proof, we show that the same argument holds in the presence of topology changes and only affects the constant  $\ell$  by a constant factor.*

**Lemma 2.** *There are two constant  $0 < p, \ell < 1$  such that  $|E(i+1)| \leq \ell|E(i)|$  with probability  $p$ .*

*Proof.* Let  $d(v, i) = |\Gamma(v, i)|$ . We say that a node  $v$  is *good* in  $G(i)$  if

$$|\{w \in \Gamma(v, i) \mid d(v, i) \geq d(w, i)\}| \geq d(v, i)/3 .$$

It is known, that at least half of the edges of any graph are adjacent to good nodes [3]. For the following part of the proof, we use techniques similar to the ones in [13]. The main idea is that a

good node is covered by a node with a smaller degree with a constant probability in any tournament and thus, a constant fraction of the edges is removed in each tournament. The subtle difference is that in the case of a dynamic network, there is no guarantee that a node adjacent to a good node remains adjacent until the end of the tournament. We tackle this by considering the auxiliary execution  $\mathcal{E}(\eta)$  and showing that the adversary essentially has to either delete a constant fraction of the edges in the graph or otherwise, we can use the previous techniques.

Let us now consider some good node  $v \in G(i)$ . Our goal is to show that at least  $1/6$  of the edges connected to  $v$  in tournament  $i$  do not exist in  $V(i+1)$  with a constant probability. Recall that we are now considering the auxiliary execution  $\mathcal{E}(\eta)$  in which nodes are deleted in the end of a tournament if they are excluded in this tournament according to  $\eta$ . We split our analysis into two cases, where the first case considers the option that in tournament  $i$ , the adversary either deletes  $v$  or at least half of the edges between  $v$  and the nodes in  $\{w \in \Gamma(v, i) \mid d(v, i) \geq d(w, i)\}$ . Notice that the adversary can do this by either directly deleting the edges or by deleting nodes in  $\{w \in \Gamma(v, i) \mid d(v, i) \geq d(w, i)\}$ . It follows that at least  $1/6$  edges connected to  $v$  do not exist in  $V(i+1)$  with probability 1.

Let us then assume for the second case that the adversary deletes at most half of the edges connected to

$$\{w \in \Gamma(v, i) \mid d(v, i) \geq d(w, i)\}$$

in tournament  $i$ . Let  $A(w, i)$  be the event that  $w$  is not deleted in tournament  $i$  and let

$$\hat{\Gamma}(v, i) = \{w \in \Gamma(v, i) \mid A(w, i) \wedge (d(w, i) \geq d(v, i))\} .$$

We say that node  $w \in \hat{\Gamma}(v, i)$  wins  $v$  in tournament  $i$  if

$$X^w(i) > \max\{X^z(i) \mid z \in \Gamma(w, i) \cup \hat{\Gamma}(v, i) - \{w\}\}$$

and denote this event by  $\text{WIN}_i(w, v)$ . The idea is that if  $w$  wins  $v$  in tournament  $i$ , then  $w$  enters the W state and  $v$  enters the L state. Furthermore, the events  $\text{WIN}_i(w, v)$  and  $\text{WIN}_i(z, v)$  are disjoint for every  $w, z \in \hat{\Gamma}(v, i), w \neq z$ .

Let us now fix a node  $w \in \hat{\Gamma}(v, i)$ . We denote the event that the maximum of  $\{X^z \mid z \in \Gamma(w, i) \cup \hat{\Gamma}(v, i)\}$  is attained at a single  $z \in \Gamma(w, i) \cup \hat{\Gamma}(v, i)$  by  $B_i(w, v)$ . Since

$$|\Gamma(w, i) \cup \hat{\Gamma}(v, i)| \leq 2d(v, i)$$

by the definition of a good node and  $X^z(i)$  are independent geometric random variables for all  $z \in \Gamma(w, i) \cup \hat{\Gamma}(v, i)$ , we get that

$$\mathbb{P}(\text{WIN}_i(w, v)) = \mathbb{P}(\text{WIN}_i(w, v) \mid B_i(w, v)) \cdot \mathbb{P}(B_i(w, v)) \geq \frac{1}{2d(v, i)} \cdot \frac{2}{3} .$$

Since  $v$  is good in  $G(i)$  and the events  $\text{WIN}_i(w, v)$  disjoint, we get that

$$\begin{aligned} \mathbb{P}(v \notin V(i+1)) &\geq \mathbb{P}\left(\bigvee_{w \in \hat{\Gamma}(v, i)} \text{WIN}_i(w, v)\right) \\ &= \sum_{w \in \hat{\Gamma}(v, i)} \mathbb{P}(\text{WIN}_i(w, v)) \\ &\geq \frac{d(v, i)}{6} \cdot \frac{1}{2d(v, i)} \cdot \frac{2}{3} = \frac{1}{18}. \end{aligned}$$

