



HAL
open science

Tight Bounds for MIS in Multichannel Radio Networks

Sebastian Daum, Fabian Kuhn

► **To cite this version:**

Sebastian Daum, Fabian Kuhn. Tight Bounds for MIS in Multichannel Radio Networks. DISC 2015, Toshimitsu Masuzawa; Koichi Wada, Oct 2015, Tokyo, Japan. 10.1007/978-3-662-48653-5_38 . hal-01207228

HAL Id: hal-01207228

<https://hal.science/hal-01207228>

Submitted on 30 Sep 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Tight Bounds for MIS in Multichannel Radio Networks

Sebastian Daum and Fabian Kuhn

Department of Computer Science, University of Freiburg, Germany
{sdaum, kuhn}@cs.uni-freiburg.de

Abstract. In [8] an algorithm has been presented that computes a maximal independent set (MIS) within $O(\log^2 n/\mathcal{F} + \log n \text{ polyloglog } n)$ rounds in an n -node multichannel radio network with \mathcal{F} communication channels. The paper uses a multichannel variant of the standard graph-based radio network model without collision detection and it assumes that the network graph is a polynomially bounded independence graph (BIG), a natural combinatorial generalization of well-known geographic families. The upper bound of [8] is known to be optimal up to the polyloglog n factor.

In this paper, we adapt this algorithm and its analysis to improve the result of [8] in two ways. Mainly, we get rid of the polyloglog n factor in the running time and we thus obtain an asymptotically optimal multichannel radio network MIS algorithm. In addition, our new analysis allows to generalize the class of graphs from those with polynomially bounded local independence to graphs where the local independence is bounded by an arbitrary function of the neighborhood radius.

1 Introduction

In recent years there has been an increased interest in algorithms for *shared spectrum networks* [27]. Nowadays, most modern wireless communication networks feature a multitude of communication frequencies [1, 2, 5]¹—and we can certainly expect this trend to continue.

In the light of this development, in the present paper, we settle the question of determining the *optimal* asymptotic time complexity of computing a maximal independent set (MIS) in the multichannel variant of the classic radio network model first introduced in [4, 7]. The task of constructing an MIS is one of the best studied problems in the area of large-scale wireless networks. On the one hand this is due to the fact that MIS (together with coloring problems) is one of the key problems to study the problem of symmetry breaking in large, decentralized systems. On the other hand an MIS provides a simple local clustering of the graph, which can be used as a building block for computing more enhanced organization structures in these networks such as, e.g., a communication backbone based on a connected dominating set [6, 19, 30]. This is specifically relevant in the context of wireless mobile ad hoc networks or sensor networks, in which devices cannot rely on already existing infrastructure to organize themselves—devices need to compute a meaningful structure by themselves to coordinate their interactions.

¹ For example, the IEEE 802.11 WLAN standard provides a channel spectrum of up to 200 (partially overlapping) channels and Bluetooth specifies 79 usable channels.

Related Work. In [3,22] Alon et al. and Luby presented a simple and efficient randomized parallel algorithm to compute MIS of a general graph. It is straightforward to a standard distributed message passing model and as a consequence, the algorithm soon became an archetype for many distributed MIS algorithms also in other—usually more limiting—settings. The model we assume here is an extension to the radio network model, for which an MIS algorithm with runtime $O(\log^2 n)$ has been presented in [23] for the class of unit disk graphs (UDGs). This algorithm has been proven to be asymptotically optimal [11] even for more basic version of the problem known as the wake-up problem in single-hop radio networks. While the UDG restriction is well-known and popular, a more general variant known as growth-bounded graphs or bounded independence graphs that contains UDGs has also become the focus of quite some research, e.g., [21, 25, 26]. In particular, in [26] it is shown that an MIS and many related structures can be compute in (asymptotically optimal) $O(\log^* n)$ rounds in such graphs.

Much of the early algorithmic research on multichannel radio networks has focused on networks with faults assuming a malicious adversary that can jam up to t of the \mathcal{F} available channels [10, 14–18, 20, 28, 29]. In addition, for fault-free networks, in [24] a series of lower bound proofs have been provided, which show that $\Omega(\log^2 n/\mathcal{F} + \log n)$ rounds are needed to solve any problem which requires communication. In [10] a new technique (called heralding) to deal with congestion in multichannel radio networks has been established to solve leader election in single-hop networks in time asymptotically matching the lower bound of [24]. This technique has been extended in [12] and [8] to solve the problems of computing an approximate minimum dominating set and an MIS, respectively. Our research here is based on this work and in particular on the MIS algorithm of [8].

Contributions. In radio network models, in almost all cases a restriction to the underlying graph model is being assumed. One of the most general ones are so-called α -bounded independence graphs, where $\alpha(r)$ is a function that limits the size of a *maximum* independent set in any r -neighborhood of the given graph. The MIS algorithm from [8] solves the MIS problem in time $O(\log^2 n/\mathcal{F} + \log n(\log \log n)^d)$ in such graphs for which α is bounded by polynomial of degree d . Here we get rid of the polyloglog factor and thus show how to close the gap to the lower bound from [24]. At the same time, we remove any restriction on the function α . We do so by adjusting the algorithm from [8]—and though the change in the algorithm is relatively small, it leads to a significantly more involved analysis.

2 Preliminaries

Algorithm and analysis of this paper are based on [8] and on the complete version [9] of [8]. In the paper, we try to be as self-contained as possible. However, due to lack of space, proofs are omitted and we focus on motivating and explaining the algorithms and the main ideas of the analysis. For a full, detailed version of the paper, we refer to [13].

Radio Network Model. We model the network as an n -node graph $G = (V, E)$. We assume that n or a polynomial upper bound on n is known by all nodes. Nodes start out dormant and are awakened/activated by an adversary. While nodes do not have access to a global clock, communication is assumed to happen in synchronous time slots

(rounds). The network comprises \mathcal{F} communication channels. In each round each node can choose to operate on one channel, either by listening or broadcasting. A node that broadcasts does not receive any message in that round, and its signal reaches all neighbors that operate on the same channel. A node v listening on some channel can decode an incoming message iff in the given round, exactly one of its neighbors broadcasts on the same channel. If two or more neighbors broadcast, their signals collide at v and v receives nothing, unable to detect this collision. A node can only operate on one channel in each round and therefore it does not learn anything about events on other channels.

Notation. In our algorithm all nodes move between a finite set of states: \mathbb{W} – *waiting*, \mathbb{D} – *decay*, \mathbb{A} – *active*, \mathbb{H}' – *herald candidate*, \mathbb{H} – *herald*, \mathbb{L}' – *leader candidate*, \mathbb{L} – *leader*, \mathbb{M} – *MIS node*, \mathbb{E} – *eliminated/dominated*. We overload this notation to also indicate the set of nodes being in a certain state, e.g., $\mathbb{A} := \{v \in V : v \text{ is in state } \mathbb{A}\}$. Since nodes change their states, in case of ambiguity, we write \mathbb{A}_r for the set of active nodes in round r . State changes always happen between rounds. We define $V_{\text{hf}} := \mathbb{A} \cup \mathbb{H}' \cup \mathbb{L}' \cup \mathbb{H} \cup \mathbb{L}$ as the nodes in the so-called herald filter.

