

Self-Stabilizing MIS Computation in the Beeping Model

George Giakkoupis^{*1}, Volker Turau^{†2}, and Isabella Ziccardi^{‡3}

¹INRIA, University of Rennes, CNRS, IRISA, France

²Institute of Telematics, Hamburg University of Technology, Germany

³Bocconi University, BIDSa, Italy

Abstract

We consider self-stabilizing algorithms to compute a Maximal Independent Set (MIS) in the extremely weak *beeping* communication model. The model consists of an anonymous network with synchronous rounds. In each round, each vertex can optionally transmit a signal to all its neighbors (beep). After the transmission of a signal, each vertex can only differentiate between no signal received, or at least one signal received. We assume that vertices have some knowledge about the topology of the network.

We revisit the not self-stabilizing algorithm proposed by Jeavons, Scott, and Xu (2013), which computes an MIS in the beeping model. We enhance this algorithm to be self-stabilizing, and explore two different variants, which differ in the knowledge about the topology available to the vertices. In the first variant, every vertex knows an upper bound on the maximum degree Δ of the graph. For this case, we prove that the proposed self-stabilizing version maintains the same run-time as the original algorithm, i.e. it stabilizes after $O(\log n)$ rounds w.h.p. on any n -vertex graph. In the second variant, each vertex only knows an upper bound on its own degree. For this case, we prove that the algorithm stabilizes after $O(\log n \cdot \log \log n)$ rounds on any n -vertex graph, w.h.p.

1 Introduction

The Maximal Independent Set (MIS) problem has a central role in the areas of parallel and distributed computing. In a graph $G = (V, E)$, an MIS is a subset of vertices $I \subseteq V$ where no two vertices in I are adjacent, and it is maximal with respect to inclusion. Recognized for its importance in the field of distributed computing since the early 1980s [21, 3], the computation of an MIS serves as a foundational subroutine in various algorithms in wireless networks, routing, and clustering [23]. The interest in the MIS problem has recently extended to biological networks, with observations of processes similar to the MIS elections in the development of the fly’s nervous system [2].

While distributed MIS algorithms are well explored in the standard synchronous message-passing models like LOCAL, CONGEST, and CONGESTED-CLIQUE [23, 19, 20, 13, 16, 12, 4, 15], recently the MIS selection was considered also within weaker communication frameworks [22, 1, 9]. Indeed, novel distributed communication models, inspired by scenarios in biological

^{*}george.giakkoupis@inria.fr

[†]turau@tuhh.de

[‡]isabella.ziccardi@unibocconi.it

cellular networks, wireless sensor networks and networks with sub-microprocessor devices, were defined. The Stone Age model, introduced by Emek and Wattenhofer, provides an abstraction of a network of randomized finite state machines that communicate with their neighbors using a fixed message alphabet based on a weak communication scheme [11]. Another related model, which is the one we consider in this paper, is the full-duplex beeping model¹, where a network of anonymous processors and synchronous rounds is considered [6]. In each round, each vertex has the option to either broadcast a signal – a beep – to all its neighbors or to remain silent. Subsequently, each vertex can determine whether it received any signals or if all its neighbors remained silent. This does not allow a vertex to differentiate which vertex emitted the signal, nor the number of signals received. We notice that a variation of this model can be defined where, instead of a single type of signal, a constant number of distinct signals exist, and the vertices can distinguish between the types of signals received. The beeping model finds motivation in scenarios such as wireless sensor networks or biological systems, where organisms can only detect proteins transmitted by neighboring entities [1]. The problem of computing an MIS was already considered in the full-duplex beeping model [18, 14, 1] and in the Stone Age model [10, 9, 11].

In both biological and wireless systems, another notable trait is their capability for self-recovery. This ability is also essential in distributed and large-scale systems, which must be able to effectively manage faults. Self-stabilizing algorithms are designed to ensure that systems can recover from any state and eventually stabilize into a valid state, maintaining stability as long as faults are absent [7, 8]. Indeed, self-stabilizing algorithms are guaranteed to converge from any initial configuration. However, only a few self-stabilizing MIS algorithms have been proposed for the aforementioned weak communication models [1, 9, 17]. In the full-duplex beeping model, Afek et al. [1] introduced a self-stabilizing algorithm that converges to an MIS in $O(\log^2 N \log n)$ rounds with high probability (w.h.p.), if all vertices know an upper bound N on the network's size n . In the same model, a constant-state algorithm was proposed in [17], which stabilizes in poly-logarithmic rounds w.h.p., albeit being efficient only for some graph families. Meanwhile, Emek et al. [9] devised an algorithm for a simplified version of the Stone Age model that is slightly stronger than the beeping communication model, which stabilizes in $O((D + \log n) \log n)$ rounds w.h.p. on any D -bounded diameter graph, where D is considered a fixed parameter. However, in this context, it would be desirable to relinquish the assumption that vertices possess global information about the network's structure.

Algorithms that do not require any knowledge of the network's topology were also proposed for the beeping model, but they strongly rely on the assumption that, at the beginning of the algorithm, the vertices are in the same fixed initial state, and hence they are not self-stabilizing. One algorithm was proposed by Afek et al. [1] for the full-duplex beeping model, which stabilizes in $O(\log^2 n)$ rounds w.h.p., without requiring vertex knowledge of the network's topology. Later, Jeavons et al. [18] improved this result by proposing an algorithm for the same model, capable of computing an MIS in any n -vertex graph in $O(\log n)$ rounds w.h.p., without requiring any vertex knowledge². Notice that these algorithms are not self-stabilizing because they also rely on the presence of phases of two rounds, implying a synchronization of the vertices modulo two.

¹This model is also called the beeping model with collision detection.

²Ghaffari provided a refined analysis for Jeavons' algorithm in [14].

1.1 Our Contribution

In this paper, we propose a self-stabilizing algorithm for computing the MIS in the full-duplex beeping model, aiming for a stabilization time of $O(\log n)$ with minimal vertex knowledge about network topology. The starting point of our work is Jeavons' algorithm in [18], which is non-self-stabilizing and converges within $O(\log n)$ rounds. We propose two variants that achieve self-stabilization and efficiency across all graph sizes. Our algorithms rely on each vertex's ability to compute a quantity $\ell_{\max}(v)$, which may require access to information like the maximum degree of the graph or its neighbors. The first variant assumes that vertices know an upper bound on the maximum degree Δ and stabilizes in $O(\log n)$ rounds, while the second variant assumes that each vertex knows an upper bound on its own degree and stabilizes in $O(\log n \log \log n)$ rounds. Additionally, we present a third algorithm for the extended beeping model with two channels, stabilizing in $O(\log n)$ time if vertices know an upper bound on the maximum degree among the 1-hop neighborhood. In summary, our contributions yield three algorithms for computing MIS in the beeping model, each highlighting different scenarios based on varying levels of vertex knowledge and beeping channels. The results just described are summarized in the following theorem.

Theorem. *Let G be a n -vertex graph. Then,*

1. *If each vertex v knows the same upper bound on the maximum degree of G , which is at most polynomial in n , then an MIS can be computed in the beeping model, in a self-stabilizing manner, within $O(\log n)$ rounds w.h.p.*
2. *If each vertex v knows an upper bound on its own degree, which is at most polynomial in n , then an MIS can be computed in the beeping model, in a self-stabilizing manner, within $O(\log n \log \log n)$ rounds w.h.p.*
3. *If each vertex v knows an upper bound on the maximum degree of all vertices in its 1-hop neighborhood, which is at most polynomial in n , then an MIS can be computed, in the beeping model with two channels, in a self-stabilizing manner, within $O(\log n)$ rounds w.h.p.*

It remains an open question whether a fast, self-stabilizing algorithm computing an MIS in the beeping model can be designed so that no information about the network topology is required to be known by the vertices.

2 The Algorithm

We assume the full-duplex beeping communication model and the starting point for our algorithm is the beeping, randomized algorithm of Jeavons et al. in [18]. Each vertex v is associated with an adaptive probability $p_t(v)$ of beeping in round t , and the algorithm works in phases, each consisting of two rounds. In the first round of each phase, each vertex v beeps with probability $p_t(v)$ and, if v beeps and all its neighbors are silent, then v joins the MIS. In the second round of each phase, vertices that joined the MIS beep and neighboring vertices hearing a beep become non-MIS vertices. Then, the newly joined MIS and non-MIS vertices remain silent for the rest of the algorithm. The crucial point leading to a $O(\log n)$ global round complexity with

high probability, is that active vertices adapt in each phase the beeping probability, initially $p_1(v) = 1/2$ for each vertex v . The value of $p_{t+1}(v)$ is decreased whenever neighboring vertices beep and is increased otherwise. In particular $p_{t+1}(v) = p_t(v)/2$ in the former case and $p_{t+1}(v) = \min\{2p_t(v), 1/2\}$ otherwise. The rationale of this behavior is twofold: to reduce the probability of neighboring vertices attempting to concurrently join the MIS, and to increase the probability of making an attempt to join the MIS in case of no concurrent attempts to do so.

This algorithm is not self-stabilizing for two reasons. First, it works just if at the beginning of the algorithm the probability of beeping of each vertex v is $p_1(v) = 1/2$, and the analysis of the convergence time strongly relies on that. Second, the presence of phases with two rounds requires that the vertices are synchronized modulo two. These reasons are also the main obstacle to making it self-stabilizing. Moreover, in self-stabilizing algorithms, vertices must be able to detect errors, e.g., when a fault forces a vertex to change its state from MIS to non-MIS, and hence stable vertices cannot be silent after they stabilized.

In order to design a self-stabilizing MIS algorithm for the full-duplex beeping model, achieving a $O(\log n)$ global round complexity w.h.p., we dispense with the idea of phases and we change the details of updating the beeping probabilities $p_t(v)$ to overcome the mentioned issues. While keeping the idea of increasing and decreasing the beeping probability depending on whether a beep was received, we refine this behavior in a significant way. As before, when a vertex v beeps while hearing no beeps at the same time it attempts to join the MIS. To signal this to neighboring vertices, vertex v keeps beeping, i.e., it sets its beeping probability $p_t(v)$ to 1. If such a vertex hears a beep in one of the following rounds, it does not immediately give up its attempt to join the MIS, but it keeps beeping with probability 1 for some fixed number rounds. Only after hearing a beep in a certain number of rounds, the vertex changes its behavior back to halving its beeping probability in every round it hears a beep. Furthermore, if the beeping probability decreases over a fixed threshold, the vertex sets its beeping probability to 0 and stops beeping. The complete code is shown in Algorithm 1.