Recalling that half of the edges are connected to good nodes, we get that  $\mathbb{E}[|E(i+1)|] < 35/36|E(i)|$ . The claim now follows by Markov's inequality.  $\square$

**Theorem 1.** *Any node participates in  $\mathcal{O}(\log n)$  tournaments before becoming passive w.h.p. and in expectation.*

*Proof.* Let  $Z = \min\{0 \leq i \in \mathbb{Z} \mid |E(i)| = 0\}$ . By Lemma 2,  $Z$  is dominated by a random variable that follows distribution

$$\mathcal{O}(\log n) + \text{NB}(\mathcal{O}(\log n), 1 - p).$$

Therefore,  $Z = \mathcal{O}(\log n)$  in expectation and w.h.p.

According to the design of the algorithm, a node that is not excluded and does not have any active neighbors in tournament  $i$  goes into state  $W$  with probability 1. Therefore, all active nodes in tournament with index  $Z$  will either become passive or be deleted. Then, by applying Lemma 1 and observing that no new edges can be added between nodes in active states of **Proportional**, we get that there are no nodes for which tournament  $Z + 1$  exists in  $\eta$ .  $\square$

The last step of the analysis of **Proportional** is to bound the number of rounds that any node  $v$  spends in an active state. Let  $V_{\text{MAX}}(v)$  denote the maximal connected component of nodes in active states of **Proportional** that contains node  $v$ . Consider the following modification of **Proportional**: before starting tournament  $i + 1$ , node  $v$  waits for every other node in  $V_{\text{MAX}}(v)$  to finish tournament  $i$  or to become excluded. Notice that any node that gets inserted into the graph after node  $v$  has exited state  $S$  will not be part of  $V_{\text{MAX}}(v)$ . We do not claim that we know how to implement such a modification, but clearly the modified process is not faster than the original one.

The length of tournament  $i$  is determined by  $\max_v\{X^v(i)\}$ . Given that the random variables  $X^v(i)$  are independent and follow the  $\text{Geom}(1/2)$  distribution, we get that  $\max_v\{X^v(i)\} \in \mathcal{O}(\log n)$  with high probability and in expectation. Combining with Observation 1, we get the following corollaries.

**Corollary 3.** *Let  $t_v \geq 1$  be the round in which node  $v$  is inserted into the graph, where  $t_v = 1$  if  $v \in G_1$ . If  $v$  is not deleted, it enters either state  $W$  or  $L$  by time  $t_v + \mathcal{O}(\log^2 n)$  w.h.p. and in expectation.*

**Corollary 4.** *The maximum runtime of any non-affected node is  $\mathcal{O}(\log^2 n)$  in expectation and w.h.p.*

## 4.2 The Quality of an MIS

Before we go to the analysis of **Greedy**, we want to point out that even though an invalid configuration can be detected locally, a single topology change can have an influence in an arbitrary subgraph.

**Intuition spotlight:** *To demonstrate this observation, consider a graph  $G$  with some node  $v$  that is connected to every other node in  $G$  and assume, that after executing **Proportional**, the MIS consists of only node  $v$ . Given that the adversary deletes node  $v$ , the MIS has to be fixed on the remainder of the graph, that can have an arbitrary topology, e.g., a large diameter. Furthermore, several changes like this can potentially happen in a chain, which implies that not all nodes start executing the fixing component at the same time.*

We begin the analysis of **Greedy** by taking a closer look at the properties of non-affected nodes in the passive states and show that if a node has a high degree, it is either unlikely for this node to be in the MIS or that the neighbors of this node have more than one MIS node in their neighborhoods. Let  $i$  be the index of the tournament in which node  $v$  enters state  $W$ . We say that node  $v$  covers node  $w \in \Gamma(v, i) \cup \{v\}$  if  $w$  entered state  $L$  in tournament  $i$ .

**Definition 2.** *The quality  $q(v)$  of node  $v$  is given by  $q(v) = |\{w \in \Gamma(v, i) \cup \{v\} \mid v \text{ covers } w\}|$ . The quality  $q(B)$  of any set of nodes  $B$  is defined as  $\sum_{v \in B} q(v)$ .*