We use $N(v)$ to denote the neighbors of v in G , while we use $N^k(v)$ to denote the set of nodes in distance at most k from v , including v itself. We also often write $N_{\mathbb{S}}(u)$ or $N_{\mathbb{S}}^k(u)$ to abbreviate $N(u) \cap \mathbb{S}$ or $N^k(u) \cap \mathbb{S}$ respectively, for some state \mathbb{S} . For $S \subseteq V$ we let $N(S) := \bigcup_{v \in S} N(v)$. We call a node v *alone* or *lonely*, if $N_{V_{\text{hf}} \cup \mathbb{M}}(u) = \emptyset$.

We say that an event A happens *with high probability (w.h.p.)*, *with decent probability*, or *with constant probability (w.c.p.)*, if it happens with probability at least $1 - n^{-c}$, $1 - \log^{-c} n$, or $\Omega(1)$, respectively, where c is a constant that can be chosen arbitrarily large. By $x \gg y$ we denote that $x > cy$ for a sufficiently large $c > 1$.

Bounded Independence. In addition to the communication characteristics of the network, we require the network graph to be a bounded independence graph (BIG) [21,25]. A graph G is called an α -*bounded independence graph* with *independence function* $\alpha : \mathbb{N} \rightarrow \mathbb{N}$, if for every node v , no independent set S of the subgraph of G induced by $N^d(v)$ exceeds cardinality $\alpha(d)$. Note that in particular, α does not depend on n and thus for every fixed d , $\alpha(d)$ is a constant. In [8], α is required to be a polynomial, whereas in this paper, we put no restrictions on α . It can easily be verified that one can always upper bound the largest independent set of the subgraph induced by $N^d(v)$ by $\alpha(2)^d$ and thus any independence function is always upper bounded by some exponential function. For simplicity we define a constant $\alpha := \alpha(2)$ and we assume that all nodes know the value of α .

Number of Channels. We assume that $\mathcal{F} = \omega(1)$ as otherwise single channel algorithms achieve the same asymptotic time bounds. For $\mathcal{F} = \omega(\log n)$ we only actually use $\Theta(\log n)$ channels since more channels do not lead to an additional asymptotic advantage. For ease of exposition we assume $\mathcal{F} = \Omega(\log \log n)$ and refer to [9] for an explanation of how to adapt the algorithm for the case $\mathcal{F} = o(\log \log n)$.

Maximal Independent Set. We say an algorithm computes an MIS in time T , if the following properties hold w.h.p. for each round r and node v (waking up in round r_v):

- (P1) v declares itself as either *dominating* ($\in \mathbb{M}$) or *dominated* ($\in \mathbb{E}$) before round $r_v + T$ and this decision is permanent.
- (P2) If v is *dominated* in round r , then $N(v) \cap \mathbb{M}_r \neq \emptyset$.
- (P3) If v is *dominating* in round r , then $N(v) \cap \mathbb{M}_r = \emptyset$.

3 Algorithm Description

Detailed pseudo-code can be found in [13]. Here we only present pseudo-code for the core structure of our MIS algorithm (Algorithm 1).

Algorithm 1 HeraldMIS—core structure

Input: $\sigma_{\oplus}, \sigma_{\ominus}, \Delta_{max}, \pi_{\ell}, \alpha, n, n_{\mathcal{D}} = \Theta(\mathcal{F}), n_{\mathcal{A}} = \Theta(\log \log n), n_{\mathcal{R}} = \Theta(\alpha(2)),$
 $\tau_{\mathbb{W}} = \Theta(\log n), \tau_{\mathbb{D}} = \Theta(\log n / \mathcal{F}), \tau_{lonely} = \Theta(\log^2 n / \mathcal{F} + \log n), \tau_{rbg} = \Theta(\log n)$
States: \mathbb{W} —waiting, \mathbb{D} —decay, \mathbb{M} —MIS node, \mathbb{E} —eliminated
 \mathbb{A} —active, \mathbb{L}/\mathbb{L}' —leader (candidate), \mathbb{H}/\mathbb{H}' —herald (candidate)
Channels: $\mathcal{R}_1, \dots, \mathcal{R}_{n_{\mathcal{R}}}$ —report, $\mathcal{D}_1, \dots, \mathcal{D}_{n_{\mathcal{D}}}$ —decay,
 $\mathcal{A}_1, \dots, \mathcal{A}_{n_{\mathcal{A}}}$ —herald, \mathcal{H} —handshake, \mathcal{G} —red-blue game

```

1: count  $\leftarrow$  0; state  $\leftarrow$   $\mathbb{W}$ ;  $\gamma \leftarrow \perp$ ; lonely  $\leftarrow \perp$ ;  $\gamma_{min} \leftarrow \log^{-24} n$ 
2: while state  $\neq$   $\mathbb{E}$  do
3:   count  $\leftarrow$  count + 1
4:   lonely  $\leftarrow$  lonely + 1
5:    $\gamma \leftarrow \min \{ \gamma \cdot \sigma_{\oplus}, 1/2 \}$ 
6:   uniformly at random pick  $q \in [0, 1), j \in \{1, \dots, n_{\mathcal{D}}\}$  and  $k \in \{1, \dots, n_{\mathcal{R}}\}$ 
7:   switch state do
8:     case  $\mathbb{W}$  or  $\mathbb{D}$ : run DFILTER ▷ stage 1—decay filter
9:     case  $\mathbb{A}$ : run HERALDPROTOCOL ▷ stage 2—herald filter
10:    case  $\mathbb{H}'$  or  $\mathbb{L}'$ : run HANDSHAKE
11:    case  $\mathbb{H}$  or  $\mathbb{L}$ : run REDBLUEGAME
12:    case  $\mathbb{M}$ : run DOMINATOR ▷ stage 3—MIS node
13:   if lonely =  $\tau_{lonely}$  then
14:     state  $\leftarrow$   $\mathbb{M}$ 
15: endWhile

```

Theorem 1. *Algorithm HERALDMIS solves MIS within $O(\log^2 n / \mathcal{F} + \log n)$ rounds.*

We first give a short summary of how the algorithm works including a summary of results or [9]. The algorithm is divided into three stages, the decay filter (states \mathbb{W} and \mathbb{D}), the herald filter (states \mathbb{A} , \mathbb{L}' , \mathbb{H}' , \mathbb{L} and \mathbb{H}), and decided nodes (states \mathbb{M} and \mathbb{E}). Nodes move forward within those stages—possibly omitting the herald filter—but never backwards. The decay filter is a powerful tool (which we use as a black box) that provides that over the full runtime of the algorithm the degree of the graph induced by nodes in the herald filter is bounded by $O(\log^3 n)$. In short, nodes first only listen for a while (\mathbb{W}), then they start broadcasting on one out of $\Theta(\mathcal{F})$ random channels with probability $1/n$ (\mathbb{D}), doubling this probability every $O(\log n / \mathcal{F})$ rounds. A node that broadcasts moves to the herald filter and a node that receives a message restarts in \mathbb{W} . The decay filter is the same as in [9] and for a detailed description and analysis we refer to [9]. Eliminated nodes (\mathbb{E}) know that they have a neighbor in the MIS and stop their protocol. MIS nodes (\mathbb{M}) try to inform their neighborhood (eliminating them), but they also actively disrupt protocols in the herald filter, causing them to fail. Apart from this, there is no influence between nodes being in different stages.

