Maximal Independent Sets in Multichannel Radio Networks

Sebastian Daum
Dept. of Comp. Science
U. of Freiburg, Germany
Freiburg, Germany
sdaum@cs.uni-freiburg.de

Mohsen Ghaffari
CSAIL
MIT, USA
Boston, Massachusetts
ghaffari@csail.mit.edu

Seth Gilbert
Dept. of Comp. Science
NUS, Singapore
Singapore, Singapore
seth.gilbert@comp.nus.edu.sg

Fabian Kuhn
Dept. of Comp. Science
U. of Freiburg, Germany
Freiburg, Germany
kuhn@cs.uni-freiburg.de

Calvin Newport
Dept. of Comp. Science
Georgetown University
Washington D.C., USA
cnewport@cs.georgetown.edu

Abstract

We present new upper bounds for fundamental problems in multichannel wireless networks. These bounds address the benefits of dynamic spectrum access, i.e., to what extent multiple communication channels can be used to improve performance. In more detail, we study a multichannel generalization of the standard graph-based wireless model without collision detection, and assume the network topology satisfies polynomially bounded independence.

Our core technical result is an algorithm that constructs a maximal independent set (MIS) in $O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds, in networks of size n with $\mathcal{F}$ channels, where the $\tilde{O}$-notation hides polynomial factors in $\log\log n$.

Moreover, we use this MIS algorithm as a subroutine to build a constant-degree connected dominating set in the same asymptotic time. Leveraging this structure, we are able to solve global broadcast and leader election within $O\big(D + \frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds, where D is the diameter of the graph, and k-message multi-message broadcast in $O\big(D + k + \frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds for unrestricted message size (with a slow down of only a log factor on the k term under the assumption of restricted message size). In all five cases above, we prove:

(i) our results hold with high probability (i.e., at least $1 - \frac{1}{n}$);
(ii) our results are within polyloglog factors of the relevant lower bounds for multichannel networks; and
(iii) our results beat the relevant lower bounds for single channel networks.

These new (near) optimal algorithms significantly expand the number of problems now known to be solvable faster in multichannel versus single channel wireless networks.

1 Introduction

Modern wireless devices rarely operate on a fixed communication channel. It is more common for them to use a wide swath of spectrum that has been subdivided into multiple independent channels (e.g., [1, 7]). This reality inspires a compelling question: When and how can we leverage the availability of multiple channels to improve the performance of wireless algorithms?

One might hope that using $\mathcal{F}$ channels you can always achieve an $\mathcal{F}$-times speed-up. For distributed algorithms, however, this goal is complicated by two factors: (a) each device can only use a single channel at a time; and (b) the size and density of the network is often unknown *a priori*. (In fact, some well-known problems, such as multihop wake-up, provably derive *no* benefit from multiple channels [12].) In this paper, we overcome these challenges to significantly increase the corpus of algorithms known to solve problems faster in multichannel versus single channel wireless networks. In more detail, we prove new randomized upper bounds for the following fundamental problems in graphs satisfying polynomial bounded independence (defined below):

(i) establishing a maximal independent set (MIS);
(ii) establishing a constant-degree connected dominating set (CDS);
(iii) broadcasting one message—or a set of messages—to every device in a network; and
(iv) electing a leader in a network.

For each of these problems, we give solutions that are within polyloglog factors of optimal in the multichannel setting, and that are faster than their corresponding lower bounds in single channel networks.

We argue that these results provide a powerful argument for wireless algorithm designers to more aggressively embrace the availability of multiple channels to gain performance.

Results. We assume a multichannel generalization of the standard graph-based wireless model [5, 9]. In each round, each node can choose a single channel to participate on from among $\mathcal{F} \geq 1$ available channels. We further assume that the graph representing our network topology satisfies polynomial bounded independence (the independence number of a radius r neighborhood is bounded by $f(r)$ for some polynomial f) [23, 27]. This assumption generalizes a variety of attempts to model the topology of wireless networks, including the widely used unit-disk graphs, quasi-unit-disk graphs, or, more generally, unit-ball graphs where the underlying metric has bounded doubling dimension [27].