**Lemma 5.** *For any node  $v$ ,  $\mathbb{E}[q(v)] \in \mathcal{O}(\log n)$ .*

*Proof.* Consider node  $v$  and tournament  $i$ . Let  $\hat{\Gamma}(v, i) \subseteq \Gamma(v, i)$  be the neighbors of  $v$  that are not excluded in tournament  $i$ . We note that  $v$  can only cover  $w \in \Gamma(v, i)$  if  $v$  is not excluded in tournament  $i$ . Let  $q(v, i)$  denote the random variable that counts the number of nodes that  $v$  covers in tournament  $i$ . Since  $X^v(i)$  and  $X^w(i)$  for all  $w \in \Gamma(v, i)$  are independent  $\text{Geom}(1/2)$  random variables, we get that

$$\begin{aligned} \mathbb{E}[q(v, i)] &\leq \mathbb{P}\left[X^v(i) \geq \max\{X^w(i) \mid w \in \hat{\Gamma}(v, i)\}\right] \cdot (|\hat{\Gamma}_i(v)| + 1) \\ &\leq 2 \cdot \mathbb{P}\left[X^v(i) > \max\{X^w(i) \mid w \in \hat{\Gamma}(v, i)\}\right] \cdot (|\hat{\Gamma}_i(v)| + 1) \leq \frac{2(|\hat{\Gamma}(v, i)| + 1)}{|\hat{\Gamma}(v, i)| + 1} = 2. \end{aligned}$$

Consider the random variable  $Z = \min\{j \in \mathbb{Z}_{\geq 0} \mid |E(j)| = 0\}$ . The total number of tournaments is at most  $Z$  and  $\mathbb{E}[q(v, i)] = \mathbb{E}[q(v, i) \mid Z \geq i] \cdot \mathbb{P}[Z \geq i]$ . Thus, by Lemma 2

$$\mathbb{E}[q(v)] = \sum_{i=0}^{\infty} \mathbb{E}[q(v, i) \mid Z \geq i] \cdot \mathbb{P}[Z \geq i] \in \mathcal{O}(1) \cdot \left(\sum_{i=0}^{\infty} \mathbb{P}[Z \geq i]\right) \subseteq \mathcal{O}(\log n).$$

□



The quality of a node  $v$  gives an upper bound on the number of nodes in the neighborhood of  $v$  that are only covered by  $v$ . Let  $v$  be a node in state  $L$ . If  $t$  is the first round such that there are no nodes adjacent to  $v$  in state  $W$ , we say that  $v$  is *released* at time  $t$ . Similarly, we say that node  $v$  in state  $W$  is released if an adjacent edge is added. Every node can only be released once, i.e., node  $v$  counts as released even if it eventually again has a neighboring node in state  $W$  or if it enters state  $W$ .

**Lemma 6.** *Let  $H$  be the set of nodes that are eventually released. Then  $\mathbb{E}[|H|] \in \mathcal{O}(C \log n)$ .*

*Proof.* Let  $V^c$  be the set of nodes that are (eventually) deleted or incident to an edge which is (eventually) added or deleted. By Lemma 5, the quality of any node is  $\mathcal{O}(\log n)$ . Therefore,  $\mathbb{E}[q(V^c)] = \sum_{v \in V^c} \mathbb{E}[q(v)] \in |V^c| \cdot \mathcal{O}(\log n)$ . Thus,  $\mathbb{E}[|H|] \in |V^c| \cdot \mathcal{O}(\log n) \subseteq \mathcal{O}(C \log n)$ .  $\square$

### 4.3 Fixing the MIS

We call a maximal contiguous sequence of rounds in which node  $v$  resides in state  $U'$  a *greedy tournament*. Unlike with **Proportional**, we index these greedy tournaments by the time this particular tournament starts. In other words, if node  $v$  resides in state  $D'$  in round  $t - 1$  and in state  $U'$  in round  $t$ , we index this tournament by  $t$ . We denote the random variable that counts the number of transitions from  $U'$  to itself in greedy tournament  $t$  by node  $v$  by  $X^v(t)$ . Notice that  $X^v(t)$  obeys  $\text{Geom}(1/2)$  unless  $v$  is deleted or an adjacent edge is added during greedy tournament  $t$ . We say that a greedy tournament  $t$  is *active* if there is at least one node  $v$  in state  $U'$  of greedy tournament  $t$ .

**Observation 2.** *There are at most  $\mathcal{O}(\log n)$  rounds in which greedy tournament  $t$  is active in expectation and w.h.p.*

*Proof.* Let  $V(t)$  be the set of nodes that participate in greedy tournament  $t$ . Recall that the number of state transitions in greedy tournament is exactly  $X^v(t)$ . Furthermore,  $X^v(t)$  follows distribution  $\text{Geom}(1/2)$ . Since  $|V(t)| \leq n$ , all  $X^w(t), w \in V(t)$  are independent and the maximum of  $X^w(t)$  for at most  $n$  variables is  $\mathcal{O}(\log n)$  w.h.p. and in expectation, the claim follows.  $\square$

**Intuition spotlight:** *In every greedy tournament  $i$ , either a node joins the MIS with a constant probability or a topology change occurs that affects some node. Furthermore, any single topology change can force at most 2 nodes to leave state  $W$ . Therefore, the sum of transitions to state  $W$  by the nodes that ever execute **Greedy**, i.e., nodes that are either released or excluded, can be at most  $\mathcal{O}(H + C)$ . Thus, there can be at most  $\mathcal{O}(H + C)$  greedy tournaments in total.*

**Lemma 7.** *The total expected number of greedy tournaments is  $\mathcal{O}(C \log n)$ .*

*Proof.* To prove the claim, we bound from above the sum of state transitions into state  $W$  by the released nodes and the excluded nodes. A node transitions away from state  $W$  at most once per

added incident edge and therefore, there can be at most  $2C$  transitions from  $W$  to  $D'$  in total. After at most  $2C + |H|$  transitions into state  $W$  by the released nodes, all released nodes are either in state  $W$  or deleted.