The focus of this paper is almost exclusively on the herald filter. It helps for understanding the complex algorithm to *only* think of the graph that is induced by nodes in the herald filter and to recall that its maximum degree is polylogarithmic in n .

The herald filter is divided into three blocks, *active state/herald protocol* (\mathbb{A}), *handshake protocol* (\mathbb{L}' and \mathbb{H}') and *red-blue protocol* (\mathbb{L} and \mathbb{H}). In the first block nodes try

to make contact with surrounding nodes. If this does indeed happen, both nodes engage in a handshake, which is only successful, if none of the two nodes neighbors any MIS node or a node in the third block. If the handshake succeeds, both nodes start a series of coin flipping games, with the purpose of ensuring that no two nodes that have become leaders (\mathbb{L}) *simultaneously*, can join the MIS. The blocks that differ from the algorithm description in [9] are the active state and the red-blue protocol, where the changes in the latter are made to compensate the changes in the more complicated active state.

There are two ways for a node v to join the MIS—either by waiting for a long time without hearing from any nearby node, or by successfully communicating with a node u during the active state, *teaming up* with it (as a leader-herald pair) and together passing through the handshake and the red-blue protocol. The farther a pair of nodes advances in these blocks, the closer its leader is to become an MIS node. We now recap and describe the behavior of a node v in the herald filter, i.e., $v \in V_{\text{hf}}$, pointing out where changes to the original algorithm occur.

Loneliness. v maintains a counter *lonely*. Whenever v hears from another node, it resets *lonely* to zero. If *lonely* ever exceeds $\tau_{\text{lonely}} = \Theta(\log^2 n / \mathcal{F} + \log n)$, then v assumes that it is *alone/lonely* in the herald filter (i.e., $N_{V_{\text{hf}} \cup \mathbb{M}}(v) = \emptyset$) and joins the MIS—w.h.p., this action is safe, i.e., should v not be alone, then the neighbors of v are far from becoming MIS nodes themselves and v has enough time to eliminate them.²

Activity. Also, v maintains an *activity* value $\gamma(v) \in [\gamma_{\text{min}}, 1/2]$, where γ_{min} , the initial value, is in $\Omega(1/\text{polylog } n)$. The activity $\gamma(v)$ solely governs the behavior of v in \mathbb{A} , however all nodes in V_{hf} maintain this value. Nodes outside V_{hf} have zero activity. By default, In each round, $\gamma(v)$ increases by a (small) constant factor $\sigma_{\oplus} > 1$ such that after $\Theta(\log \log n)$ rounds it reaches the maximum possible value $1/2$. However, whenever v receives a message from a neighboring *leader* or *herald*, then v reduces $\gamma(v)$ by a (large) constant factor $\sigma_{\ominus} \gg \sigma_{\oplus}$. This is a change to the original algorithm, where $\gamma(v)$ could only increase. The reason for this change is the following. Leaders are nodes that likely become MIS nodes, and if they do then they eliminate their neighbors anyway. For safety reasons a leader l needs to wait for $\Theta(\log n)$ rounds before it may join the MIS. During that time, if l 's neighbors keep high activity values, progress might be hindered in a $\delta' = O(\log \log n)$ neighborhood of l , which is why in [9] an $\alpha(\delta') = O(\text{polyloglog } n)$ speed loss had to be accepted and also why the algorithm of [9] only works for polynomially bounded interference graphs. By also reducing activity values, progress can be guaranteed even in a close neighborhood of a leader-herald pair and nodes can join the MIS in a much more pipelined fashion. At the same time, 'unjustified' activity value reductions only cause 'minor damage' that can be mitigated.

Herald Protocol. A node v in the active state (\mathbb{A}) participates in the *herald protocol* with probability $\gamma(v) \in [\gamma_{\text{min}}, 1/2]$, otherwise it tries to learn of nearby leaders, heralds or MIS nodes, by listening to one of *constantly* many *report channels* $\mathcal{R}_1, \dots, \mathcal{R}_{n_{\mathcal{R}}}$, $n_{\mathcal{R}} \geq 3\alpha^2$. If v participates in the herald protocol, then it chooses a channel \mathcal{A}_i from $\mathcal{A}_1, \dots, \mathcal{A}_{n_{\mathcal{A}}}$ with probability 2^{-i} . It then listens on \mathcal{A}_i with probability $\pi_{\ell} \leq 1/10$ or broadcasts its *ID* otherwise.³ If v listens, but receives nothing, nothing happens and v

² In [9] there existed some component called *loneliness support block*, operating on its own set of channels $\mathcal{S}_1, \dots, \mathcal{S}_{n_{\mathcal{S}}}$; this block and its channels have been removed.

³ We want to note that π_{ℓ} is a constant parameter that we can choose arbitrarily.

stays in \mathbb{A} . Should v receive the message of another node u on \mathcal{A}_i , then in the next round it engages with u in the *handshake protocol* as a *herald candidate* (\mathbb{H}'), in the hope of moving forward to the *red-blue protocol* together with u . Should v choose to broadcast, then it deterministically pursues the handshake as a *leader candidate* (\mathbb{L}'), hoping that some other node u has heard its message and joins in for the handshake.

Handshake and Red-Blue Protocol. For a detailed description and analysis we refer to [13] and [9]. In short, a node $h \in \mathbb{H}'$ that received a message in the herald protocol sends for two rounds on the *handshake channel* \mathcal{H} , then listens twice, and sends again for two rounds. A node $l \in \mathbb{L}'$ that was sending before acts reversely, i.e., it listens, sends and listens. Only if a node receives *all* expected messages it moves forward to the red-blue protocol, otherwise it returns to \mathbb{A} . The handshake can only possibly be completed if a pair of exactly one broadcaster and one receiver participates.

The *red-blue protocol* is a repetition of $\tau_{rbg}/8 = \Theta(\log n)$ *red-blue games* of 8 rounds each. In odd rounds, both nodes l and h of the leader-herald pair send a blocking signal on \mathcal{H} , preventing nearby handshakes to succeed. At the beginning of each game, the leader l randomly picks blue or red. If it picks red, it will send a message on channel \mathcal{G} in round 2 and it listens on \mathcal{G} in round 4, for blue it acts symmetrically. In round 6, l sends the index k' of the meeting channel for the next red-blue game on a previously decided meeting channel \mathcal{R}_k .⁴ In round 8 it listens on $\mathcal{R}_{k'}$. The herald h on the other hand sends a message in both rounds 2 and 4. It listens in round 6 to update the meeting channel and in round 8 it sends a message on $\mathcal{R}_{k'}$.

By design of the handshake and the blocking signals of odd rounds in the red-blue protocol, a leader l can neighbor a leader or herald of a different pair *only* if that other node moved to the red-blue protocol simultaneously or with a 2-round shift. If l does have such a neighbor, at some point it will not hear its herald in round 2/4, when it listens. l then aborts the red-blue protocol, notifies h in round 6 and returns to \mathbb{A} . The messages sent by l/h in round 6/8 also have the purpose of letting nearby listening active nodes reduce their activity values. An *isolated pair* on the other hand cannot be knocked out anymore⁵ and after $\tau_{rbg} = \Theta(\log n)$ rounds the pair can assume that w.h.p. there is no other conflicting pair nearby and the leader joins the MIS.