Algorithm 1: Self-stabilizing version of Jeavons, Scott and Xu’s Algorithm [18]

state: $\ell \in \{-\ell_{\max}(v), \dots, \ell_{\max}(v)\}$

in each round $t = 1, 2, \dots$ **do**

- if** $\ell < \ell_{\max}(v)$ **then**
 - $\text{beep} \leftarrow \text{true}$ with probability $\min\{2^{-\ell}, 1\}$ and $\text{beep} \leftarrow \text{false}$ otherwise
- else** $\text{beep} \leftarrow \text{false}$
- if** beep **then** send signal to all neighbors
receive any signals sent by neighbors
- if** *any signal received* **then**
 - $\ell \leftarrow \min\{\ell + 1, \ell_{\max}(v)\}$
- else if** beep **then**
 - $\ell \leftarrow -\ell_{\max}(v)$
- else** $\ell \leftarrow \max\{\ell - 1, 1\}$

To implement the described behavior, each vertex maintains an integral state variable ℓ , which we call *level*. The value of ℓ for vertex v is in the range $-\ell_{\max}(v), \dots, \ell_{\max}(v)$, where $\ell_{\max}(v)$ is a fixed value that depends on the vertex’s knowledge of some graph parameters. We will see that this value has a strong influence on the analysis of the stabilization time. The value

of $\ell_t(v)$ of vertex v in round t implies the beeping probability $p_t(v)$ of v similar to an activation function in an artificial neural network (see Figure 1). As long as $\ell_t(v) \leq 0$ vertex v beeps and $p_t(v) = 1$, if $\ell_t(v) = \ell_{\max}(v)$ it stops beeping and $p_t(v) = 0$, otherwise $p_t(v) = 2^{-\ell_t(v)}$.

In each round t each vertex v updates the value of $\ell_t(v)$ as follows. If v hears a beep then its level increases $\ell_{t+1}(v) = \min\{\ell_t(v) + 1, \ell_{\max}(v)\}$. Otherwise, $\ell_{t+1}(v) = \max\{\ell_t(v) - 1, 1\}$ unless v was beeping in round t , in this case $\ell_{t+1}(v) = -\ell_{\max}(v)$. Note that the only way the level of a vertex v can decrease below 0 is if v beeps without beeping neighbors. We observe that Algorithm 1 is self-stabilizing if its convergence is guaranteed for every initial value of the levels.

The update rules of the algorithm guarantee that, once the level's value of a vertex v is $-\ell_{\max}(v)$ and each of v 's neighbors w has level's value $\ell_{\max}(w)$, then v is such that $p_t(v) = 1$ and all the neighbors u of v are such that $p_t(u) = 0$. This guarantees that v and its neighbors will not change their status as long as no faults occur, and hence they are stable. In this case, v will be a MIS vertex and the neighbors become non-MIS vertices. Also, this strategy allows all vertices to detect faults and react accordingly. But foremost, it allows to determine the stabilization time.




Figure 1: Beeping probability $p_t(v)$ of v based on value of $\ell_t(v)$.

The result and the analysis of the algorithm depend on the values $\ell_{\max}(v)$ of each vertex v , which in turn depends on the knowledge available to each vertex v . We state the results in the following theorems, and notice that we denote with $\deg_2(v) = \max_{u \in N(v) \cup \{v\}} \deg(u)$ the maximum degree in the 1-hop neighborhood of v .

Theorem 1. *For any n -vertex graph G , Algorithm 1 computes an MIS, starting from an arbitrary configuration, within $O(\log n)$ rounds w.h.p., provided that $\ell_{\max}(v) = \ell_{\max} \in [\log \Delta + c_1, c_2 \log n]$ for each vertex v and constants $c_1 \geq 15$ and $c_2 > 0$.*

Theorem 2. *For any n -vertex graph G , Algorithm 1 computes an MIS, starting from an arbitrary configuration, within $O(\log n \cdot \log \log n)$ rounds w.h.p., provided that $\ell_{\max}(v) \in [2 \log \deg(v) + c_1, c_2 \log n]$ for each vertex v and constants $c_1 \geq 30$ and $c_2 > 0$.*

Corollary 3. *There exists a variant of Algorithm 1 for the beeping model with two beeping channels such that, for any n -vertex graph G , it computes an MIS, starting from an arbitrary configuration, within $O(\log n)$ rounds w.h.p., provided that $\ell_{\max}(v) \in [2 \log \deg_2(v) + c_1, c_2 \log n]$, for each vertex v and any constants $c_1 \geq 15$ and $c_2 > 0$.*

Roadmap

The rest of the paper is organized as follows. Section 3 contains notations, preliminary definitions, the statements of the two key Lemmas 7 and 8, and an analysis outline. In Section 4 we give the proof of Theorem 1, and in Section 5 the proof of Theorem 2. The proofs of the key Lemmas 7 and 8 are in Section 6. Finally, in Section 7 we consider the model with two beeping channels and prove Corollary 3, and in Section 8 we summarize the conclusions and open problems.

3 Stabilization Time: Definitions and Analysis Outline

Let $G = (V, E)$ be a graph with n vertices. For each vertex $v \in V$, $N(v)$ denotes the set of v 's neighbors in G , and $\deg(v) = |N(v)|$ is the degree of v . Also, $N^+(v) = N(v) \cup \{v\}$ is the set of v 's neighbors and v itself. Let $\deg_2(v) = \max_{u \in N^+(v)} \deg(u)$ the maximum degree of all the vertices in $N^+(v)$.

We introduce a few random variables that are used to describe the random process generated by the execution of Algorithm 1. If we denote with $\ell_t(v)$ the level of vertex $v \in V$ at the beginning of round $t \geq 1$, the process is completely described by the random variables $\{\ell_t(v), v \in V\}_{t \geq 1}$. We denote with \mathcal{F}_t the filtration of the process until step t , which in particular gives us the values $\{\ell_t(v)\}_{v \in V}$.

We notice that in Algorithm 1 a vertex $v \in V$ is *stable* and permanently added to the MIS prior to round t if $\ell_t(v) = -\ell_{\max}(v)$ and, for all $u \in N(v)$, $\ell_t(u) = \ell_{\max}(u)$. Hence, if we define

$$\mu_t(v) = \min_{u \in N(v)} \frac{\ell_t(u)}{\ell_{\max}(u)},$$

which has value in $[-1, 1]$, we have that the set of vertices that have been added to the final MIS set before round t is defined by

$$I_t = \{v \in V : \ell_t(v) = -\ell_{\max}(v) \wedge \mu_t(v) = 1\}.$$

Moreover, the set of all stable vertices at the beginning of round t consists of the vertices in the MIS and their neighbors, so we define $S_t = I_t \cup N(I_t)$. We notice that the set of stable vertices is increasing in t , i.e. for each $t \geq 1$ we have that $S_t \subseteq S_{t+1}$. For any vertex $v \in V$, we denote with $p_t(v)$ the probability that v beeps during round t , which is

$$p_t(v) = \begin{cases} 1 & \text{if } \ell_t(v) \leq 0 \\ 2^{-\ell_t(v)} & \text{if } 0 < \ell_t(v) < \ell_{\max}(v) \\ 0 & \text{if } \ell_t(v) = \ell_{\max}(v). \end{cases}$$

We also denote with $b_t(v)$ a Bernoulli random variable which takes value 1 if v beeps in round t , i.e., $\mathbf{E}[b_t(v)] = p_t(v)$. We define $B_t(v) = \sum_{u \in N(v)} b_t(u)$ as the number of v 's neighbors that beep in round t and $d_t(v) = \mathbf{E}[B_t(v)] = \sum_{u \in N(v)} p_t(u)$ as the expected number of beeping neighbors of v in round t . We notice the following important lemma.

Lemma 4. *For any $t > \max_{w \in V} \ell_{\max}(w)$ and for any $v \in V$, we have $\ell_t(v) > 0$ or $\mu_t(v) > 0$.*

Proof. Let t_0 be the first round such that $\ell_{t_0}(v) > 0$ or $\mu_{t_0}(v) > 0$. First, we will prove that this condition keeps holding for all rounds $t \geq t_0$. Then, we will prove that $t_0 \leq \max_{w \in V} \ell_{\max}(w) + 1$.

Consider any round $t \geq t_0$ and assume that $\ell_t(v) > 0$ or $\mu_t(v) > 0$. This implies that $\mu_{t+1}(v) > 0$ or $\ell_{t+1}(v) > 0$. Indeed, assume that $\mu_t(v) \leq 0$. Then $\ell_t(v) > 0$ and at least one neighbor of v beeps in round t . Thus, $\ell_{t+1}(v) = \min\{\ell_t(v) + 1, \ell_{\max}(v)\} \geq \ell_t(v) > 0$, i.e., the condition of the lemma holds in round $t + 1$. Next consider the case that $\mu_t(v) > 0$. If v beeps in round t then all neighbors increase their value for ℓ , i.e., $\mu_{t+1}(v) \geq \mu_t(v) > 0$. If v does not beep in round t then $\ell_t(v) > 0$. Indeed, if no neighbor of v beeps then $\ell_{t+1}(v) = \max\{\ell_t(v) - 1, 1\} > 0$, and if at least one neighbor of v beeps then $\ell_{t+1}(v) = \min\{\ell_t(v) + 1, \ell_{\max}(v)\} \geq \ell_t(v) > 0$, i.e.,

the condition of the lemma holds in round $t + 1$.