We say that a greedy tournament  $t$  is *unclean* if either, some node  $v$  in an active state of greedy tournament  $t$  is deleted or an edge is added adjacent to  $v$ . Otherwise, we say that greedy tournament  $t$  is *clean*. Let  $V(t)$  be the nodes that participate in greedy tournament  $t$ , i.e., the nodes that reside in state  $U'$  in round  $t$  and in state  $D'$  in round  $t - 1$ . Assuming that greedy tournament  $t$  is clean and recalling that the random variables  $X^v(t)$  are independent, the probability that the maximum of  $X^v(t), v \in V(t)$  is attained in a single node is at least  $1/3$ . Let  $Y'_k$  be the random variable that counts the number of clean greedy tournaments we need until the sum of state transitions to  $W$  by the released nodes is  $k$ . It is easy to see that  $Y'_k$  is dominated by random variable  $Y_k$  that obeys distribution  $k + \text{NB}(k, 2/3)$ .

Now we set  $k = 2C + |H|$  and get that  $\mathbb{E}[Y_k] \in \mathcal{O}(k) \subseteq \mathcal{O}(C + |H|)$ . By Lemma 6,  $|H| \in \mathcal{O}(C \log n)$  in expectation. Thus, we get that  $\mathbb{E}[Y'_k] \in \mathcal{O}(C + |H|) \subseteq \mathcal{O}(C \log n)$ . Finally, we observe that there can be at most  $2C$  unclean greedy tournaments and thus, the total expected number of greedy tournaments is  $\mathcal{O}(C \log n) + \mathcal{O}(C) = \mathcal{O}(C \log n)$ .  $\square$

**Observation 3.** *Let  $k$  be the total number of greedy tournaments. There are at most  $2k + 2C$  non-silent rounds without either *i*) at least one node in an active state of **Proportional** or *ii*) an active greedy tournament.*

*Proof.* Consider a non-silent round  $t$  and assume that in rounds  $t$  and  $t + 1$  no changes occur, that no greedy tournaments are active, and that no node is in an active state of **Proportional**. Notice that no node  $v$  can be in state  $D'$  in round  $t$ . Otherwise,  $v$  transitions to state  $U'$  in round  $t + 1$  and thus, greedy tournament  $t + 1$  is active in round  $t + 1$ . Then, in graph  $G_t$ , either *i*) there exists at least one pair of adjacent nodes  $v, w$  that are in state  $W$  or *ii*) there is a node  $v'$  that is in state  $L$  and none of its neighbors are in state  $W$ .

In case *ii*), node  $v'$  can have only letters  $S$  and  $L$  in its ports in round  $t$ . Therefore, the logic of the  $L$  state ensures that node  $v'$  is in state  $D'$  in round  $t + 1$ . Now, we get that greedy tournament  $t + 2$  is active in round  $t + 2$ .

In case *i*), both nodes  $v$  and  $w$  have at least one letter  $S$  or  $W$  in their ports in round  $t$ . Therefore, the logic of the  $W$  state ensures that, both  $v$  and  $w$  reside in state  $D'$  in round  $t + 1$ . Again we get that a greedy tournament has to be active in round  $t + 2$ . Thus, given no topology changes, there can be at most 2 successive non-silent rounds without an active greedy tournament or at least one node active in **Proportional** and the claim follows.  $\square$

Next we show that our MIS algorithm, that consists of **Proportional** and **Greedy**, fulfills the correctness properties given in the model section. Then, we establish our main result, i.e., that our

MIS algorithm is indeed effectively confining.

**Lemma 8.** *Our MIS algorithm is correct.*


*Proof.* The correctness property (C3) follows directly from the logic of the states  $W$  and  $L$ . To obtain property (C1), we first observe that if the nodes are in an incorrect output configuration in round  $r$ , there has to be either some node in state  $W$  with another node in state  $W$  in its neighborhood or some node along with all of its neighbors in state  $L$ . In both cases, the logic of states  $W$  and  $L$  ensure that the configuration, given no topology changes, is a non-output configuration in round  $r + 1$ . Given Corollary 3, we get that eventually, no active nodes in **Proportional** remain once there are no more topology changes. Combining this with Observation 3 and Lemma 7, we get property (C1).