The handshake did not change and the red-blue protocol has been extended by 2 rounds—round 8 now gives heralds also the possibility to reach their neighbors.

Summary of Changes. Compared to the algorithm in [9], the following three things have changed. The *loneliness support block* of the algorithm in [9] is not executed anymore, except for maintaining the counter *lonely*. Also, the threshold τ_{lonely} has been lowered to $\Theta(\log^2 n/\mathcal{F} + \log n)$ to reflect the new runtime of the algorithm. The main change is that nodes reduce their activity γ if they hear from a nearby leader or herald. The change in the red-blue game is an addition of 2 rounds: the seventh round is just a copy of rounds 1, 3 and 5; the eighth round gives the herald of the pair a possibility to notify nearby active nodes in order to reduce their activity values—so far only leaders and MIS nodes were able to reach out to their neighbors.

Note that while the algorithm itself has barely changed, the analysis needed to be extended significantly in order to achieve the optimal runtime.

⁴ The very first meeting channel is fixed by l during the handshake.

⁵ except by an MIS node, but that already implies progress

4 Analysis

4.1 Approach

To prove that Algorithm 1 indeed solves MIS in the given time bounds, we take the following approach. In [9] it was proven that the graph, induced by nodes that passed through the decay filter, has maximum degree $\Delta_{max} = O(\log^3 n)$. A node u in the herald filter ($u \in V_{\text{hf}}$) enters the MIS either if it assumes to be alone, or if it manages to create and maintain a leader-herald bond with a neighboring node for $\tau_{rbg} = \Theta(\log n)$ rounds. Once $u \in V_{\text{hf}}$, it either enters the MIS due to assumed loneliness; or if u has a neighbor in V_{hf} , then within radius $\delta := \delta_\alpha = \Theta(\log \log n)$ soon a leader-herald pair is created that maintains its bond for τ_{rbg} rounds.⁶ So far this is the same as in [9]. There, however, a stagnation of up to τ_{rbg} rounds might follow before the next isolated leader-herald pair or MIS node gets created in $N^\delta(u)$. Considering that up to $\alpha(\delta)$ nodes in $N^\delta(u)$ can enter the MIS before u or one of its neighbors enters itself, the runtime of the herald filter is $O(\tau_{rbg}\alpha(\delta))$, or $O(\log n \text{ polyloglog } n)$ if α is polynomial.

In the present paper, by decreasing activity levels of nodes neighboring leader-herald pairs, the stagnation that can be caused by leaders on their way to join the MIS does not last for longer than $O(\log \log n)$ rounds in expectation. This allows the creation of isolated leader-herald pairs in $N^\delta(u)$ in a pipelined manner, reducing the expected runtime of the herald filter to $O(\alpha(\delta) \log \log n) = O(\alpha^\delta \log \log n)$. Unlike in [9], here we also can choose δ as an arbitrarily small value in $\Theta(\log \log n)$ without increasing the runtime by more than constant factors. Choosing $\delta < \log \log n / \log \alpha$ and a Chernoff argument bounds the runtime of the herald filter by $O(\log n)$ with high probability.

In more detail, let u be a node that enters the herald filter in round t_u . For the sake of contradiction assume that u is not decided by time $t_u + \tau_{runtime}$. If u stays lonely, it enters the MIS eventually in $\tau_{lonely} \ll \tau_{runtime}$ rounds. Note that for u to move from being non-lonely to lonely, some node in $N^2(u)$ must have entered the MIS shortly before that and eliminated all neighbors that u had in V_{hf} . This can happen at most α^2 times and thus the time u spends lonely is at most $\alpha^2 \tau_{lonely} \ll \tau_{runtime}$. Hence, assume that u is not lonely, i.e., has a neighbor u' , and that no node in $N^2(u)$ joins the MIS. We show that then most of the time both u and u' have a high activity value γ .

The following argumentation motivates this. For a node u to decrease $\gamma(u)$, it must neighbor a pair. Let us call isolated pairs (in which the leader does not neighbor another leader or herald) *good pairs* and the others *bad pairs*. Conditioning on the event of a pair being created, there is a constant probability that it is a good pair. This can be considered progress, as it guarantees one of two things: Within $O(\log n)$ rounds either the leader of the good pair itself enters the MIS or a neighbor of this pair does. In the opposite case of bad pairs being created, in expectation these remain bad pairs only for a constant number of rounds. Moreover, w.h.p., there are no more than $O(\log n)$ rounds *in total* in which bad pairs exist in $N^3(u)$ after t_u , also causing at most $O(\log n)$ rounds of u and u' having an activity value below $1/2$. Adjusting parameters we get that for some $\tau_{progress} = O(\tau_{lonely})$ and an arbitrarily small constant ε , for $(1 - \varepsilon)\tau_{progress}$ rounds in $[t_u, t_u + \tau_{progress}]$ the activity values of both u and u' are $1/2$.

⁶ “Soon” indeed means in $O(1)$ rounds in expectation, as long as $\mathcal{F} = \Omega(\log \log n)$.

Furthermore, all pairs, good and bad, inform their neighbors. By the definition of good pairs, the leaders of these form an independent set. With our choice of δ thus at most $O(\sqrt{\log n})$ good pairs exist in $N^\delta(u)$. We argue that the activity values of nodes neighboring a pair that participated in the red-blue protocol for $\Omega(\log \log n)$ rounds (which almost surely holds for good pairs), are below $\gamma_{low} := \Theta(1/\text{polylog } n)$ with some decent probability (i.e., $1 - \log^{\Omega(1)} n$). The total number of nodes in $N^\delta(u)$ becoming part of a good pair in $[t_u, t_u + \tau_{progress}]$ is $O(\sqrt{\log n})$ and hence the total amount of nodes neighboring good pairs in that time is $O(\sqrt{\log n} \Delta_{max})$. A union bound and a Chernoff bound provide that the total amount of rounds in which any node v in $N^\delta(u)$ neighboring a good pair has an $\gamma(v) > \gamma_{low}$ is less than $\varepsilon \tau_{progress}$.

Together with the previous claim we get that in $(1 - 2\varepsilon)\tau_{progress}$ rounds in $[t_u, t_u + \tau_{progress}]$ both conditions are true: $\gamma(u) = \gamma(u') = 1/2$ and *all* good pairs in $N^\delta(u)$ “silenced” their neighbors—i.e., all their neighbors have activity below γ_{low} . Let us call a round with this property *promising for u*. Without going into detail, we can show that now within distance δ there exists a node w with the property of being so-called η -fat, i.e., w ’s neighborhood is at least roughly as active as that of any of its neighbors’. Fatness implies that w.c.p. two nodes l and h in $N^1(w)$ become a *good* leader-herald pair. As said before, such a pair reduces the activity values of its neighbors rather quickly, which causes the property of η -fatness to move away from w to another node in $N^\delta(u)$ and we can repeat the argument. If a bad pair is created, then η -fatness might shortly fade, but is restored quickly, so we can almost omit this case. Again using Chernoff tail bounds, we show that at some point u itself becomes η -fat and now the creation of an MIS node in $N^2(u)$ is inevitably.