Assume that $\ell_0(v) \leq 0$ and $\mu_0(v) \leq 0$. Then, in the first round v and all its neighbors beep. Hence, all these vertices increment their level by 1, i.e., $\ell_1(v) = \ell_0(v) + 1$ and $\mu_1(v) = \min_{u \in N(v)} \frac{\ell_0(u)+1}{\ell_{\max}(u)}$. Since $-\ell_{\max}(u) \leq \ell_0(u)$ for all vertices $u \in V$, there exists $t_0 \leq \max_{u \in N^+(v)} \ell_{\max}(u) + 1$, such that $\ell_{t_0}(v) > 0$ or $\mu_{t_0}(v) > 0$. This completes the proof. \square

Lemma 4 implies that in order to prove that our algorithm stabilizes within $O(\log n)$ rounds we can assume that $\ell_t(v) > 0$ or $\mu_t(v) > 0$ for all rounds $t \geq 0$. This is because $\max_{w \in V} \ell_{\max}(w) \in O(\log n)$. Hence, we can ignore the initial $\max_{w \in V} \ell_{\max}(w)$ rounds and start our analysis after those rounds. In particular, $\ell_t(u) \leq 0$ implies $\mu_t(u) > 0$.

We define a vertex to be *prominent* if it has negative or zero level, and a round to be *platinum* for some vertex v if some of v 's neighbors is prominent.

Definition 5 (Prominent Vertices and Platinum Rounds). A vertex $u \in V$ is *prominent* in round t if $\ell_t(u) \leq 0$. The set of prominent vertices in round t is denoted with PM_t . Moreover, we say that round t is a *platinum round of vertex v* if there is a prominent vertex in $N^+(v)$, i.e. $u \in N^+(v) \cap PM_t$. We denote with $P_{t,k}(v)$ the number of platinum rounds of vertex v during rounds $\{t, \dots, t+k\}$.

We notice that, since we assume $t > \max_{w \in V} \ell_{\max}(w)$, then Lemma 4 implies that each platinum round t is such that $\mu_t(u) > 0$, so all u 's neighbors are silent with nonzero probability during round t . Remember that the only possibility for the level of vertex u to become less or equal to 0 is when u beeps while no neighbor of u is beeping. This leads to the next lemma.

Lemma 6. *If $t > \max_{w \in V} \ell_{\max}(w)$ is a platinum round for vertex v there exists a vertex $u \in N^+(v)$ and a round t' with $t - \ell_{\max}(u) \leq t' \leq t$ in which u was beeping without beeping neighbors and $\ell_{t'+1}(u) = -\ell_{\max}(u)$.*

We define, for any $v \in V$ and $t \geq 1$, the quantities

$$\eta_t(v) = \sum_{u \in N(v) \setminus S_t} 2^{-\ell_{\max}(u)} \quad \text{and} \quad \eta'_t(v) = \sum_{\substack{u \in N(v) \setminus S_t: \\ \ell_{\max}(u) > \ell_{\max}(v)}} 2^{-\ell_{\max}(v)}.$$

For the moment, the definitions of $\eta_t(v)$ and $\eta'_t(v)$ are rather technical, but they will be used to upper bound the value of $d_{t+1}(v)$. We notice that $\eta_t(v) \geq \eta'_t(v)$ and also that $\eta_t(v)$ and $\eta'_t(v)$ are both decreasing in t , since $S_t \subseteq S_{t+1}$.

The following two lemmas are the key to prove Theorems 1 and 2 and Corollary 3, their proofs are deferred to Section 6. For a fixed $v \in V$ the next lemma tells us how many rounds we have to wait in order to have a platinum round of v .

Lemma 7 (Lower Bound on Platinum Rounds). *Assume that $\ell_{\max}(w) \geq \log \deg(w) + 4$ for all $w \in V$. Consider a vertex $v \in V$ and a round $t > \max_{w \in V} \ell_{\max}(w)$ such that t is not a platinum round of v , and $\eta_t(v) \leq 0.0001$. Let $\tau^{(v)}(t) = \min\{m \geq 0 : P_{t,m}(v) \geq 1\}$ and $\gamma = e^{-30}$. Then for any $k \geq 70\ell_{\max}(v)$*

$$\Pr \left[\tau^{(v)}(t) \geq k \mid \mathcal{F}_t \right] \leq e^{-\gamma^k}.$$

We notice that, if $\ell_{\max}(w)$ is constant over the vertices $w \in V$, i.e. $\ell_{\max}(w) = \ell_{\max}$ for every vertex $w \in V$, then the existence of a platinum round t of v such that $t > \ell_{\max}$ is by Lemma 6

sufficient to guarantee that v will be stable at the latest in round $t + \ell_{\max}$. Indeed, Lemma 6 implies the existence of a round $1 \leq t' \leq t$ and a vertex $u \in N^+(v)$ such that u was beeping in round t' without beeping neighbors, and so $\ell_{t'+1}(u) = -\ell_{\max}$. This implies that u beeps in the following ℓ_{\max} rounds, during which all neighbors of u will increase their level until they reach maximum level ℓ_{\max} . This implies that u is such that $\ell_{t+\ell_{\max}}(u) \leq 0$ and $\mu_{t+\ell_{\max}}(u) = 1$, and hence $u, v \in S_{t+\ell_{\max}}$ and $u \in I_{t+\ell_{\max}}$.

However, when $\ell_{\max}(w)$ is not constant, the analysis becomes considerably more complicated, since the existence of a platinum round of v does not imply necessarily the subsequent stabilization of v . We remark that a round t is *platinum for vertex v* if there is a prominent vertex $u \in N^+(v)$. It is immediate that $I_t \subseteq PM_t$. Consider now some round $t > \max_{w \in V} \ell_{\max}(w)$ which is platinum for v , and let $u \in N^+(v)$ be a prominent vertex. After round t , two things may happen:

- (i) In some round $t + m$ with $m \geq 1$, u is no longer prominent, and hence $u \notin I_{t+m}$ and u, v may not be stable in round $t + m$;
- (ii) In some round $t + m$ with $m \geq 0$, vertex u is prominent and all its neighbors have reached the maximum level, i.e. $\mu_{t+m}(u) = 1$, so and $u \in I_{t+m}$ and $u, v \in S_{t+m}$.

In the next lemma, we characterize the distribution of the rounds t for the above two cases. We define

$$\begin{aligned}\sigma_{\text{out}}^{(u)}(t) &= \min\{m \geq 1 : u \notin PM_{t+m}\}, \\ \sigma_{\text{in}}^{(u)}(t) &= \min\{m \geq 0 : u \in I_{t+m}\}, \\ \sigma^{(u)}(t) &= \min\{\sigma_{\text{out}}^{(u)}(t), \sigma_{\text{in}}^{(u)}(t)\}.\end{aligned}$$

Lemma 8 (Stopping Times for Platinum Rounds). *Assume that $\ell_{\max}(w) \geq \log \deg(w) + 4$ for all $w \in V$. Consider a round $t > \max_{w \in V} \ell_{\max}(w)$ and a vertex $u \in PM_t$. Then, for any $x \geq 0$,*

- (a) $\Pr \left[\sigma^{(u)}(t) = \sigma_{\text{in}}^{(u)}(t) \wedge \sigma^{(u)}(t) < \max_{w \in N(u)} \ell_{\max}(w) \mid \mathcal{F}_t \right] \geq 3^{-\eta'_t(u)};$
- (b) $\Pr \left[\sigma^{(u)}(t) = \sigma_{\text{out}}(t) \wedge \sigma^{(u)}(t) > \ell_{\max}(u) + x \mid \mathcal{F}_t \right] \leq \eta'_t(u) 2^{-x}.$

3.1 Overview of the analysis

We first give an overview of the proofs of Lemmas 7 and 8, and then we will see how to use these results to prove Theorems 1 and 2 and Corollary 3. The proof of Lemma 7 has as a starting point the proof in [14], but then it develops differently. First, as in [14], we define a further type of round called *golden round*, which has constant probability of becoming platinum in the subsequent round. We prove that, for any vertex v in any fixed interval of rounds of length $k = \Omega(\ell_{\max}(v))$, we have a constant fraction of golden rounds with probability at least $1 - e^{-\Omega(k)}$, conditioned on the absence of platinum rounds during that time interval. To prove the latter, as in [14], we analyze the development of the function $d_t(v)$ – the expected number of beeping neighbors of v in round t – during this time. Note that platinum rounds and the conditioning were not considered in [14] and are essential in our proof and setting.

The proof of Lemma 8 relies on Lemma 6. Assuming that u is prominent at time t , we characterize the probabilities with which, after round t , u reaches again a positive level or stabilizes.

From Lemma 6 there exists a round $t - \ell_{\max}(u) \leq t' \leq t$ where we have that $\ell_{t'+1}(u) = -\ell_{\max}(u)$ and then trivially $d_{t'+1}(u) = \sum_{w \in N(u)} p_{t'+1}(u) \leq \deg(u)$. Then, in the subsequent $\ell_{\max}(u)$ rounds, vertex u keeps beeping regardless the behavior of the vertices in $N(u)$, and this implies that $d_{t+1+\ell_{\max}(u)}(u) \leq \eta'_{t+1}(u)$. We will see that this implies that the vertices in $N(u)$ will reach their maximal level with probability at least $3^{-\eta'_{t+1}(u)}$, and so in this case the platinum round leads to the stabilization of u . On the other hand, part (b) of the lemma follows from the observation that, after the first $\ell_{\max}(u)$ rounds after t , the probability that some vertex in $N(u)$ beeps decreases in each round of a constant factor.

Theorem 1 and Corollary 3 follow from an observation, already stated above, which is that if $\ell_{\max}(w)$ is constant over w then, for each vertex v , one platinum round is sufficient to guarantee the stabilization of v . Moreover, the choices of $\ell_{\max}(w)$ specified in the theorems guarantee that $\eta_t(v) \leq 0.0001$ for every v and $t \geq 1$, and so Lemma 7 can always be used for each non-platinum round t , and implies that we have to wait at most $O(\log n)$ rounds to have a platinum round for each vertex v , w.h.p.