Consider some node  $v$  in an active state in round  $r$ . According to the design of our protocol, if an edge is added adjacent to  $v$  in round  $r$ , then  $v$  will not reside in state  $W$  in round  $r + 1$ . We observe that due to the logic of the excluding transitions, i.e., the transitions to state  $D'$  when reading the letter  $S$ , node  $v$  that participates in a tournament of **Proportional** cannot have any neighbors that participate in a greedy tournament. Furthermore, node  $v$  can only enter state  $W$  by winning a tournament (either in **Proportional** or in **Greedy**). Therefore, it cannot be the case that both  $v$  and some neighbor  $w$  of  $v$  enter state  $W$  in round  $r + 1$ . Furthermore, node  $v$  can only enter state  $L$  in round  $r + 1$  if it has a neighbor in state  $W$  in round  $r$ . Thus, if node  $v$  is in a non-output state in round  $r$ , it will either be in a non-output state, enter state  $W$  and none of its neighbors will, or enter state  $L$  and at least one of its neighbors will enter state  $W$  in round  $r + 1$ .

Consider then the case that node  $v$  is in an output state in round  $r$ . Due to the logic of state  $W$ , node  $v$  only remains in state  $W$  if there are no letters  $S$  or  $W$  in its ports. Since a port that corresponds to a neighbor in state  $W$  can only contain letters  $W$  or  $S$ ,  $v$  cannot have a neighbor in state  $W$  in round  $r$ . Due to the logic of state  $L$ ,  $v$  only remains in  $L$  if there is at least one neighbor in state  $W$ . Therefore, either the configuration in round  $r + 1$  is not an output configuration or the configuration is a correct output configuration and we have obtained property (C2). □

**Theorem 3.** *The global runtime of our MIS algorithm is  $\mathcal{O}((C + 1) \log^2 n)$  in expectation.*

*Proof.* By Lemma 7, we have  $\mathcal{O}(C \log n)$  active greedy tournaments in expectation and by Observation 2 each greedy tournament takes  $\mathcal{O}(\log n)$  rounds w.h.p. and in expectation. Thus, in expectation, we have  $\mathcal{O}(C \log^2 n)$  non-silent rounds where a greedy tournament is active. Combining this with Corollary 3 and Observation 3 and by linearity of expectation, the total expected runtime is  $\mathcal{O}(\log^2 n) + \mathcal{O}(C \log n) + \mathcal{O}(C \log^2 n) \subseteq \mathcal{O}((C + 1) \log^2 n)$ . □



**Figure 4:** One node in each of the connected components has to be in the MIS. Since the nodes in the components are indistinguishable and execute the same algorithm, the probability of  $u_i$  being in the MIS for any  $i$  is  $1/3$ .

## 5 Lower Bound

The runtime of our algorithm might seem rather slow at the first glance, since it is linear in the number of topology changes. In this section, we show that one cannot get rid of the linear dependency, i.e., there are graphs where the runtime of any algorithm grows at least linearly with the number of changes. In particular, we construct a graph where the runtime of any algorithm is  $\Omega(C)$ .

***Intuition spotlight:** Even in a graph with two connected nodes, any algorithm has to perform some sort of symmetry breaking. Furthermore, if the degrees of node  $u$  and its neighbors are small, it is likely that deleting  $u$  leaves some of its neighbors in an invalid state.*

Let  $G^\ell$  be a graph that consists of  $n$  nodes and  $n = 3\ell$ . The graph consists of  $\ell$  components  $B_i$ , where  $B_i$  is a 3-clique for each  $1 \leq i \leq \ell$ . In addition, let  $u_1, \dots, u_\ell$  be a set of nodes such that  $u_i \in B_i$  for every  $i$ . For illustration, see Figure 4.

**Theorem 4.** *There exists a graph  $G$  and an schedule of updates such that the runtime of any algorithm on  $G$  is  $\Omega(C)$  in expectation and w.h.p.*

*Proof.* Consider graph  $G^\ell$  and the following adversarial strategy. In round  $2i - 1 \geq 1$ , the adversary deletes node  $u_i$ , i.e., one of the nodes in component  $B_i$ . Our goal is to show that at least a constant fraction of the first  $2C$  rounds are non-silent for any  $C \leq n/3$ . Consider round  $2i$  and component  $B_i$ . Since the nodes in component  $B_i$  form a clique, their views are identical. Therefore, according to any algorithm that computes an MIS, their probability to join the MIS is equal, i.e.,  $1/3$ . Thus, the probability that the nodes in  $B_i - \{u_i\}$  are not in the MIS in any round  $2j \leq 2i$  is at least  $1/3$ .