We summarize again. Once an MIS node or good pair arises in constant distance from u , we are done. In an $\Omega(\log^2 n/\mathcal{F} + \log n)$ interval, u is mostly in a promising state. W.c.p. every $O(1)$ rounds a node in $N^\delta(u)$ becomes part of a good pair or joins the MIS. In expectation, within $\Theta(\log \log n)$ rounds MIS nodes eliminate their neighbors and good pairs *silence* theirs. After any of those events happen, we measure the time until u is in a promising state again. Using Chernoff over $O(\sqrt{\log n})$ such random variables results in needing at most $O(\log n)$ time, thus, by then u must be covered.

4.2 Guarantees from the Decay Filter

Due to lack of space we refer to [9] and [13] for a more detailed description of what the decay filter accomplishes.⁷ But we informally state the two main results. For each node v the decay filter guarantees that over the runtime of $\Theta(\log^2 n/\mathcal{F} + \log n)$ rounds, (1) v or one of its neighbors enters the herald filter, but (2) no more than $\Delta_{max} = O(\log^3 n)$ nodes in $N^1(v)$ do.

From now on we only look at the graph G' induced by $V' := V_{hf} \cup \mathbb{M}$, induced by non-eliminated nodes that made it past the decay filter. *All* notations are tied to this subgraph, though we omit this in our notations, i.e., $N(u)$ means the neighborhood of u in G' . Instead, if we need to consider nodes from the states \mathbb{W} and \mathbb{D} , then we explicitly say so and show this e.g. by writing $N_G(u)$.

⁷ The underlying algorithm has been first used and analyzed in [11], in a slightly more restrictive graph model and in [8] it was shown that it also works in BIGs.

4.3 Definitions for the Herald Filter

Practically all parameters (including the above mentioned Δ_{max}) depend in one way or another on the bound on independence, i.e., on α , but in most cases those dependencies are captured in the hidden constants of those asymptotic bounds.

For our analysis of a node u that enters the herald filter, we observe a specific $\delta = \Theta(\log \log n)$ neighborhood $N^\delta(u)$ of u . We set

$$\delta := \delta_\alpha := \frac{\log \log n}{2 \log \alpha} = \Theta(\log \log n), \quad (1)$$

i.e., $\alpha^\delta = (2^{\log \alpha})^{\frac{\log \log n}{2 \log \alpha}} = \sqrt{\log n}$. The choice of δ guarantees that any independent set in a δ -neighborhood is of size at most $\sqrt{\log n}$.

Our main goal is to show quick progress in $N^\delta(u)$. Progress is clearly achieved if an MIS node arises, but due to the way a node can become an MIS node, we also consider the creation of an isolated leader-herald pair progress (more precisely, the *leader* of the pair needs to be isolated from other nodes in \mathbb{L} or \mathbb{H}), as the leader will eventually join the MIS (or be knocked out permanently by a newly created MIS node).

Definition 1. (Good Pair, Bad Pair) Consider a leader-herald pair (l, h) in round r . We say (l, h) is a good pair in round r if none of the neighbors of l (other than h) is (1) in state \mathbb{L} or (2) in state \mathbb{H} or (3) is a herald candidate in round 5 or 6 of its respective handshake. Otherwise we say that (l, h) is a bad pair.

Note that the definition of a good/bad pair is independent of possibly neighboring MIS nodes. MIS nodes existing already for 4 rounds prevent the creation of leader-herald pairs in their neighborhood completely. If on the other hand a new MIS node appears next to a leader-herald pair (which is w.h.p. only possible through the loneliness route), then we have progress in a close neighborhood. Also, note that only the leader of the pair must be 'isolated'. There are two reasons for this: (1) only leaders join the MIS (2) by protocol design the *herald* of a pair can only receive messages from MIS nodes or its own leader—not by other leaders (not even in round 6) nor other heralds. This is due to the fact that any neighboring heralds act completely synchronously and a leader neighboring a non-paired herald is ahead by precisely 2 rounds.⁸ Note also that bad pairs can become good, but not vice versa. This is because all leaders and heralds prevent the creation of further leaders/heralds in their neighborhood.

Definition 2. (Activity Mass) For a node u we define $\Gamma(u) := \sum_{v \in N^1(u)} \gamma(v)$. We call this the activity sum or activity mass of node u . Furthermore we let $\Gamma^\circ(u) := \Gamma(u) - \gamma(u) = \sum_{v \in N(u)} \gamma(v)$. In some cases we are only interested in the activity mass of active nodes and then we have $\Gamma_{\mathbb{A}}(u) := \sum_{v \in N_{\mathbb{A}}^1(u)} \gamma(v)$ and $\Gamma_{\mathbb{A}}^\circ(u)$ is defined analogously. Also

$$\gamma_{min} := \log^{-24} n = \Theta(1/\text{polylog } n), \quad (2)$$

$$\gamma_{low} := \sqrt{\gamma_{min}} = \log^{-12} n. \quad (3)$$

⁸ Cf. Lemma 8.10 in [9] and the actions of heralds/leaders in rounds 6/8 of a red-blue game.

Definition 3. (Fatness) We call a node u $\hat{\eta}$ -fat for some value $\hat{\eta} \in (0, 1)$, if it holds that $\Gamma(u) \geq \hat{\eta} \cdot \max_{v \in N(u)} \{\Gamma(v)\}$.

In simple words, in terms of activity mass, u is (at least) in the same 'league' as its neighboring nodes. Using this we choose a specific fatness parameter $\eta < 1$:

$$\eta = \eta_\alpha := \alpha^{-8} \leq \alpha^{-2 \frac{\log \Delta_{\max}}{\log \log n}} \quad (4)$$

The choice of η assures that a chain of activity sums $(\Gamma(v_i))_{i \geq 1}$ of nodes v_i on a path v_1, v_2, v_3, \dots with $\Gamma(v_i) \geq \eta^{-1} \Gamma(v_{i-1})$ and $\Gamma(v_1) \geq 2$ has length at most δ , because

$$(\eta^{-1})^\delta = (\alpha^{-8})^{-\frac{\log \log n}{2 \log \alpha}} \geq (2^{\log \alpha})^{2 \frac{\log \Delta_{\max}}{\log \log n} \frac{\log \log n}{2 \log \alpha}} = \Delta_{\max}^{\gamma(u) \leq 1/2} \max_{u \in V_{\text{hf}}} \Gamma(u). \quad (5)$$

The algorithm needs to know a few more parameters. σ_\oplus and σ_\ominus govern the changes in a node's activity level. The former is a small constant, greater than, but close to 1. In most rounds a node u increases $\gamma(u)$ by σ_\oplus . σ_\ominus is a much larger factor used for decreasing activity, large enough to undo many previous increments, but still in $O(1)$.

$$\sigma_\oplus := 2^{6/(1000\bar{m})} > 1 \quad (6)$$

$$\sigma_\ominus := \sigma_\oplus^{20\bar{m}} = 2^{12/100} > 1 \quad (7)$$

\bar{m} is a large enough constant that depends on $n_{\mathcal{R}}$, but assuming that $n_{\mathcal{R}} \geq 3\alpha^2$, $\bar{m} \geq 2^{16} n_{\mathcal{R}}$ suffices. Since $\gamma_{\min} = \log^{-24} n$, $167\bar{m} \log \log n = \Theta(\log \log n)$ consecutive increments raise a node's activity value to $1/2$. Analogously, $\Theta(\log \log n)$ decrements decrease it to its minimal value γ_{\min} .