The proof of Theorem 2 is harder. In this case, we can have several sequences of consecutive platinum rounds, intermittent by sequences of consecutive non-platinum rounds, until we reach a platinum round leading to the stabilization of the vertex. The analysis relies on two main parts:

- (1) We split the vertices in $O(\log \log n)$ sets V_i . Before analyzing the stabilization of a vertex $v \in V_i$, we wait for round T_i until all vertices in $\cup_{j < i} V_j$ have stabilized. The sets V_i are defined according to the values $\ell_{\max}(v)$ of the vertices. According to the definition of T_i , we can apply, for each round $t \geq T_i$, Lemmas 7 and 8 to vertices in V_i .
- (2) We then prove that, after round T_i , each vertex $v \in V_i$ stabilize in $O(\log n)$ additional rounds w.h.p. The analysis of the latter statement strongly relies on Lemmas 7 and 8, which characterize the length of the times interval of the non-platinum rounds, of the platinum rounds, and the number of platinum rounds not leading to the stabilization of vertex v .

4 Warm Up: Knowledge of the Maximum Degree Δ

The following proof is a warm-up for the general case. It is directly implied by Lemma 7 and the choice of $\ell_{\max}(v)$.

Proof of Theorem 1. As already mentioned, since $\ell_{\max}(v)$ is defined independently of v , each vertex v just requires a single platinum round to become stable in at most ℓ_{\max} rounds. Indeed, for each $v \in V$ and each $t \geq 1$,

$$\eta_t(v) \leq \sum_{u \in N(v)} 2^{-\log \Delta - 15} \leq 2^{-15} \leq 0.0001 \quad \text{and} \quad \eta'_t(v) = 0.$$

This implies that, if $t = 2\ell_{\max}$ and $v \in V$, we have $\eta_t(v) \leq 0.0001$. Hence, by Lemma 7, if we take $m = 2\gamma^{-1} \log n$ (where γ is defined in Lemma 7), we have that $\Pr[\tau^{(v)}(t) \leq m \mid \mathcal{F}_t] \geq 1 - 1/n^2$, and so $P_{t,m}(v) \geq 1$ with probability at least $1 - 1/n^2$. Then, from Lemma 8(a), and since $\eta'_{t+m}(v) = 0$, we have that, given \mathcal{F}_{m+t} , the vertex v is stable after at most ℓ_{\max} rounds with probability 1. Hence, vertex v is stable with probability $1 - 1/n^2$ after $t + m + \ell_{\max}$ rounds,

and since $\ell_{\max} = O(\log n)$ we have that $t + m + \ell_{\max} = O(\log n)$. The theorem follows from the union bound over all the vertices. \square

5 Knowledge of the degrees: Proof of Theorem 2

In this section we prove Theorem 2. First, we prove the following lemma.

Lemma 9. *Assume that $\ell_{\max}(w) \geq \log \deg(w)$ for every $w \in V$ and that $\max_{w \in V} \ell_{\max}(w) \leq c \log n$ for some $c = O(1)$. Consider a vertex $v \in V$ and a round $t > \max_{w \in V} \ell_{\max}(w)$ such that $\eta_t(v) \leq 0.0001$ and such that, for each $u \in N(v) \setminus S_t$, $\ell_{\max}(v) \leq 2\ell_{\max}(u)$. Then, there exists a constant $M = O(1)$ such that $\Pr[v \in S_{t+m} \mid \mathcal{F}_t] \geq 1 - 1/n^2$, provided $m = M \log n$.*

Proof. We fix the execution up to the end of round t , so we do not have to condition probabilities on \mathcal{F}_t . We consider the sequence of rounds (which may also be infinite, with $J = +\infty$)

$$t \leq t + \tau_1 \leq t + \tau_1 + \sigma_1 = m_1 + t \leq t + m_1 + \tau_2 \leq \dots \leq t + m_{J-1} + \tau_J \leq t + m_{J-1} + \tau_J + \sigma_J = t + m_J$$

and the corresponding sequence of vertices $v_1, v_2, \dots, v_J \in N^+(v) \setminus S_t$ such that

1. $t + m_{i-1} + \tau_i$ is platinum for v and $v_i \in PM_{t+m_{i-1}+\tau_i} \cap N^+(v)$ for each $i = 1, \dots, J$;
2. $m_i = m_{i-1} + \tau_i + \sigma_i$ is such that $v_i \notin PM_{t+m_i}$ for each $i = 1, \dots, J-1$;
3. $J = \min\{h \geq 1 : v_h \in I_{m_h+t}\}$, hence $v_J \in I_{t+m_J}$ and $v \in S_{t+m_J}$. If v never stabilizes, we define $J = +\infty$ and the sequence v_1, v_2, \dots has infinite length.

We observe that σ_i and τ_i are defined such that

$$\tau_i = \tau^{(v_i)}(t + m_{i-1}), \quad \sigma_i = \sigma^{(v_i)}(t + m_{i-1} + \tau_i) \quad \text{and} \quad \sigma_J = \sigma_{\text{in}}^{(v_J)}(t + m_{J-1} + \tau_J). \quad (1)$$

The proof is organized as follows: we first show that $\sum_{i=1}^J \ell_{\max}(v_i) \leq 7 \log n$ and $\sigma_J \leq \max_{w \in V} \ell_{\max}(w)$ with probability at least $1 - 1/n^3$, and secondly that, under the latter event, $\sum_{i=1}^J \sigma_i + \tau_i \leq M \log n$ for some $M = \Theta(1)$ and with probability at least $1 - 1/n^3$.

We notice that

$$\Pr \left[\sum_{i=1}^J \ell_{\max}(v_i) \geq 7 \log n \vee \sigma_J \geq \max_{w \in V} \ell_{\max}(w) \right] \quad (2)$$

$$\leq \Pr \left[(J \geq h \wedge \sum_{i=1}^h \ell_{\max}(v_i) \geq 7 \log n) \vee \sigma_J \geq \max_{w \in V} \ell_{\max}(w) \right], \quad (3)$$

where we define $h = \min\{j \geq 1 : \sum_{i=1}^j \ell_{\max}(v_i) \geq 7 \log n\}$. We notice that, since $\min_{v \in V} \ell_{\max}(v) \geq 1$, for the minimality of h we have that $h \leq 7 \log n$. Lemma 8 (a) and (1) yield that the above probability is bounded by

$$\Pr \left[\bigcap_{i=1}^h \{v_i \notin PM_{t+i} \vee \sigma_i \geq \max_{w \in V} \ell_{\max}(w)\} \right] \leq \prod_{i=1}^h \left(1 - 3^{-\eta'_t(v_i)} \right) \leq 2 \prod_{i=1}^h \eta'_t(v_i). \quad (4)$$

Moreover, we have that

$$\eta'_t(v_i) \leq \sum_{\substack{w \in N(v_i): \\ \ell_{\max}(w) > \ell_{\max}(v_i)}} 2^{-\ell_{\max}(v_i)} \leq \deg(v_i) \cdot 2^{-\ell_{\max}(v_i)} \leq \frac{\deg(v_i)}{2^{\ell_{\max}(v_i)/2}} 2^{-\ell_{\max}(v_i)/2} \leq 2^{-\ell_{\max}(v_i)/2},$$

where the last inequality follows from the fact that $\ell_{\max}(v_i) \geq 2 \log \deg(v_i)$. Hence, from the latter inequality, (3) and (4) we have that

$$\Pr \left[\sum_{i=1}^J \ell_{\max}(v_i) \geq 7 \log n \vee \sigma_J \geq \max_{w \in V} \ell_{\max}(w) \right] \leq 2 \prod_{i=1}^h \eta'_t(v_i) \leq 2^{-\sum_{i=1}^h \ell_{\max}(v_i)/2+1} \leq \frac{1}{n^3}, \quad (5)$$

where the last inequality follows from the fact that $\sum_{i=1}^h \ell_{\max}(v_i) \geq 7 \log n$. We will refer to the event in the first expression in (5) as $\neg A$, and for the rest of the proof we assume that A is happening. We recall that the event A implies that $J \leq 7 \log n$.

Next we show that provided A happens for some $M = O(1)$ we have

$$m_J - t = \sum_{i=1}^J \tau_i + \sigma_i \leq M \log n$$

with probability at least $1 - 1/n^3$ by proving separately that

$$(i) \quad \Pr \left[\sum_{i=1}^J \tau_i \geq \frac{M \log n}{2} \mid A \right] \leq \frac{1}{2n^3};$$

$$(ii) \quad \Pr \left[\sum_{i=1}^J \sigma_i \geq \frac{M \log n}{2} \mid A \right] \leq \frac{1}{2n^3}.$$

The facts (i) and (ii), together with (4), prove the lemma. Indeed, if $m = M \log n$, we have that

$$\begin{aligned} \Pr[v \in S_{t+m}] &\geq \Pr \left[\sum_{i=1}^J (\sigma_i + \tau_i) \leq m \right] \geq \Pr \left[\sum_{i=1}^J \sigma_i + \tau_i \leq m \mid A \right] \Pr[A] \\ &\geq \left(1 - \frac{1}{n^3} \right)^2 \geq 1 - \frac{1}{n^2}, \end{aligned}$$

where the last inequality follows by (i) and (ii), since

$$\begin{aligned} \Pr \left[\sum_{i=1}^J \sigma_i + \tau_i \geq M \log n \mid A \right] &\leq \Pr \left[\sum_{i=1}^J \sigma_i \geq M \log n / 2 \mid A \right] + \Pr \left[\sum_{i=1}^J \tau_i \geq M \log n / 2 \mid A \right] \\ &\leq \frac{1}{n^3}, \end{aligned}$$

and from (4) that implies that $\Pr[A] \geq 1 - 1/n^3$. Now we proceed with the proof of (i) and (ii).