Let  $X_i$  then be the indicator random variable, where  $X_i = 0$  if round  $i$  is silent and 1 otherwise and let  $X = \sum_i^\infty X_i$  be the random variable that counts the number of non-silent rounds. Since any round  $1 \leq 2i \leq 2C$  is non-silent with at least probability  $1/3$ , we get that  $\mathbb{E}[X] \geq (1/3)C$ . Now by applying a Chernoff bound, we get that

$$\mathbb{P}[X < 1/2\mathbb{E}[X]] = \mathbb{P}[X < (1/2) \cdot (1/3)C] \leq 2^{-C/12}.$$

Since  $C = n/3$ , we get that  $\mathbb{P}[X < 1/2\mathbb{E}[X]] \in \mathcal{O}(n^{-k})$  for any constant  $k$  and thus, the claim follows.  $\square$

## 6 Concluding remarks

As stated in Sec. 1, this paper makes the first step in the study of recovery from dynamic changes under the Stone Age model and we believe that exploring this research direction further raises some fascinating open questions: How does the complexity of the MIS problem in dynamic graphs changes under various aspects of asynchrony (e.g., asynchronous wake-up or even fully asynchronous message passing)? Does there exist an efficient Stone Age MIS algorithm if the communication model is modified so that the node receives its own transmissions (depicted by allowing the graph to have self-loops)? What would be the effect of augmenting the problem definition with self-stabilization requirements?

## References

- [1] Yehuda Afek, Noga Alon, Ziv Bar-Joseph, Alejandro Cornejo, Bernhard Haeupler, and Fabian Kuhn. Beeping a Maximal Independent Set. In *Proceedings of the 25th International Conference on Distributed Computing (DISC)*, pages 32–50, 2011.
- [2] Yehuda Afek, Noga Alon, Omer Barad, Eran Hornstein, Naama Barkai, and Ziv Bar-Joseph. A Biological Solution to a Fundamental Distributed Computing Problem. *Science*, 331(6014):183–185, 2011.
- [3] Noga Alon, László Babai, and Alon Itai. A Fast and Simple Randomized Parallel Algorithm for the Maximal Independent Set Problem. *Journal of Algorithms*, 7:567–583, December 1986.
- [4] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in Networks of Passively Mobile Finite-state Sensors. *Distributed Computing*, pages 235–253, 2006.
- [5] James Aspnes and Eric Ruppert. An introduction to population protocols. In Benoît Garbinato, Hugo Miranda, and Luís Rodrigues, editors, *Middleware for Network Eccentric and Mobile Applications*, pages 97–120. Springer-Verlag, 2009.
- [6] Baruch Awerbuch, Boaz Patt-Shamir, David Peleg, and Michael Saks. Adapting to Asynchronous Dynamic Networks (Extended Abstract). In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing (STOC)*, pages 557–570, 1992.
- [7] Baruch Awerbuch and Michael Sipser. Dynamic Networks Are as Fast as Static Networks. In *29th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 206–220, 1988.
- [8] Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. The Locality of Distributed Symmetry Breaking. In *Proceedings of the 53rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 321–330, 2012.

- [9] I. Chlamtac and S. Kutten. On Broadcasting in Radio Networks—Problem Analysis and Protocol Design. *Communications, IEEE Transactions on [legacy, pre - 1988]*, 33(12):1240–1246, 1985.
- [10] Richard Cole and Uzi Vishkin. Deterministic Coin Tossing with Applications to Optimal Parallel List Ranking. *Information and Control*, 70(1):32–53, jul 1986.
- [11] Alejandro Cornejo and Fabian Kuhn. Deploying Wireless Networks with Beeps. In *Proceedings of the 24th International Conference on Distributed Computing (DISC)*, pages 148–162, 2010.
- [12] Sebastian Daum, Mohsen Ghaffari, Seth Gilbert, Fabian Kuhn, and Calvin Newport. Maximal Independent Sets in Multichannel Radio Networks. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 335–344, 2013.
- [13] Yuval Emek and Roger Wattenhofer. Stone Age Distributed Computing. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 137–146, 2013.
- [14] Jasmin Fisher, Thomas A. Henzinger, Maria Mateescu, and Nir Piterman. Bounded Asynchrony: Concurrency for Modeling Cell-Cell Interactions. In Jasmin Fisher, editor, *FMSB*, volume 5054 of *Lecture Notes in Computer Science*, pages 17–32. Springer, 2008.
- [15] Roland Flury and Roger Wattenhofer. Slotted Programming for Sensor Networks. In *IPSN*, April 2010.
- [16] Martin Gardner. The Fantastic Combinations of John Conway’s New Solitaire Game ‘Life’. *Scientific American*, 223(4):120–123, 1970.
- [17] Mohsen Ghaffari. An Improved Distributed Algorithm for Maximal Independent Set. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 270–277, 2016.
- [18] Sukumar Ghosh, Arobinda Gupta, Ted Herman, and Srimam Pemmaraju. Fault-containing Self-stabilizing Algorithms. In *Proceedings of the 15th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 45–54, 1996.
- [19] Thomas Hayes, Jared Saia, and Amitabh Trehan. The Forgiving Graph: A Distributed Data Structure for Low Stretch Under Adversarial Attack. In *Proceedings of the 28th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 121–130, 2009.
- [20] Michael König and Roger Wattenhofer. On Local Fixing. In *17th International Conference on Principles of Distributed Systems (OPODIS)*, pages 191–205, 2013.
- [21] Amos Korman. Improved Compact Routing Schemes for Dynamic Trees. In *Proceedings of the 27th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 185–194, 2008.