The primary technical result of the paper is an algorithm that constructs an MIS in $O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds—where $\tilde{O}$ hides polynomial factors in $\log\log n$—with high probability[1]. This algorithm consists of two main pieces: a "decay filter" that reduces the number of nodes competing in each "area" to $O(\text{polylog}\, n)$, and a "herald filter" that leverages multiple channels to efficiently further reduce the nodes down to a constant

[1]We use the phrase *high probability* to indicate a probability at least $1 - \frac{1}{n^c}$, for some arbitrary constant $c \geq 1$.

number per area.

Much of the complexity resides in the herald filter, where we reduce the number of contenders to join the MIS from $O(\text{polylog}\, n)$ to $O(1)$. Part of the complexity comes from asynchrony: new arrivals and neighboring regions can force existing nodes to "restart," preventing progress toward the MIS. Another part of the complexity comes from the fact that randomized symmetry breaking works well over large populations, but less predictably as the number of participants gets small.

To put this result in context, in the single channel model, building an MIS requires $\Theta(\log^2 n)$ time [12, 19, 21, 25]. Based on the lower bound techniques developed in [11, 12, 14, 21], we show in Section 4 that in bounded independence graphs (and even in unit-disk graphs) any MIS algorithm requires at least $\Omega\big(\frac{\log^2 n}{\mathcal{F}} + \log n\big)$ rounds in a network with $\mathcal{F}$ channels. Our algorithm matches this multichannel lower bound up to polyloglog factors and beats the single channel lower bound. The lower bound also implies that even if the number of channels is arbitrarily large, solving the MIS problem still requires at least $\Omega(\log n)$ rounds, and our upper bound achieves almost the same time with just $\Theta(\log n)$ many channels.

Having developed an MIS algorithm, we use it as a subroutine to build a constant-degree CDS, with high probability, also in $O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds. The key challenge here is to efficiently—i.e., in $o(\log^2 n)$ time—identify short paths between nearby MIS nodes, even while the MIS subroutine is ongoing. We then leverage the overlay provided by our CDS algorithm to solve global broadcast and leader election (with synchronous starts) in $O\big(D + \frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds, and k-message multi-message broadcast in $O\big(D+k+\frac{\log^2 n}{\mathcal{F}}\big)+\tilde{O}(\log n)$ rounds for unrestricted message size (with a slow down of only a log factor on the k term under the assumption of restricted message size). These bounds (nearly) match the relevant $\Omega\big(D + \frac{\log^2 n}{\mathcal{F}}\big)$ bound for multichannel networks [17], and beat the relevant $\Omega(D + \log^2 n)$ lower bound for single channel networks [3].

Related Work. There has been much research on algorithms for graph-based *single channel* wireless network models, beginning with Chlamtac and Kutten [9] in the centralized setting and Bar-Yehuda et al. [5] in the distributed setting. The problem of finding an MIS in a distributed fashion has been studied extensively for a standard message passing model (i.e., *without* collisions). On general network topologies, an MIS can be built in $O\big(\min\big\{\log n, \sqrt{\log n}\log\Delta\big\}\big)$, where Δ is the largest degree of the network graph [2, 6, 24]. For bounded independence graphs, this is improved to $O(\log^* n)$ [28]. For single-channel radio networks (i.e., *with* collisions), without collision detection that satisfy the unit disk graph property, it has been shown that $O(\log^2 n)$ rounds are sufficient [25]. Using a reduction from the single-hop wake-up problem, this bound was shown tight [12, 19, 21, 25].