Proof of inequality (i) From (1) and Lemma 7, we have that each τ_i is stochastically dominated by $70\ell_{\max}(v) + X_i$, where X_i is a geometric random variable with parameter $p = 1 - e^{-\gamma}$,

where $\gamma = e^{-30}$. Then, we have that, conditioning on the event A ,

$$\begin{aligned} \Pr \left[\sum_{i=1}^J X_i \geq \frac{J}{p} + \frac{4 \log n}{p^2} \right] &= \Pr \left[\text{Bin} \left(\frac{J}{p} + \frac{4 \log n}{p^2}, p \right) \leq J \right] \leq \exp \left(\left(\frac{4 \log n}{p} \right)^2 \frac{1}{\frac{J}{p} + \frac{4 \log n}{p^2}} \right) \quad (6) \\ &\leq \exp \left(\frac{16 \log^2 n}{Jp + 4 \log n} \right) \leq \frac{1}{2n^3}, \end{aligned}$$

where the first inequality follows from Lemma 15, and the last inequality follows from the fact that $p = 1 - e^{-\gamma}$ and that $J \leq 7 \log n$ under the event A . This implies that, with probability at least $1 - 1/(2n^3)$,

$$\sum_{i=1}^J \tau_i \leq \sum_{i=1}^J 70 \ell_{\max}(v) + X_i \leq 70J \cdot \ell_{\max}(v) + \frac{J}{p} + \frac{4 \log n}{p^2} \leq \frac{M \log n}{2},$$

where we define $M/2 = 70J \ell_{\max}(v) \log^{-1} n + 9/p + 4/p^2 + c = O(1)$. Indeed, we have that, since $v_i \in N^+(v) \setminus S_t$, then $\ell_{\max}(v) \leq \ell_{\max}(v_i)$ for each $i = 1, \dots, J$, so

$$70h \ell_{\max}(v) \leq 70 \cdot 140 \sum_{i=1}^J \ell_{\max}(v_i) \leq 1260 \log n,$$

where the last inequality follows from the fact that we are under the event A .

Proof of inequality (ii) From (1) and Lemma 8 (b) we have that, for each $i = 1, \dots, J-1$, the random variable σ_i is stochastically dominated by $\ell_{\max}(u) + Y_i$, where Y_i is a geometric random variable with parameter $1/2$. As in (6), under A we have that, using again Lemma 15,

$$\Pr \left[\sum_{i=1}^{J-1} Y_i \geq 2J + 8 \log n \right] \leq \exp \left(\frac{16 \log^2 n}{J/2 + 4 \log n} \right) \leq \frac{1}{2n^3},$$

where the last inequality follows since $J \leq 7 \log n$ under event A . Since A implies that $\sigma_J \leq \max_{w \in N(u)} \ell_{\max}(w) \leq c \log n$, we have that, conditioning on A , with probability at least $1 - 1/(2n^3)$

$$\sum_{i=1}^J \sigma_i \leq \sum_{i=1}^{J-1} (Y_i + \ell_{\max}(v_i)) + \sigma_J \leq 2J + 8 \log n + c \log n \leq \frac{M \log n}{2}. \quad \square$$

We now can proceed with the proof of Theorem 2.

Proof of Theorem 2. We have that $2 \log \deg(w) + 30 \leq \ell_{\max}(w) \leq c_2 \log n$ for every $w \in V$ and some $c_2 = O(1)$. For each $i = 1, \dots, c_2 \log \log n$, we consider the following subsets of vertices

$$V_i = \{v \in V : \ell_{\max}(v) \in [2^i, 2^{i+1}]\}.$$

Let T_i be the round until all the vertices in $\cup_{j \leq i} V_j$ are stabilized, i.e.

$$T_i = \min\{t' \geq 1 : \cup_{j \leq i} V_j \subseteq S_{t'}\}.$$

We have that, for each vertex $v \in V_{i+1}$ and each $t \geq T_i$,

$$2 \ell_{\max}(u) \geq \ell_{\max}(v) \quad \forall u \in N^+(v) \setminus S_t.$$

Indeed, we have $u \notin \cup_{j \leq i} V_j$ since $u \notin S_t$ and $t \geq T_i$. Hence, $\ell_{\max}(u) \geq 2^{i+1}$. Since $v \in V_{i+1}$, $\ell_{\max}(v) \leq 2^{i+2}$ and so $2\ell_{\max}(u) \geq \ell_{\max}(v)$. We also have, for each $t \geq T_i$ and each $v \in V_{i+1}$

$$\eta_t(v) \leq \sum_{u \in N(v) \setminus S_t} 2^{-\ell_{\max}(u)} \leq \sum_{u \in N(v) \setminus S_t} 2^{-\ell_{\max}(v)/2} \leq \sum_{u \in N(v) \setminus S_t} \frac{1}{\deg(v)} 2^{-15} \leq 0.0001,$$

where the last inequality follows from the fact that $2\ell_{\max}(u) \geq \ell_{\max}(v)$, and the last inequality since $\ell_{\max}(v) \geq 2 \log \deg(v) + 30$.

We can now apply Lemma 9, if $t \geq \max\{T_i, \max_{w \in V} \ell_{\max}(w) + 1\}$, to all vertices $v \in V_{i+1}$, obtaining (with an union bound over all the vertices in V_i) the existence of a round $m_i = O(\log n)$ such that $\Pr[V_{i+1} \subseteq S_{t+m_i}] \geq 1 - 1/n$. Applying this argument iteratively for each $i = 1, \dots, c_2 \log \log n$, we obtain the existence of a round $m = \sum_{i=1}^{\log \log n} m_i = O(\log n \cdot \log \log n)$ such that all vertices are stable w.h.p. at round m . \square

6 Proof of Key Lemmas

6.1 Lower Bound on Platinum Rounds (Proof of Lemma 7)

Before proving Lemma 7, we introduce some definitions and preliminary lemmas.

Definition 10 (Light vertices). A vertex $v \in V$ is called *light in round t* if $\mu_t(v) > 0 \wedge (d_t(v) \leq 10 \vee \ell_t(v) \leq 0)$. We denote with L_t the set of light vertices at round t and with $H_t = V \setminus L_t$ the set of *heavy (non-light) vertices at round t* .

Intuitively, a light vertex v has a positive, constant probability of not receiving a beep signal during round t and, hence, if $p_t(v)$ is large enough, v has a constant probability of beeping without beeping neighbors during round t . We remark that the condition $\mu_t(v) > 0$ is necessary since, if $\mu_t(v) = 0$, the vertex v hears a beep during round t with probability 1.

We denote with $d_t^L(v) = \sum_{u \in N(v) \cap L_t} p_t(u)$ the expected number of light and beeping neighbors of v in round t , and with $d_t^H(v) = \sum_{u \in N(v) \cap H_t} p_t(u)$ the expected number of beeping heavy neighbors of v in round t . We notice that $d_t(v) = d_t^L(v) + d_t^H(v)$.

Definition 11 (Golden Rounds). Round t is a *golden round of vertex v* if one of the following happens:

- (a) $\ell_t(v) \leq 1$ and $d_t(v) \leq 0.02$;
- (b) $d_t^L(v) > 0.001$.

We denote with $G_{t,k}(v)$ the number of golden rounds of vertex v in the rounds $\{t, \dots, t+k\}$.

In the next section, we will give a lower bound on the number of golden rounds.

6.1.1 Lower Bound on Golden Rounds

Lemma 12. Assume that $\ell_{\max}(w) \geq \log \deg(w) + 4$ for all $w \in V$. Consider a vertex $v \in V$ and a round $t > \max_{w \in V} \ell_{\max}(w)$ such that t is not a platinum round of v , and $\eta_t(v) \leq 0.0001$. Let $\tau^{(v)}(t)$ be defined as in Lemma 7. Then, we have that, for any $k \geq 70 \cdot \ell_{\max}(v)$,

$$\Pr \left[G_{t,k}(v) \leq 0.05k \wedge \tau^{(v)}(t) > k \mid \mathcal{F}_t \right] \leq e^{-k/100}.$$

We notice that, if round t is not a platinum round of v , every round $s \in [t, \tau^{(v)}(t)]$ is also not a platinum round of v , since the only way a vertex in $N^+(v)$ can take a negative level is by beeping without beeping neighbors, and $\tau^{(v)}(t)$ is the first round that happens. The proof of Lemma 12 strongly relies on the following lemma.

Lemma 13. *Let $v \in V$ and $t > \max_{w \in V} \ell_{\max}(w)$ such that round t is not a platinum round of v and $\eta_t(v) \leq 0.0001$.*

- (a) *If $d_t(v) \leq 0.02$, then with probability at least 0.97, we have that $\ell_{t+1}(v) \leq \max\{1, \ell_t(v) - 1\}$.*
- (b) *If $d_t(v) > 0.01$ and $d_t^L(v) < 0.01d_t(v)$, then with probability at least 0.97, we have that $d_{t+1}(v) < 0.6d_t(v)$ or that $t + 1$ is a platinum round for v .*

Proof. We fix the execution up to the end of round t , so we do not have to condition probabilities on \mathcal{F}_t . In what follows, we prove separately the two items.

Proof of (a) Since $d_t(v) = \sum_{u \in N(v)} p_t(u) \leq 0.02$ it follows that $p_t(u) \leq \frac{1}{2}$ for all $u \in N(v)$. Thus, the probability that no neighbor of v beeps is at least $\prod_{u \in N(v)} (1 - p_t(u)) \geq 4^{-d_t(v)} \geq 0.97$. Hence, $\Pr[\ell_{t+1}(v) \leq \max\{\ell_t(v) - 1, 1\}] \geq 0.97$.

Proof of (b) Since t is not a platinum round of v , we have that for each $u \in N^+(v)$, $\ell_t(u) \geq 1$. Moreover, we notice that there may be in round t a beeping vertex $u \in N^+(v)$ that in does not receive a signal, and so $\ell_{t+1}(u) = -\ell_{\max}(u) \leq 0$.