- [22] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. What Cannot be Computed Locally! In *Proceedings of the 23th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 300–309, 2004.
- [23] Fabian Kuhn, Stefan Schmid, and Roger Wattenhofer. A Self-repairing Peer-to-peer System Resilient to Dynamic Adversarial Churn. In *Proceedings of the 4th International Conference on Peer-to-Peer Systems (IPTPS)*, pages 13–23, 2005.
- [24] Shay Kutten and David Peleg. Fault-Local Distributed Mending. *Journal of Algorithms*, 30(1):144–165, 1999.
- [25] Shay Kutten and David Peleg. Tight Fault Locality. *SIAM Journal on Computing*, 30(1):247–268, 2000.
- [26] Xiaozhou Li, Jayadev Misra, and C.Greg Plaxton. Active and Concurrent Topology Maintenance. In Rachid Guerraoui, editor, *Distributed Computing*, volume 3274 of *Lecture Notes in Computer Science*, pages 320–334. Springer Berlin Heidelberg, 2004.
- [27] Ji-Cherng Lin and Tetz Huang. An Efficient Fault-Containing Self-Stabilizing Algorithm for Finding a Maximal Independent Set. *IEEE Transactions on Parallel and Distributed Systems*, 14(8):742–754, 2003.
- [28] Nati Linial. Locality in Distributed Graph Algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.
- [29] Michael Luby. A Simple Parallel Algorithm for the Maximal Independent Set Problem. *SIAM J. Comput.*, 15:1036–1055, November 1986.
- [30] Navneet Malpani, Jennifer L. Welch, and Nitin Waidya. Leader Election Algorithms for Mobile Ad Hoc Networks. In *Proceedings of the 4th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIAL-M)*, pages 96–103, 2003.
- [31] Carsten Marr and Marc-Thorsten Hütt. Outer-totalistic Cellular Automata on Graphs. *Physics Letters A*, 373(5):546–549, 2009.
- [32] Othon Michail, Ioannis Chatzigiannakis, and Paul G. Spirakis. *New Models for Population Protocols*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2011.
- [33] Alessandro Panconesi and Aravind Srinivasan. On the Complexity of Distributed Network Decomposition. *Journal of Algorithms*, pages 356–374, 1996.
- [34] David Peleg. *Distributed Computing: a Locality-sensitive Approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.

- [35] Johannes Schneider and Roger Wattenhofer. A Log-star Distributed Maximal Independent Set Algorithm for Growth-bounded Graphs. In *Proceedings of the 27th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 35–44, 2008.
- [36] Alex Scott, Peter Jeavons, and Lei Xu. Feedback from Nature: an Optimal Distributed Algorithm for Maximal Independent Set Selection. In *Proceedings of the 32nd Symposium on Principles of Distributed Computing (PODC)*, pages 147–156, 2013.
- [37] Sandeep Shukla, Daniel Rosenkrantz, and S. S. Ravi. Observations on Self-Stabilizing Graph Algorithms for Anonymous Networks (Extended Abstract). In *Proceedings of the Second Workshop on Self-Stabilizing Systems*, pages 1–15, 1995.
- [38] L. G. Valiant. Parallel Computation. In *7th IBM Symp. on Math. Foundations of Computer Science*, 1982.
- [39] John von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, Champaign, IL, USA, 1966.
- [40] Jennifer Walter, Jennifer L. Welch, and Nitin Vaidya. A Mutual Exclusion Algorithm for Ad Hoc Mobile Networks. *Wireless Networks*, 7:585–600, 1998.
- [41] Stephen Wolfram. *A new kind of science*. Wolfram Media, Champaign, Illinois, 2002.



## A Pseudo-Locality

In this section, we show that the runtime of node  $u$  is independent of topology updates further than  $1 + \log n$  distance from  $u$ . We say that an update occurs in distance  $d$  from node  $u$  if it involves a node within the  $d$ -hop neighborhood of  $u$ , where involving node  $u$  indicates that either  $u$  is inserted/deleted or an edge adjacent to  $u$  is inserted/deleted. Let  $N_r^i(u)$  denote the (inclusive)  $i$ -hop neighborhood of node  $u$  in round  $r$ . Then, we denote the number of topology changes that involve the nodes in

$$\bigcup_{r=0}^{\infty} N_r^{1+\log n}(u) = N_{\infty}^{1+\log n}(u)$$

by  $C_u$ . In other words,  $C_u$  counts the number of changes that involve nodes that have been within  $1 + \log n$  hops from  $u$  in some round.