Also two time thresholds $\tau_{\text{rbg}} = \Theta(\log n)$ and $\tau_{\text{lonely}} = \Theta(\log^2 n / \mathcal{F} + \log n)$ are needed by the algorithm. τ_{rbg} is the number of rounds a node spends in the red-blue protocol, and it is a multiple of 8. If a node $u \in V_{\text{hf}}$ does not receive a single message for τ_{lonely} consecutive rounds, while being in the herald filter, a u deduces that it is alone or all its neighbors got eliminated, and joins the MIS. In our analysis we use further time thresholds $\tau_{\text{notification}} = \Theta(\log n)$, $\tau_{\text{progress}} = \Theta(\log^2 n / \mathcal{F} + \log n)$ and $\tau_{\text{runtime}} = \Theta(\log^2 n / \mathcal{F} + \log n)$, for which the following inequality chain holds:

$$\tau_{\text{runtime}} \gg \tau_{\text{lonely}} \gg \tau_{\text{progress}} \gg \tau_{\text{rbg}} \gg \tau_{\text{notification}}$$

$\tau_{\text{notification}}$ is the maximum time needed for an MIS node to notify, w.h.p., all its neighbors. If a node u is not lonely, then, w.h.p., significant progress is achieved in less than τ_{progress} rounds; more precisely, an MIS node is created in $N^{O(1)}(u)$. W.h.p., τ_{runtime} is the maximum time a node spends in the herald filter before it gets decided.

4.4 Candidate Election—Nodes in States \mathbb{A} (and \mathbb{L}')

In this section we want to establish a few facts about how nodes can transit from state \mathbb{A} to state \mathbb{L}' or \mathbb{H}' , respectively. Note that nodes can switch between states \mathbb{A} and \mathbb{L}' without communication, but to get towards any of the three states \mathbb{H}' , \mathbb{L} and \mathbb{H} , communication is mandatory.

For some node u , constant k , round r and an index $i \in \{1, \dots, n_{\mathcal{A}}\}$ we can show that with probability $1 - \Omega(\alpha^k \pi_\ell)$ no node in $N_{\mathbb{A}}^k(u)$ receives anything on any herald

election channel \mathcal{A}_i . If we condition on certain events tied to index i , like knowing (by peeking at random bits of some nodes) that some nodes *do not* operate on \mathcal{A}_i , *do* operate on \mathcal{A}_i , *do* operate *and* broadcast or listen on \mathcal{A}_i , then the statement does still hold true for all other channels. I.e., regardless of conditioning on aforementioned events, with probability $1 - \Omega(\pi_\ell)$ no node in $N_{\mathbb{A}}^k(u)$ receives anything on any channel $\mathcal{A}_j \neq \mathcal{A}_i$. Thus, by choosing π_ℓ small, we can decrease the chances of herald creation.

For lack of space we refer to [13] for a proper lemma statement.

Under certain conditions the creation of herald candidates can be lower bounded. However, for our algorithm to work, we not only need to prove that they are created, but that this happens in *solitude*, i.e., in a close neighborhood no other herald candidates are created. The next lemma—an adaption of Lemma 8.8 from [9]—ensures this.

Lemma 1. *Let t be a round in which for a node $u \in \mathbb{A}$ the following holds:*

- *there is no herald candidate in $N^2(u)$,*
- *all nodes $v \in N^2(u)$ that neighbor a herald or leader, have $\gamma(v) \leq \gamma_{low}$,*
- *all nodes in $N^2(u)$ neighboring MIS nodes are eliminated,*
- *$\Gamma(u) \geq 1$,*

If in addition it holds that either

- (a) *$\Gamma(u) < 5\alpha$, u is $\frac{1}{5\alpha}$ -fat and $\gamma(u) = \frac{1}{2}$, or*
- (b) *$\Gamma(u) \geq 5\alpha$ and u is η -fat,*

then by the end of round $t' \in [t, t + 7]$, with probability $\Omega(\pi_\ell)$ either a node in $N^2(u)$ joins the MIS or a good pair $(l, h) \in (\mathbb{L} \cap N^1(u)) \times (\mathbb{H} \cap N^1(u))$ is created.

Let us start with an intuition of this Lemma. The basic intention is to show that if u is η -fat, then w.c.p. in constant many rounds a *good* leader-herald pair with both endpoints in $N^1(u)$ arises—for this u itself does not have to have a high γ value, i.e., u does not need to be a likely part of the pair. The lemma lists many requirements. We show that shortly after a node v moves to the herald filter, in distance $\delta = O(\log \log n)$ most of the time there exists a node u that satisfies these conditions. We also show that if an isolated pair is created in $N^1(u)$, those requirements are again satisfied $O(\text{polyloglog } n)$ rounds later (in expectation) by another node u' in this δ -neighborhood of v .

We want to point out that in the neighborhood of a fat node u with $\Gamma(u)$ at least one, w.c.p. “good things” happen (i.e., the creation of MIS nodes or good leader-herald pairs) within constant many rounds, *even if* there are herald candidates nearby or *even if* some nodes neighboring *bad* leaders/heralds have high γ values. In other words, the first two requirements could be omitted. Instead we use other results to show that from those relaxed conditions one can get to the tighter ones listed here w.c.p. in constant many rounds. We argue in Subsection 4.5, that every time an isolated pair is created, the algorithm achieves progress, as it guarantees the creation of an MIS node nearby—even if this event is delayed by $O(\log n)$ rounds.

Therefore, Lemma 1 “promises” progress in the proximity of a fat node. However, we have no such statement for areas without fat nodes. Indeed an excessive creation of bad pairs in such areas can even cause problems for our argumentation. The next result, which is a key result within the whole analysis, implies that *if* a pair is created at all, then w.c.p. this pair is good. This allows us to proof later in Lemma 4 that nodes in the herald filter are practically always very active in the candidate election process—unless they already neighbor an MIS node or a good pair.

Lemma 2. *Let r be a round in which node u is in state \mathbb{A} and $N_{\mathbb{A}}(u) \neq \emptyset$. Let B^u be the event that at the end of round r , u moves to state \mathbb{H}' due to receiving a message from some node $v \in N_{\mathbb{A}}(u)$ on some channel $\mathcal{A}_{\bar{\lambda}}$. Further, let $D^u \subseteq B^u$ be the event that B^u holds and in addition no other node $v' \in N^3(v) \setminus \{u\}$ receives any message on channel $\mathcal{A}_{\bar{\lambda}}$ in round r . It holds that*

$$\mathbf{P}(B^u) = \begin{cases} O\left(\pi_\ell \frac{\gamma(u)}{\Gamma_{\mathbb{A}}^\circ(u)}\right) & \Gamma_{\mathbb{A}}^\circ(u) > 2 \\ O(\pi_\ell \gamma(u) \Gamma_{\mathbb{A}}^\circ(u)) & \Gamma_{\mathbb{A}}^\circ(u) \leq 2 \end{cases}, \quad (8)$$

$$\mathbf{P}(D^u) = \begin{cases} \Omega\left(\pi_\ell \frac{\gamma(u)}{\Gamma_{\mathbb{A}}^\circ(u)}\right) & \Gamma_{\mathbb{A}}^\circ(u) > 2 \\ \Omega(\pi_\ell \gamma(u) \Gamma_{\mathbb{A}}^\circ(u)) & \Gamma_{\mathbb{A}}^\circ(u) \leq 2 \end{cases}. \quad (9)$$

A simple corollary is—using the remarks in the beginning of Subsection 4.4—that if u gets a message from v and there are no leaders or heralds nearby, then w.c.p. u and v form a good leader-herald pair after 6 rounds.