For any vertex $u \in N^+(v)$, we have the following upper bounds for $p_{t+1}(u)$ (recall that $\ell_t(u) > 0$ since t is not a platinum round of v):

$$p_{t+1}(u) \leq \begin{cases} 2^{-\ell_{\max}(u)+1} & \text{if } \ell_t(u) = \ell_{\max}(u) \text{ and } u \notin S_t \\ 0 & \text{if } \ell_t(u) = \ell_{\max}(u) \text{ and } u \in S_t \\ \frac{p_t(u)}{2} & \text{if } B_t(u) \geq 1 \text{ and } \ell_t(v) \neq \ell_{\max}(v) \\ 2p_t(u) & \text{if } B_t(u) = b_t(u) = 0 \text{ and } \ell_t(u) \neq \ell_{\max}(u) \\ 1 & \text{if } B_t(u) = 0, b_t(u) = 1 \text{ and } \ell_t(u) \neq \ell_{\max}(u) \end{cases}$$

The latter case, i.e. when $B_t(u) = 0$ and $b_t(u) = 1$ implies that $t + 1$ is a platinum round for v , and that $\ell_{t+1}(u) = -\ell_{\max}(u)$. Define $J_{t+1}(v)$ the set of such vertices, i.e. the set of vertices in $N(v)$ beeping in round t without beeping neighbors. Then,

$$\begin{aligned} d_{t+1}(v) &\leq \sum_{\substack{u \in N(v) \setminus S_t: \\ \ell_t(u) = \ell_{\max}(u)}} 2^{-\ell_{\max}(u)+1} + \sum_{\substack{u \in N(v): \\ B_t(u) \geq 1, \ell_t(u) \neq \ell_{\max}(u)}} \frac{p_t(u)}{2} + \sum_{\substack{u \in N(v): \\ B_t(u) = b_t(u) = 0}} 2p_t(u) + J_{t+1}(v) \\ &\leq 2\eta_t(v) + \sum_{u \in N(v) \cap H_t} p_t(u) \left(\frac{1}{2} + 2 \cdot \mathbb{1}_{B_t(u)=0} \right) + \sum_{u \in N(v) \cap L_t} 2p_t(u) + J_{t+1}(v). \end{aligned}$$

We notice that, since $d_t^L(v) = \sum_{u \in N(v) \cap L_t} p_t(u) \leq 0.01d_t(v)$ and $\eta_t(v) \leq 0.001$, we have that

$$d_{t+1}(v) \leq 0.0002 + 0.02d_t(v) + J_{t+1}(v) + \sum_{u \in N(v) \cap H_t} p_t(u) \left(\frac{1}{2} + 2 \cdot \mathbb{1}_{B_t(u)=0} \right)$$

We want to bound, for each $u \in N(v) \cap H_t$, the probability that $B_t(u) = 0$. Since $u \in H_t$, then $d_t(u) \geq 10$ or $\mu_t(u) = 0$. In the latter case, we know u has some neighbors $u' \in N(u)$ with $p_t(u') = 1$. Hence, we have that $\Pr[B_t(u) = 0] = 0$. In the former case, we have that none of u 's neighbors is beeping with probability at most

$$\prod_{w \in N(u)} (1 - p_t(w)) \leq e^{-d_t(u)} \leq e^{-10}.$$

Hence, we have that, for each $u \in H_t$, $\Pr[B_t(u) = 0] \leq e^{-10}$. So,

$$\mathbf{E} \left[\sum_{u \in N(v) \cap H_t} 2p_t(u) \mathbb{1}_{B_t(u)=0} \right] \leq \sum_{u \in N(v) \cap H_t} 2p_t(u) e^{-10}.$$

Markov's Inequality implies that $\sum_{u \in N(v) \cap H_t} 2p_t(u) \mathbb{1}_{B_t(u)=0} \leq 0.01 \sum_{u \in N(v) \cap H_t} 2p_t(u)$ with probability at least $1 - \frac{e^{-10}}{0.01} \geq 0.97$. Thus, with probability at least 0.97, we have that

$$d_{t+1}(v) \leq 0.0002 + 0.02d_t(v) + J_{t+1}(v) + 0.5d_t(v) + 0.02d_t(v) \leq 0.6d_t(v) + J_{t+1}(v), \quad (7)$$

where the last inequality follows by noticing that $d_t(v) > 0.01$ and hence $0.0002 < 0.02d_t(v)$. It follows the lemma, since we have that (7) implies that either $d_{t+1}(v) < 0.6d_t(v)$, or $J_{t+1}(v) > 0$, and hence $t + 1$ is a platinum round for v . \square

We are now ready to prove Lemma 12.

Proof of Lemma 12. Fix a vertex $v \in V$. We consider $k \geq 70\ell_{\max}(v)$ consecutive rounds, starting from a round t which is not a platinum round of v . Since $\eta_t(v)$ is decreasing in t , in all rounds $t + m$ with $m \geq 0$ we will have $\eta_{t+m}(v) < 0.0001$. We consider the following sets of rounds

$$D_{t,k}(v) = \{0 \leq m \leq k : d_{t+m}(v) > 0.2\}$$

$$E_{t,k}(v) = \{0 \leq m \leq k : d_{t+m}(v) > 0.1 \text{ and } d_{t+m}^L(v) \geq 0.1d_{t+m}(v)\}$$

$$F_{t,k}(v) = \{0 \leq m \leq k : d_{t+m}(v) > 0.1 \text{ and } d_{t+m}^L(v) < 0.1d_{t+m}(v)\}$$

$$H_{t,k}(v) = \{0 \leq m \leq k : d_{t+m}(v) < 0.2 \text{ and } \ell_{t+m}(v) \leq 1\}$$

We say that in some round t' we have a *wrong move* if none of the following conditions occurs

- (a) t or $t + 1$ is a platinum round of v ;
- (b) $\eta_t(v) > 0.0001$;
- (c) $d_t(v) \leq 0.01$ or $d_t(v) > 0.02$;
- (d) $d_t^L(v) \geq 0.01d_t(v)$;
- (e) $d_{t+1}(v) < 0.6d_t(v)$;
- (f) $\ell_{t+1}(v) \leq \max\{\ell_t(v) - 1, 1\}$;

From Lemma 13, we have that a vertex makes a wrong move with probability at most 0.03. Since the randomness of each round is independent of the others, we know by Chernoff's bound, that

in the rounds $\{t, t+1, \dots, t+k\}$ there are at most $0.04k$ wrong moves with probability at least $1 - e^{-k/100}$, and we will refer to this event with B .

In the rest of the proof, we assume that B happens, and we will see that it implies, deterministically, that $\tau^{(v)}(t) \leq k$ or that $G_{t,k}(v) \geq 0.1k$. So, we assume that $\tau^{(v)}(t) > k$ and we will prove that, under event B , this implies that $G_{t,k}(v) \geq 0.1k$. We remark that, if $\tau^{(v)}(t) > k$, for each $0 \leq m \leq k$, we have that $d_{t+m+1}(v) \leq 2d_{t+m}(v)$.

In what follows, we will prove that:

- (i) if $E_{t,k}(v) < 0.05k$, then $D_{t,k}(v) < 0.25k$
- (ii) if $D_{t,k}(v) < 0.25k$, then $H_{t,k}(v) > 0.06k$.

Proof of (i) We denote with $D'_{t,k}(v) = \{0 \leq m \leq k : d_{t+m}(v) > 0.1\}$ and we notice that $D'_{t,k}(v) = E_{t,k}(v) \cup F_{t,k}(v)$. Also let $h = |D_{t,k}(v)|$ and $h' = |D'_{t,k}(v)|$. Since the number of wrong moves is bounded by $0.04k$, and since $E_{t,k}(v) < 0.05k$ the number of rounds in $D'_{t,k}(v)$ in which $d_{t+m}(v)$ can double is at most $0.09k$, and in the rest of the rounds it will decrease of a factor of 0.6 .

In order to keep $d_{t+m}(v) > 0.2$, in a consecutive interval of rounds in $D'_{t,k}(v)$, the number of increasing moves must be at least $\log_{0.5}(0.6) > 0.7$ times the number of decreasing moves, and at most $\log_{5/3}(10d_t(v)) \leq 2 \log(10 \deg(v)) \leq 2 \log \deg(v) + 8$ decreases are used to decrease the initial value of $d_t(v)$ below 0.1 . Hence, the total number of rounds in $D_{t,k}(v)$ is at most

$$\begin{aligned} 0.09k + \frac{0.09}{0.7}k + 2 \log(\deg(v)) + 8 &\leq 0.22k + 2 \log \deg(v) + 8 \\ &\leq 0.22k + 2\ell_{\max}(v) \\ &\leq 0.22k + 0.03k \\ &= 0.25k. \end{aligned}$$

Proof of (ii) Since $|D_{t,k}(v)| \leq 0.25k$, the set $D_{t,k}^C(v) = \{0 \leq m \leq k : d_{t+m}(v) \leq 0.2\}$ contains at least $0.75k$ rounds. The number of wrong moves is bounded by $0.04k$, and in those rounds a wrong moves implies that $\ell_{t+m+1}(v) = \min\{\ell_{t+m}(v) + 1, \ell_{\max}(v)\}$. Moreover, we have that in the $0.25k$ rounds in which $d_{t+m}(v) > 0.2$, $\ell_{t+m+1}(v) \leq \min\{\ell_{t+m}(v) + 1, \ell_{\max}(v)\}$. Hence, $\ell_t(v)$ can increase in at most $0.29k$ rounds. If we do not have a wrong move, then $\ell_{t+m+1}(v) = \min\{1, \ell_{t+m}(v) - 1\}$, since we assumed that $\tau^{(v)}(t) > k$. Since $D_{t,k}^C(v)$ has at least $0.58k$ elements, and since the number of wrong moves is bounded by $0.29k$, the number of moves in which $\ell_{t+m}(v)$ decreases is at least $0.58k - 0.04k$, and the number of rounds in which $\ell_{t+m}(v)$ increases is at most $0.29k$. Hence, the number of increases is at least 1.8 times the number of decreases.

Denote the number of rounds in $D_{t,k}^C(v)$ where $\ell_{t+m}(v)$ decreases by U and those where it increases by D . Thus, $D + U \geq 0.75k$. In the worst case, each round with an increase follows a round with a decrease. Then, we still have $0.75k - 2D$ rounds with an increase left. Then, $0.73k - 2D = U - D \geq U/5 \geq 0.08k$. As it takes at most $\ell_{\max}(v)$ for $p_{t+m}(v)$ to reach $1/2$ we can say that we have at least $0.08k - \ell_{\max}(v) \geq 0.06k$ rounds where $\ell_{t+m}(v) = 1$ and $d_{t+m}(v) < 0.2$, hence in $H_{t,k}(v)$. \square

6.1.2 From Golden to Platinum Rounds

We first notice that, for each golden round t of v , the round $t+1$ is platinum for v with constant probability. Indeed, we have the following lemma.