We start by using Lemma 5 to bound the expected number of released nodes in  $N_{\infty}^{\log n}(u)$ . The quality of the nodes in  $N_{\infty}^{1+\log n}(u)$  gives an upper bound to the number of released nodes in  $N_{\infty}^{\log n}(u)$ . The proof for the following lemma is analogous to the proof of Lemma 6 and therefore omitted.

**Lemma 9.** *Let  $H_u$  be the released nodes in  $N_{\infty}^{\log n}(u)$ . Then  $\mathbb{E}[|H_u|] \in \mathcal{O}(C_u \log n)$ .*

Similarly to the proof of the global runtime, we wish to bound the number of greedy tournaments before either all released nodes in  $N_{\infty}^{\log n}(u)$  become permanently passive or are deleted. We say that a node *participates* in greedy tournament  $t$  if it enters state  $U'$  from  $D'$  at time  $t$ . We say that greedy tournament  $i$  is (locally) unclean, if there is a node  $v \in N_{\infty}^{\log n}(u)$  in an active state of greedy tournament  $i$  while  $v$  is deleted or an edge is added adjacent to  $v$ . Otherwise, greedy tournament  $i$  is clean. We show that if at least one node in  $N_{\infty}^1(u)$  participates in greedy tournament  $i$  and greedy tournament  $i$  is (locally) clean, then with a constant probability, at least one node  $v \in N_{\infty}^{\log n}(u)$  wins in greedy tournament  $i$ . Recall that winning indicates that  $v$  enters state  $W$ .

**Lemma 10.** *Consider a clean greedy tournament  $t$  where at least one node in  $N_{\infty}^1(u)$  participates. Then there is some  $v \in N_{\infty}^{\log n}(u)$  s.t.  $v$  wins in greedy tournament  $t$  with probability at least  $1/6$ .*

*Proof.* Let  $N^i(u, t) \subseteq N_{\infty}^i(u)$  be the set of nodes (within  $i$  hops from  $u$ ) that participate in greedy tournament  $t$ . To prove the claim, we first show that if  $|N^i(u, t)| \leq 2|N^{i-1}(u, t)|$  for any  $i$ , then there is a winning node in  $N^i(u, t) \subseteq N_{\infty}^i(u)$  with probability at least  $1/6$ . Assume  $|N^i(u, t)| \leq 2|N^{i-1}(u, t)|$  and let  $A(u, t)$  be the event that the maximum of  $\{X^u(t) \mid u \in N^i(u, t)\}$  is attained in a single node. Since greedy tournament  $t$  is clean, the probability of  $A(u, t)$  is at least  $1/3$ . As the coin tosses in  $N^i(u, t)$  obey the same distribution, the maximum is attained equally likely in all the nodes. Therefore, a node from  $N^{i-1}(u, t)$  wins with probability

$$\frac{1}{3} \frac{|N^{i-1}(u, t)|}{|N^i(u, t)|} \geq \frac{1}{6}.$$

Assume now for contradiction that  $|N^i(u, t)| > 2|N^{i-1}(u, t)|$  for all  $0 \leq i \leq \log n$ . It follows that  $|N^{\log n}(u, t)| > 2^{\log n} = n$ , which is a contradiction. Therefore, for any clean greedy tournament  $t$ , there is always a node in  $N^{\log n}(u, t) \subseteq N_{\infty}^{\log n}(u)$  that wins in greedy tournament  $t$  with probability at least  $1/6$ .  $\square$

Now with an argument analogous to the one of Lemma 7, we get that the expected number of clean greedy tournaments in which at least one node in  $N_{\infty}^1(u)$  participates is  $\mathcal{O}(C_u \log n)$ . By applying Observation 2, we get that the number of rounds in which any node in  $N_{\infty}^1(u)$  participates in any greedy tournament is  $\mathcal{O}(C_u \log^2 n)$ . Since there can be at most  $2C_u$  unclean greedy tournaments, the number of rounds in which some node in  $N_{\infty}^1(u)$  participates in a greedy tournament that is not clean is bounded by  $C_u \cdot \mathcal{O}(\log n)$ . According to the design of the  $D'$  state, there cannot be 2 successive rounds in which  $u$  is active in **Greedy** and no node in  $N_{\infty}^1(u)$  participates in any greedy tournament or are in an active state of **Proportional**. Finally, recalling that a non-affected node becomes passive in  $\mathcal{O}(\log^2 n)$  rounds and by Corollary 3, we get the following theorem.

**Theorem 5.** *The expected runtime is  $\mathcal{O}((C_u + 1) \log^2 n)$  for any node  $v$ .*