4.5 Handshake & Red-Blue Protocol—States \mathbb{L}' , \mathbb{H}' , \mathbb{L} , \mathbb{H}

We next very shortly recap and summarize the effects of the so-called *Handshake* and *Red-Blue Game*, but for detailed information please see [9, 13].

Foremost, the handshake cannot be passed by two nodes l and h , if l also reached another herald h' . But the handshake *also* guarantees, that if two leader-herald pairs neighbor each other during the red-blue protocol, i.e., one node from one pair neighbors a node of another pair, then both pairs conducted the handshake at the *same* time (or with an offset of 2 rounds).

The red-blue protocol grants that if a good pair (l, h) is executing the protocol, then a new MIS node arises nearby within $\tau_{rbg} = \Theta(\log n)$ rounds (usually the leader l). On the other hand, any bad pair remains bad for only $O(1)$ rounds in expectation (note that good pairs can never turn bad).

4.6 Joining the MIS—Nodes in States \mathbb{M} and \mathbb{E}

Property 1 (P). The set \mathbb{M} is an independent set at all times.

This intuitive assumption is needed for some of the upcoming statements; it is clearly true at the beginning of the algorithm, when $\mathbb{M} = \emptyset$. Lemma 7 shows that if (P) is violated, then w.h.p. a contradiction occurs. The next lemma makes sure that nodes in $N(\mathbb{M})$ soon learn of their coverage.

Lemma 3. *Assume (P) holds. Let v be a node that enters state \mathbb{M} at time t . Let w be a node in $N_G(v)$ that is awake at time $t' \geq t$ and, if $w \in \mathbb{L} \cup \mathbb{H}$, that it is at most in round $\frac{9}{10}\tau_{rbg}$ of its corresponding red-blue protocol. Then by time $t' + \tau_{notification} = t' + O(\log n)$, w.h.p., w is in state \mathbb{E} .⁹*

⁹ Note that this lemma also considers nodes w from the decay filter.

4.7 Progress and Runtime

Lemma 8.13 of [9] shows that once a good leader-herald pair (l, h) is created, its leader (or another node in $N^2(l)$) joins the MIS within $O(\log n)$ rounds. Also, within close proximity of *fat* nodes (which exist in any δ -neighborhood of nodes in V_{hf}) w.c.p. such solitary pairs are created every $O(1)$ rounds. In the algorithm of [9] it might happen that after a good pair (l, h) is created $N^\delta(u)$, the only fat node in $N^\delta(u)$ is close to (l, h) . A good pair blocks the creation of other pairs around, so progress might be stalled until l joins the MIS, eliminating its neighbors (and therefore their activity) and finally, forcing the local condition of fatness to move to a different area of the graph.

Here we changed the algorithm to take care of this potential stagnancy issue. We want the attribute of fatness to move away from a good pair long before the leader joins the MIS. More precisely, a node not neighboring good pairs should become fat within $o(\log n)$ rounds. For this we require good pairs to reduce the activity levels of their neighborhoods. However, a leader-herald pair does not know whether it forms a good pair or a bad one before the $\tau_{\text{rbg}} = \Theta(\log n)$ red-blue games are over. The idea to deal with this difficulty is the following. Good pairs manage in expectation within $O(\log \log n)$ rounds to reduce their neighborhood's activity far enough such that most of the time those nodes can be considered inactive. Bad pairs, however, last for only a constant number of rounds in expectation, and are created rarely enough¹⁰ for affected nodes to recover their lost activity quickly. In other words, the longer a node is a leader, the more likely it is indeed a good one.

Careful analysis allows to transform these observations into high probability results. In the following $\gamma(u, t)$ denotes the activity level of node u in round t . Also, let ε be a constant smaller than one—about 0.1 is sufficiently small for the proofs.

Lemma 4. *Let t be a time at which a node $u \notin N^1(\mathbb{M})$ is in the herald filter. Then, w.h.p., one of following holds:*

- (a) *Within $\tau_{\text{progress}} = O(\log^2 n / \mathcal{F} + \log n)$ rounds, $u \in N^1(\mathbb{M})$, or*
- (b) *$|\{t' \in [t + 1, t + \tau_{\text{progress}}] : \gamma(u, t') = 1/2\}| \geq (1 - \varepsilon)\tau_{\text{progress}}$.*

Next we upper bound the number of rounds in which *any* neighbors of good pairs within distance δ from u manage to exceed the activity threshold γ_{low} .

Definition 4. *For a node u and a round r let $I(u, r)$ be the event that*

- *all nodes $x \in N^\delta(u)$, which neighbor an MIS node, are in state \mathbb{E} , and*
- *all nodes $x \in N^\delta(u)$, which neighbor a good herald h or good leader l , $h, l \in (\mathbb{H} \cup \mathbb{L}) \setminus N(\mathbb{M})$, have $\gamma(x) \leq \gamma_{\text{low}} = \sqrt{\gamma_{\text{min}}} = \log^{-12} n$ and are neither bad leaders nor bad heralds.*

Lemma 5. *Assume that (P) holds. Further, let \bar{r} be a round in which node u is in the herald filter and set $J := [\bar{r} + 1, \bar{r} + \tau_{\text{progress}}]$. Then, w.h.p., one of the following holds:*

- *Within $\tau_{\text{progress}} = O(\log n)$ rounds, there is an MIS node in $N^1(u)$, or*
- *$|\{r \in J : I(u, r) \text{ holds}\}| \geq (1 - \varepsilon)\tau_{\text{progress}}$.*

¹⁰ controlled by reducing the parameter π_ℓ

Lemma 6. *Assume that (P) holds. Let t_u be a round in which a node $u \notin N^1(\mathbb{M})$ has a neighbor $u' \notin N^1(\mathbb{M})$ in the herald filter. Then, w.h.p., within $\tau_{progress} = O(\log^2 n/\mathcal{F} + \log n)$ rounds a node in $N^1(\{u, u'\})$ joins the MIS.*

Lemma 7. *W.h.p., property (P) is not violated throughout the runtime of the algorithm.*

Now we have everything at hand to prove Theorem 1.

Proof Sketch (of Theorem 1). As stated earlier, the runtime of the decay filter is within $O(\log^2 n/\mathcal{F} + \log n)$, i.e., for every node u in the decay filter, by that time one node $v \in N_G^1(u)$ enters the herald filter. We also know that over the course of $O(\log^2 n)$ rounds the maximum degree of the graph G' induced by all nodes in the herald filter is at most $O(\log^3 n)$.