Lemma 14. *Let $t > \max_{w \in V} \ell_{\max}(w)$ be not a platinum round of v , and let $\tau^{(v)}(t)$ be defined as in Lemma 7. Then, we have that, for each $t \leq s < \tau^{(v)}(t)$ which is golden for v , $\Pr[\tau^{(v)}(t) = s+1 \mid \mathcal{F}_s] \geq \gamma$, where $\gamma \geq e^{-27}$.*

Proof. Since s is such that $t \leq s < \tau^{(v)}(t)$, then s is not a platinum round of v , and so every vertex $u \in N^+(V)$ is such that $\ell_s(v) \geq 1$. In what follows, we prove that, with constant probability $\gamma > 0$, during round s there is a vertex $u \in N^+(v)$ such that $B_s(u) = 0$ and $b_s(u) = 1$. This implies that $\ell_{s+1}(u) = -\ell_{\max}(u)$ and that $\mu_{s+1}(u) > 0$, hence that $s+1$ is platinum for v . Since s is golden for v , we have that (a) or (b) of Definition 11 holds.

First, assume that (a) holds, thus $\ell_s(v) \leq 1$ and $d_s(v) \leq 0.02$. In this case, with probability at least 0.48, we have that $B_s(v) = 0$ and $b_s(v) = 1$ and so $s+1$ is platinum for v . Indeed, the expected number of beeping neighbors of v during round s is $d_s(v) \leq 0.02$. Therefore, for Markov's inequality, $\Pr[B_s(v) \geq 1] \leq 0.02$, and v beeps with probability at least $1/2$, and then the level of v is updated to zero with probability at least $\frac{1}{2} \cdot 0.98 > 0.48$.

We now assume that round s satisfies (b), therefore that $d_s^L(v) \geq 0.001$. We will prove that, in such types of rounds, with probability at least a constant γ , there is a beeping vertex $u \in N(v)$ with no beeping neighbors during round s . Let $k = |N(v) \cap L_s|$ be the number of light neighbors of v , and denote $\{w_1, \dots, w_k\} = N(v) \cap L_s$. We remark that all light vertices w are such that $\mu_s(u) > 0$ for each $u \in N(w)$ and $d_s(w_i) \leq 10$. We define \mathcal{E}_i the event indicating that vertex w_i is beeping during round s , and denote with $\mathcal{E} = \cup_i \mathcal{E}_i$. We have that,

$$\Pr[\mathcal{E}] \geq 1 - \prod_{w \in N(v) \cap L_s} (1 - p_t(w)) \geq 1 - e^{-\sum_{w \in N(v) \cap L_s} p_s(w)} = 1 - e^{-d_s^L(v)} \geq 1 - e^{-0.001}.$$

Suppose that \mathcal{E} occurs, and let j be the smallest index such that \mathcal{E}_j occurs, i.e. $\bar{\mathcal{E}}_1 \cap \bar{\mathcal{E}}_2 \cap \dots \cap \mathcal{E}_j$ occurs. Then, if $G_j = N(w_j) \setminus \{w_1, \dots, w_{j-1}\}$, then the probability that no neighbor of w_j beeps is at least

$$\prod_{u \in G_j} (1 - p_s(u)) \geq \prod_{u \in N(w_j)} (1 - p_s(u)) \geq e^{-2d_s(w)} \geq e^{-20}.$$

where the first inequality follows from the fact that, since w_j is light, $\mu_s(w_j) > 0$ and so each $u \in N(w_j)$ is such that $\ell_s(u) \geq 1$. Combining this with the previous inequality, we have that a vertex $w \in N(v)$ with $d_s(w) \leq 10$ is beeping with no beeping neighbors with probability at least $e^{-20}(1 - e^{-0.001}) > e^{-27} = \gamma$. \square

We now will see how Lemma 7 is implied by Lemmas 12 and 14.

Proof of Lemma 7. We fix the execution up to the end of round t , so we do not have to condition on \mathcal{F}_t . We have that

$$\begin{aligned} \Pr[\tau^{(v)}(t) > k] &= \Pr[\tau^{(v)}(t) > k \wedge G_{t,k}(v) > 0.05k] + \Pr[\tau^{(v)}(t) > k \wedge G_{t,k}(v) \leq 0.05k] \\ &\leq \Pr[\tau^{(v)}(t) > k \wedge G_{t,k}(v) > 0.05k] + e^{-k/100} \end{aligned} \tag{8}$$

$$\leq (1 - e^{-27})^{0.05k} + e^{-k/100} \leq 2e^{-30k}, \tag{9}$$

where (8) follows from Lemma 12, and (9) follows from Lemma 14. \square

6.2 Stopping Times for Platinum Rounds (Proof of Lemma 8)

Proof of Lemma 8. We prove the two items separately.

Proof of (a) If the vertex u is such that $\ell_t(u) = -\ell_{\max}(u) + \ell \leq 0$, then this implies that, since $t > \max_{w \in V} \ell_{\max}(w)$, then $\mu_t(u) > 0$ and

$$d_t(u) = \sum_{w \in N(u) \setminus S_t} 2^{-(\ell_{\max}(u) - \ell)}$$

and also that, in the subsequent ℓ rounds, the vertex u is beeping and the level of u 's neighbors increases at each step. Therefore, we have that for each $w \in N(u)$, $\ell_{t+\ell}(w) = \min\{\ell_{\max}(w), \ell_t(w) + \ell_{\max}(u)\}$ and, moreover

$$d_{t+\ell}(u) = \sum_{\substack{w \in N(u) \setminus S_t: \\ \ell_{\max}(w) > \ell_{\max}(u)}} 2^{-\ell_{\max}(u)} = \eta'_t(u).$$

We notice that, if $\ell_{\max}(u) \geq \ell_{\max}(w)$ for each $w \in N(u)$, then we have that $d_{t+\ell}(u) = 0$ and hence $\Pr[\sigma^{(u)}(t) = \sigma_{\text{in}}^{(u)}(t) \wedge \sigma^{(u)}(t) \leq \ell_{\max}(u)] = 1$, and this proves (a) when $\eta'_t(u) = 0$. If otherwise $\eta'_t(u) > 0$, we can define

$$\ell' = \max_{w \in N(u)} \ell_{\max}(w) - \ell_{\max}(u)$$

which is such that $0 < \ell' \leq \max_{w \in N(u)} \ell_{\max}(w) - \ell_{\max}(u)$, and we have that $(\sigma^{(u)}(t) = \sigma_{\text{in}}^{(u)}(t)) \wedge (\sigma^{(u)}(t) \leq \ell')$ happens with probability at least

$$\begin{aligned} \prod_{i=1}^{\ell'} \prod_{\substack{w \in N(u): \\ \ell_{\max}(w) > \ell_{\max}(u)}} \left(1 - \frac{p_t(w)}{2^{i+\ell_{\max}(u)}}\right) &\geq \prod_{i=1}^{\ell'} \prod_{\substack{w \in N(u): \\ \ell_{\max}(w) > \ell_{\max}(u)}} 3^{-p_t(w)2^{-(i+\ell_{\max}(u))}} \\ &\geq \prod_{i=1}^{\ell'} 3^{-\eta'_t(u)2^{-i}} \geq 3^{-\eta'_t(u)}, \end{aligned}$$

where the first inequality follows from the fact that, for each $w \in N(u)$, $p_t(w)/2^{\ell_{\max}(u)} \leq 2^{-4}$.

Proof of (b) We observe that, for each $x \geq 0$, we have that

$$\Pr[\sigma^{(u)}(t) \neq \sigma_{\text{in}}^{(u)}(t) \wedge \sigma^{(u)}(t) > \ell_{\max}(u) + x] \leq \Pr[\sigma^{(u)}(t) = \sigma_{\text{out}}^{(u)}(t) \mid \sigma^{(u)}(t) > \ell_{\max}(u) + x].$$

Since we have that the event $\sigma^{(u)}(t) > \ell_{\max}(u) + x$ implies that, for each $w \in N(u)$, $p_{t+\ell_{\max}(u)+x}(w) \leq 2^{-(\ell_{\max}(u)+x)}$, we have that

$$\begin{aligned} \Pr[\sigma^{(u)}(t) = \sigma_{\text{out}}^{(u)}(t) \mid \sigma^{(u)}(t) > \ell_{\max}(u) + x] &\leq 1 - \prod_{w \in N(u) \setminus S_t} \left(1 - 2^{-(\ell_{\max}(u)+x)}\right) \\ &\leq \sum_{w \in N(u) \setminus S_t} 2^{-(\ell_{\max}(u)+x)} \leq \eta'_t(u)2^{-x}. \quad \square \end{aligned}$$

7 Two Beeping Channels

One reason that makes the MIS algorithm of Jeavons et al. [18] not self-stabilizing is the usage of phases consisting of two rounds. This allows a newly joined MIS vertex to signal this event to all neighbors in the subsequent round. Afterwards, these vertices become passive, i.e., no longer participate in the algorithm. Thus, a vertex v that newly joined the MIS cannot be forced by a neighbor that is unaware that v joined the MIS to leave the MIS again in the successive round. This problem can be circumvented if a second beeping channel is available, since we can let the vertices joining the MIS beep on that channel. Indeed, beginning in the round immediately following the round vertex v joined the MIS, it signals in every round on this second channel. Neighbors of v take this as an opportunity to become non-MIS vertices. This second channel and the corresponding behavior can be easily integrated into Algorithm 1 (see Algorithm 2). The range of state variable $\ell(v)$ is restricted to $[0, \ell_{\max}(v)]$, where $\ell(v) = 0$ (resp. $\ell(v) = \ell_{\max}(v)$) implies that v is in the MIS (resp. not in the MIS). If a vertex v which is enabled to signal with $beep_1$ receives neither signal from a neighbor then it sets $\ell(v)$ to 0 and signals $beep_2$ in all future rounds. Vertices receiving a $beep_2$ signal set $\ell(v)$ to $\ell_{\max}(v)$ and refrain from beeping in future rounds. We end this section by giving the proof of Corollary 3.