To our knowledge, the use of a connected dominating set (CDS) as a wireless network backbone was first described in [18]. It is well-known (and already described in [18] for the case of unit disk graphs) that a CDS can be constructed by first computing a small dominating set (in the case of bounded independence graphs, an MIS provides such a small dominating set), and then connecting the nodes of the dominating set through 2 and 3 hop paths. In a bounded independence graph, connecting all MIS nodes at distance at most 3 by a short path leads to a CDS where the graph induced by the CDS is connected and has bounded degree. The MIS algorithm of [25] combined with the CDS algorithm of [8] (which assumes an MIS as a precondition) provides a constant-degree CDS in $O(\log^2 n)$ rounds in the radio network model with synchronous starts (i.e., where all nodes start during the same round).[2]

The study of algorithms for multichannel wireless networks is more recent. Initially, much of the focus in multichannel networks was providing increased fault-tolerance: even if some of the channels were faulty, the computation would proceed. This basic model of unreliable multichannel wireless communication, often called *t-disrupted*, was introduced in [15], and has since been well-studied; e.g., [13–16, 20, 26, 29, 30].

We previously tackled the problem of leader election in single-hop networks (i.e., the diameter is 1) [11], where we solved the problem in $O\big(\frac{\log^2 n}{\mathcal{F}} + \log n\big)$ rounds. These techniques did not directly translate to multihop networks. We also have studied the problem of broadcast in multihop networks [17]. In this case, we assumed that nodes had access to collision detection, showing how to leverage this information to solve broadcast in $O\big(\big(D + \log n\big)\big(\log \mathcal{F} + \frac{\log n}{\mathcal{F}}\big)\big)$. For $\mathcal{F} = \log n$, this yields results similar to this paper, i.e., $O(D) + \tilde{O}(\log n)$. The results are hard to compare, however, as [17] assumes collision detection (which we do not), but we assume bounded independence (which [17] does not).

Finally, we have studied the problem of wake-up and approximating a minimum dominating set (MDS) in a multihop network with a topology that satisfies a clique decomposition assumption [12]. For the MDS problem, we achieved a constant-factor approximation of an MDS, in expectation, requiring $O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds. We found the technique could not easily be extended to achieve the strict independence of an MIS (with high probability) or tolerate the more general bounded independence assumption (instead of a clique decomposition assumption).

2 Preliminaries

Radio Network Model. We consider a multichannel variant of the standard graph-based radio network model [5]. The network is modeled as an n-node graph $G = (V, E)$. Each node knows n or a polynomial upper bound on n. There are $\mathcal{F}$ communication channels. Time is divided into synchronized slots, i.e., rounds. For the purpose of analysis, we imagine a global round numbering, but nodes do not have access to this global time. In each round, each node can choose one of the $\mathcal{F}$ channels to operate on; it can either transmit or listen on the channel. A node u that listens on a channel C receives a message from a neighbor v if and only if node v is transmitting on C and no other neighbor of u

[2]The MIS result of [25] does not require the synchronous start property, but the CDS piece from [8] does.

is transmitting on C. If two or more neighbors of u transmit on C, or if no neighbor of u transmits on C, then u receives silence. That is, we assume there is no collision detection available. A node that transmits does not receive anything. Notably, a node that operates on channel C in a given round learns nothing about events on channels other than C in that round.

Notation. For a subset of nodes $S \subseteq G$, we use $N_G^d(S)$ to denote the set $\{u \mid \exists v \in S,\ dist_G(u,v) \leq d\}$, where $dist_G(u,v)$ is the shortest distance between u and v in graph G. When $|S| = 1$, e.g., $S = \{v\}$, we use $N_G^d(v)$ to mean $N_G^d(\{v\})$. We use $N_G(v)$ to denote the neighbors of v, i.e., $N_G(v) = N_G^1(v) \setminus \{v\}$. When clear from the context, we omit the subscript G. In later sections, we describe algorithms in which nodes can be in various states, e.g.: $\mathbb{A}$, $\mathbb{H}'$, $\mathbb{H}$, $\mathbb{L}'$, $\mathbb{L}$, $\mathbb{M}$, $\mathbb{E}$. Where appropriate, we slightly abuse notation and use the state names to denote the set of nodes in a given state, e.g., $\mathbb{A}$ to denote the set $\{v \in V : v \text{ is in state } \mathbb{A}\}$. We sometimes study $N^d(u) \cap \mathbb{A}$ and write $N_{\mathbb{A}}^d(u)$. When referring to a local variable X of a node u, we write $X(u)$. If the round number is not clear from the context, we denote $X(u)$ in round r as $X(u,r)$.