Let thus u be a node that enters the herald filter. If it stays lonely for $\tau_{lonely} = \Theta(\log^2 n/\mathcal{F} + \log n)$ rounds, then u joins the MIS and we are done. Hence assume that u does hear from a neighboring node u' in the herald filter before τ_{lonely} rounds have passed. We can now apply Lemma 6 to get an MIS node v created within $\tau_{progress} = O(\log^2 n/\mathcal{F} + \log n)$ rounds. It either neighbors u , in which case within $\tau_{notification} = O(\log n)$ rounds u is decided w.h.p., or it neighbors u' , which is also then eliminated in $\tau_{notification}$ rounds. That way u can become lonely again. However, since an MIS node has been created in $N^2(u)$, this can happen at most α^2 times. Thus, at most $\tau_{runtime} = 2\alpha^2\tau_{lonely}$ rounds after u entered the herald filter, u is decided.

References

1. 802.11, I.: Wireless LAN MAC and Physical Layer Specifications (March 2012)
2. Alliance, Z.: Zigbee specification. ZigBee Document 053474r06 1 (2005)
3. Alon, N., Babai, L., Itai, A.: A Fast and Simple Randomized Parallel Algorithm for the Maximal Independent Set Problem. *Journal of Algorithms* (1986)
4. Bar-Yehuda, R., Goldreich, O., Itai, A.: On the Time-Complexity of Broadcast in Multi-Hop Radio Networks: An Exponential Gap Between Determinism and Randomization. *Journal of Computer and System Sciences* 45(1), 104–126 (1992)
5. Bluetooth Consortium: Bluetooth Specification Version 4.2 (December 2014)
6. Censor-Hillel, K., Gilbert, S., Kuhn, F., Lynch, N., Newport, C.: Structuring Unreliable Radio Networks. In: *Proc. ACM Symp. on Principles of Distr. Comp. (PODC)*. pp. 79–88 (2011)
7. Chlamtac, I., Kutten, S.: On Broadcasting in Radio Networks—Problem Analysis and Protocol Design. *IEEE Transactions on Communications* 33(12), 1240–1246 (1985)
8. Daum, S., Ghaffari, M., Gilbert, S., Kuhn, F., Newport, C.: Maximal Independent Sets in Multichannel Radio Networks. In: *Proc. ACM Symp. on Principles of Distr. Comp. (PODC)* (2013)
9. Daum, S., Ghaffari, M., Gilbert, S., Kuhn, F., Newport, C.: Maximal Independent Sets in Multichannel Radio Networks. *Tech. Rep. 275*, U. of Freiburg, Dept. of Computer Sc. (2013)
10. Daum, S., Gilbert, S., Kuhn, F., Newport, C.: Leader Election in Shared Spectrum Networks. In: *Proc. ACM Symp. on Principles of Distr. Comp. (PODC)* (2012)
11. Daum, S., Kuhn, F., Newport, C.: Efficient Symmetry Breaking in Multi-Channel Radio Networks. In: Aguilera, M.K. (ed.) *DISC, LNCS*, vol. 7611, pp. 238–252. Springer, Heidelberg (2012)

12. Daum, S., Kuhn, F., Newport, C.: Efficient Symmetry Breaking in Multi-Channel Radio Networks. Tech. Rep. 271, U. of Freiburg, Dept. of Computer Science (2012)
13. Daum, S., Kuhn, F.: Tight Bounds for MIS in Multichannel Radio Networks. CoRR abs/1508.04390 (2015), <http://arxiv.org/abs/1508.04390>
14. Dolev, S., Gilbert, S., Guerraoui, R., Kuhn, F., Newport, C.: The wireless synchronization problem. In: Proc. ACM Symp. on Principles of Distr. Comp. (PODC). pp. 190–199 (2009)
15. Dolev, S., Gilbert, S., Khabbazian, M., Newport, C.: Leveraging Channel Diversity to Gain Efficiency and Robustness for Wireless Broadcast. In: Peleg, D. (ed.) DISC, LNCS, vol. 6950, pp. 252–267. Springer, Heidelberg (2011)
16. Dolev, S., Gilbert, S., Guerraoui, R., Kowalski, D.R., Newport, C., Kuhn, F., Lynch, N.: Reliable Distributed Computations on Unreliable Radio Channels. In: Proc. MobiHoc S^3 Workshop (2009)
17. Dolev, S., Gilbert, S., Guerraoui, R., Newport, C.: Gossiping in a Multi-Channel Radio Network: An Oblivious Approach to Coping with Malicious Interference. In: Pelc, A. (ed.) DISC, LNCS, vol. 4731, pp. 208–222. Springer, Heidelberg (2007)
18. Dolev, S., Gilbert, S., Guerraoui, R., Newport, C.: Secure Communication Over Radio Channels. In: Proc. ACM Symp. on Principles of Distr. Comp. (PODC) (2008)
19. Ephremides, A., Wieselthier, J.E., Baker, D.J.: A Design Concept for Reliable Mobile Radio Networks with Frequency Hopping Signaling. Proc. of the IEEE 75(56–73) (1987)
20. Gilbert, S., Guerraoui, R., Kowalski, D., Newport, C.: Interference-Resilient Information Exchange. In: Proc. IEEE Conf. on Computer Communications (INFOCOM) (2009)
21. Kuhn, F., Moscibroda, T., Nieberg, T., Wattenhofer, R.: Fast Deterministic Distributed Maximal Independent Set Computation on Growth-Bounded Graphs. In: Fraigniaud, P. (ed.) DISC, LNCS, vol. 3724, pp. 273–287. Springer, Heidelberg (2005)
22. Luby, M.: A Simple Parallel Algorithm for the Maximal Independent Set Problem. SIAM Journal on Computing 15(4), 1036–1053 (1986)
23. Moscibroda, T., Wattenhofer, R.: Maximal independent sets in radio networks. In: Proc. ACM Symp. on Principles of Distr. Comp. (PODC). pp. 148–157 (2005)
24. Newport, C.: Radio Network Lower Bounds Made Easy. In: Kuhn, F. (ed.) DISC, LNCS, vol. 8784, pp. 258–272. Springer, Heidelberg (2014)
25. Schmid, S., Wattenhofer, R.: Algorithmic models for sensor networks. In: Proc. Workshop on Parallel and Distributed Real-Time Systems (WPDRTS). pp. 1–11 (2006)
26. Schneider, J., Wattenhofer, R.: A Log-Star Distributed Maximal Independent Set Algorithm for Growth-Bounded Graphs. In: Proc. ACM Symp. on Principles of Distr. Comp. (PODC). pp. 35–44 (2008)
27. Sherman, M., Mody, A., Martinez, R., Rodriguez, C., Reddy, R.: IEEE Standards Supporting Cognitive Radio and Networks, Dynamic Spectrum Access, and Coexistence. IEEE Communications Magazine 46(7), 72–79 (2008)
28. Strasser, M., Pöpper, C., Capkun, S.: Efficient Uncoordinated FHSS Anti-jamming Communication. In: Proc. ACM Symp. on Mobile Ad Hoc Netw. and Comp. (MOBIHOC) (2009)
29. Strasser, M., Pöpper, C., Capkun, S., Cagalj, M.: Jamming-resistant Key Establishment using Uncoordinated Frequency Hopping. In: Proc. IEEE Symp. on Sec. and Privacy (S&P) (2008)
30. Wan, P.J., Alzoubi, K.M., Frieder, O.: Distributed Construction of Connected Dominating Set in Wireless Ad Hoc Networks. In: Proc. IEEE Conf. on Computer Communications (INFOCOM) (2002)