Algorithm 2: Self-stabilizing version of Jeavons, Scott and Xu's Algorithm with two beeping channels [18]

state: $\ell \in \{0, \dots, \ell_{\max}(v)\}$

in each round $t = 1, 2, \dots$ **do**

- if** $0 < \ell < \ell_{\max}(v)$ **then**
 - $beep_1 \leftarrow \text{true}$ with probability $2^{-\ell}$ and $beep_1 \leftarrow \text{false}$ otherwise
- else** $beep_1 \leftarrow \text{false}$
- $beep_2 \leftarrow \ell = 0$
- if** $beep_1$ or $beep_2$ **then** send the corresponding signal to all neighbors
receive any signals sent by neighbors
- if** $beep_2$ signal received **then**
 - $\ell \leftarrow \ell_{\max}(v)$;
- else if** $beep_1$ signal received **then**
 - $\ell \leftarrow \min\{\ell + 1, \ell_{\max}(v)\}$
- else if** $beep_1$ **then**
 - $\ell \leftarrow 0$
- else if** $beep_2 = \text{false}$ **then**
 - $\ell \leftarrow \max\{\ell - 1, 1\}$

Proof of Corollary 3. We consider Algorithm 2 and we notice that the update rule of ℓ of the non-stable vertices is the same of Algorithm 1, and hence we can still use Lemma 7, since it relies just on the update rule for ℓ . Note the commonality between the two algorithms: In Algorithm 1 if the level of a vertex is 0 or lower then it is guaranteed that it sends a beep. In Algorithm 2 a vertex sends a $beep_2$ signal if and only if its level is 0.

We will prove that the termination time of Algorithm 2 is $O(\log n)$, if we take $\ell_{\max}(v) \geq$

$2 \log \deg_2(v) + 15$ for every $v \in V$. We first notice that, in this case, we have that

$$\eta_1(v) \leq \sum_{u \in N(v)} 2^{-2 \log \deg_2(u) - 15} \leq \sum_{u \in N(v)} \frac{1}{\deg^2(v)} 2^{-15} \leq 0.0001,$$

and hence, for each $t \geq 1$ and $v \in V$ we have that $\eta_t(v) \leq 0.0001$.

We notice that, for a vertex $v \in V$ to stabilize in Algorithm 2, it suffice to have a platinum round for v . Hence, from Lemma 7 we have that each vertex stabilizes in time $O(\log n)$ with probability at least $1 - 1/n^2$, and from an union bound over all the vertices follows the theorem. \square

8 Conclusion

The design of self-stabilizing MIS selection algorithms has received a lot of attention for all types of communication models. For the LOCAL model, there exists a deterministic synchronous MIS algorithm stabilizing in $O(\Delta + \log^* n)$ rounds [5]. A randomized algorithm stabilizing w.h.p. in $O(\log n)$ rounds is due to Turau [24]. Self-stabilizing algorithms for weaker communication models are also available [1, 9], both assume vertices to have information about the topology and stabilize in $O(\log^2 N \log n)$ resp. $O((D + \log n) \log n)$ rounds, where N is an upper bound on the size of the graph, and D is an upper bound on its diameter.

In this paper, we describe a new randomized self-stabilizing MIS algorithm using the beeping model requiring each vertex to have only limited knowledge about the topology that comes close to the $O(\log n)$ time bound. The algorithm is motivated by the non self-stabilizing algorithm of Jeavons et al. [18]. To transform it into a self-stabilizing algorithm we had to overcome two issues: Firstly, the known initial configuration and secondly, the phase concept. We prove that the global knowledge of the maximum degree is sufficient to obtain a $O(\log n)$ self-stabilizing algorithm. If we rely on the local knowledge of the vertex degree, the algorithm stabilizes in time $O(\log n \cdot \log \log n)$. It is an open question if this upper bound is tight, or whether the analysis can be improved to obtain the upper bound $O(\log n)$.

We remark that, for a beeping model with two channels, we can easily implement the phases with two rounds with the presence of two beeping channels, and we prove that, in such a case, a self-stabilizing MIS algorithm terminating in $O(\log n)$ relies on the (almost) local knowledge of the 2-hop neighbors. It is natural to ask whether the local knowledge can be completely removed, obtaining an algorithm for the beeping model (with one or two channels) that computes an MIS in a self-stabilizing way.

References

- [1] Yehuda Afek, Noga Alon, Ziv Bar-Joseph, Alejandro Cornejo, Bernhard Haeupler, and Fabian Kuhn. Beeping a maximal independent set. *Distributed Comput.*, 26(4):195–208, 2013. doi:10.1007/s00446-012-0175-7.
- [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. doi:10.1126/science.1193210.

- [3] Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms*, 7(4):567–583, 1986. doi:10.1016/0196-6774(86)90019-2.
- [4] Alkida Balliu, Sebastian Brandt, Juho Hirvonen, Dennis Olivetti, Mikaël Rabie, and Jukka Suomela. Lower bounds for maximal matchings and maximal independent sets. *J. ACM*, 68(5):39:1–39:30, 2021. doi:10.1145/3461458.
- [5] Leonid Barenboim, Michael Elkin, and Uri Goldenberg. Locally-Iterative distributed $(\Delta + 1)$ -coloring and applications. *J. ACM*, 69(1):5:1–5:26, 2022. doi:10.1145/3486625.
- [6] Alejandro Cornejo and Fabian Kuhn. Deploying wireless networks with beeps. In *Proc. 24th International Symposium on Distributed Computing, DISC*, pages 148–162, 2010. doi:10.1007/978-3-642-15763-9_15.
- [7] Edsger W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Commun. ACM*, 17(11):643–644, 1974. doi:10.1145/361179.361202.
- [8] Shlomi Dolev. *Self-Stabilization*. MIT Press, 2000. URL: <http://www.cs.bgu.ac.il/%7Edolev/book/book.html>.
- [9] Yuval Emek and Eyal Keren. A thin self-stabilizing asynchronous unison algorithm with applications to fault tolerant biological networks. In *Proc. ACM Symposium on Principles of Distributed Computing, PODC*, pages 93–102. ACM, 2021. doi:10.1145/3465084.3467922.
- [10] Yuval Emek and Jara Uitto. Dynamic networks of finite state machines. *Theor. Comput. Sci.*, 810:58–71, 2020. doi:10.1016/J.TCS.2017.05.025.
- [11] Yuval Emek and Roger Wattenhofer. Stone age distributed computing. In *Proc. ACM Symposium on Principles of Distributed Computing, PODC*, pages 137–146, 2013. doi:10.1145/2484239.2484244.
- [12] Salwa Faour, Mohsen Ghaffari, Christoph Grunau, Fabian Kuhn, and Václav Rozhon. Local distributed rounding: Generalized to mis, matching, set cover, and beyond. In *Proc. ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 4409–4447. SIAM, 2023. URL: <https://doi.org/10.1137/1.9781611977554.ch168>, doi:10.1137/1.9781611977554.CH168.
- [13] Mohsen Ghaffari. An improved distributed algorithm for maximal independent set. In *Proc. Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 270–277. SIAM, 2016. doi:10.1137/1.9781611974331.ch20.
- [14] Mohsen Ghaffari. Distributed MIS via all-to-all communication. In *Proc. ACM Symposium on Principles of Distributed Computing, PODC*, pages 141–149, 2017. doi:10.1145/3087801.3087830.
- [15] Mohsen Ghaffari. Distributed maximal independent set using small messages. In *Proc. Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 805–820. SIAM, 2019. doi:10.1137/1.9781611975482.50.

- [16] Mohsen Ghaffari, Christoph Grunau, and Václav Rozhon. Improved deterministic network decomposition. In *Proc. ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 2904–2923. SIAM, 2021. doi:10.1137/1.9781611976465.173.
- [17] George Giakkoupis and Isabella Ziccardi. Distributed self-stabilizing MIS with few states and weak communication. In *Proc. ACM Symposium on Principles of Distributed Computing, PODC*, pages 310–320, 2023. doi:10.1145/3583668.3594581.
- [18] Peter Jeavons, Alex Scott, and Lei Xu. Feedback from nature: Simple randomised distributed algorithms for maximal independent set selection and greedy colouring. *Distributed Comput.*, 29(5):377–393, 2016. doi:10.1007/S00446-016-0269-8.
- [19] Nathan Linial. Distributive graph algorithms-global solutions from local data. In *Proc. 28th Annual Symposium on Foundations of Computer Science, FOCS*, pages 331–335, 1987. doi:10.1109/SFCS.1987.20.
- [20] Zvi Lotker, Elan Pavlov, Boaz Patt-Shamir, and David Peleg. MST construction in $o(\log \log n)$ communication rounds. In *Proc. Fifteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures, SPAA*, pages 94–100, 2003. doi:10.1145/777412.777428.
- [21] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15(4):1036–1053, 1986. doi:10.1137/0215074.
- [22] Thomas Moscibroda and Roger Wattenhofer. Maximal independent sets in radio networks. In *Proc. Twenty-Fourth Annual ACM Symposium on Principles of Distributed Computing, PODC*, pages 148–157, 2005. doi:10.1145/1073814.1073842.
- [23] David Peleg. *Distributed computing: A locality-sensitive approach*. SIAM, 2000.
- [24] Volker Turau. Making randomized algorithms self-stabilizing. In *Proc. 26th International Colloquium on Structural Information and Communication Complexity, SIROCCO*, pages 309–324, 2019. doi:10.1007/978-3-030-24922-9_21.

A Tools

Lemma 15. *Let X_1, \dots, X_n be a sequence of i.i.d. geometric random variables with success probability p . Then, we have that*

$$\Pr \left[\sum_{i=1}^n X_i \geq k \right] = \Pr [Bin(k, p) \leq n].$$

Proof. Asking that $\sum_{i=1}^n X_i \geq k$ is like asking that, in k Bernoulli trials, we have less than n successes. \square

Theorem 16 (Chernoff's Inequality). *Let $X = \sum_{i=1}^n X_i$, where X_i with $i \in [n]$ are independently distributed in $[0, 1]$. Let $\mu = \mathbf{E}[X]$ and $\mu_- \leq \mu \leq \mu_+$. Then:*

(a) *for every $t > 0$*

$$\Pr[X > \mu_+ + t] \leq e^{-2t^2/n} \quad \text{and} \quad \Pr[X < \mu_- - t] \leq e^{-2t^2/n};$$

(b) *for $\varepsilon > 0$*

$$\Pr[X > (1 + \varepsilon)\mu_+] \leq e^{-\frac{\varepsilon^2}{2}\mu_+} \quad \text{and} \quad \Pr[X < (1 - \varepsilon)\mu_-] \leq e^{-\frac{\varepsilon^2}{2}\mu_-}.$$