Bounded Independence. We assume that the network graph G is a bounded independence graph as introduced and described in [23, 27]. Formally, any independent set $S \subseteq N_G^d(v)$ for any node v has size at most $\alpha(d)$, where $\alpha(d)$ is a polynomial function in d and (in particular) independent of n. Hence, any subgraph induced by a subset of a neighborhood $N_G^d(v)$ for $d = O(1)$ has only constant size independent sets.

Probability Notation. Consider an event A, a constant c, and a variable k. If $\mathbf{P}(A) \geq 1 - e^{-ck}$, then we say that A happens with very high probability with regard to k (w.v.h.p.(k)). If $\mathbf{P}(A) \geq 1 - k^{-c}$, then we say A happens with high probability with regard to k (w.h.p.(k)), and if A happens w.h.p.(n), then we simply say A happens with high probability (w.h.p.). Finally, w.c.p. abbreviates *'with constant probability'*.

Number of Channels. We assume $\omega(1)$ channels are available; otherwise there are existing algorithms that solve the problem in the same asymptotic time frame. If $\omega(\log n)$ channels are available, we restrict the usage to $\Theta(\log n)$, as there is no benefit from using more—in Section 4, we show that computing an MIS requires $\Omega\big(\frac{\log^2 n}{\mathcal{F}} + \log n\big)$ rounds. Solely for ease of exposition, we assume a minimum number of $\Omega(\log \log n)$ channels for all descriptions and proofs in this paper; this is not a requirement for the algorithm to work. We explain in Appendix B how to adapt our algorithms to work in a setting with $o(\log \log n)$ channels.

3 Problem Statement

We study randomized algorithms for the following problems, with high probability:

Maximal Independent Set. We say that an algorithm solves MIS in time T, if the following three properties hold:

(P1) Each node v that wakes up in round r declares itself as either *dominating* or *dominated* by round $r' \in [r, r+T]$ and this decision is permanent.

(P2) For each round r and node v, if v is *dominated* in round r, then v has at least one *dominating* neighbor in that round.

(P3) For each round r and node v, if v is *dominating* in round r, then v does not have any neighboring *dominating* node in that round.

Connected Dominating Set. We say that an algorithm solves (constant-degree) CDS in time T, if the following four properties hold:

(P1) Each node v that wakes up in round r declares itself as either *dominating* or *dominated* by round $r' \in [r, r+T]$ and this decision is permanent.

(P2) For each round r and node v, if v is *dominated* in round r, then v has at least one *dominating* neighbor in round r.

(P3) For each round r and node v, if v is *dominating* in round r, then v has at most $O(1)$ *dominating* neighbors in that round.

(P4) For each round r and each *connected component* C in the graph induced by nodes awake in round $r - T$, the *dominating* nodes in C form a connected subgraph within C.

Other Problems. We also consider **global broadcast**, where a node starts with a message, and **multi-message broadcast**, where k nodes start with a message; in both cases the algorithm succeeds when every node in the network has received the message(s). Finally, we consider leader election, which terminates when exactly one node has declared itself the leader (and no future nodes declare themselves the leader).

4 Lower Bound

In this section we present a lower bound for solving MIS in a radio network that satisfies the underlying model of this paper as presented in Section 2, i.e., collisions are assumed as well as no collision detection. This proves the (near) optimality of the MIS algorithm we present in Sections 5, 6 and 7.

Theorem 4.1. *Any algorithm that solves the MIS problem in a radio network with $\mathcal{F}$ channels, and that has at least a constant success probability requires at least $\Omega\big(\frac{\log^2 n}{\mathcal{F}} + \log n\big)$ rounds.*

Proof. We show that the lower bound even applies in a single-hop network, i.e., if the network graph G is the complete graph K_n. In this case, the problem of computing an MIS is equivalent to the leader election problem. We have shown in [11] that for a leader election protocol that is successful with probability $1 - \varepsilon$ it holds that with probability at least $1 - 3\varepsilon$, at least one message is transmitted successfully. We can therefore obtain a lower bound on the MIS problem by lower bounding the number of rounds that are needed until in a complete graph at least some node successfully receives a message from another node. In the single channel scenario, this problem is known as the wake-up problem (see e.g., [21]).

Throughout the proof, we assume that every node chooses an *ID* independently and uniformly at random from a large enough domain so that w.h.p. *ID*s are unique. Note that picking *ID*s from $\{1, \ldots, n^3\}$ suffices. As nodes pick *ID*s independently, they also behave independently during the protocol (as long as no message is received). Note also that since n only appears logarithmically in the lower bounds, this is equivalent to assuming that nodes are labeled with *ID*s $1, \ldots, n$ and $n^{1\frac{1}{3}}$ nodes are randomly picked to be woken up by the adversary.

We first show that it is sufficient to prove the following weaker lower bound: Solving the problem with success probability $1 - \varepsilon$ requires at least

$$\Omega\left(\frac{\log n}{\mathcal{F}} \cdot \log \frac{1}{\varepsilon} + \log \frac{1}{\varepsilon}\right) \quad \text{rounds.} \tag{1}$$

Thus assume that this weaker bound holds. As n only appears in the logarithm, the bound is also true if there are only say $\sqrt{n}$ nodes. Thus consider a network that consists of $\sqrt{n}$ (independent) cliques of size $\sqrt{n}$. Solving the MIS problem means that we have to solve leader election for each of these $\sqrt{n}$ instances and since we assume that *ID*s are picked independently for each node, the (probabilistic) executions of the $\sqrt{n}$ leader election instances are also independent. If each of the $\sqrt{n}$ instances has failure probability $\delta > 0$, at least one of the instances then fails with probability at least $1-(1-\delta)^{\sqrt{n}}$. To make this constant, we have to choose $\delta = \frac{1}{\Omega(\sqrt{n})}$ and thus the claimed lower bound for MIS follows.

Thus it now suffices to show that this weaker bound holds. For the $\Omega\big(\frac{\log n}{\mathcal{F}} \log \frac{1}{\varepsilon}\big)$ part of the lower bound in (1), we use a reduction from the single channel problem to the multichannel problem to show that at least $\Omega\big(\frac{\log^2(n)}{\mathcal{F}}\big)$ rounds are needed until at least one message is received successfully. It is known that solving the wake-up problem with success probability $1 - \varepsilon$ requires at least $\Omega\left(\log(n) \log\left(\frac{1}{\varepsilon}\right)\right)$ rounds on a single channel [12,19,21]. The lower bound of [12,19,21] only uses asynchronous wake-up in a weak sense, in particular the bound does also hold if nodes are woken up synchronously. Thus, at time 0, let an adversary wake up a subset of the nodes and no other nodes are woken up until (after $\Omega\left(\log(n) \log\left(\frac{1}{\varepsilon}\right)\right)$ rounds) the first message is successfully received. All nodes that participate in the wake-up protocol are therefore synchronized. This allows to simulate $\mathcal{F}$ channels in $\mathcal{F}$ rounds on a single channel. Formally, a multichannel protocol is turned into a single channel algorithm as follows. Every round of the protocol with $\mathcal{F}$ channels is broken into a block of $\mathcal{F}$ consecutive rounds in the single channel algorithm. Round i of the block is used to do all the communication that uses channel i in the multichannel protocol. The reduction therefore immediately gives the first term of the lower bound in (1).

To show a lower bound of $\Omega\big(\log \frac{1}{\varepsilon}\big)$, consider the leader election problem in a network consisting of only two nodes u and v, connected by an edge $\{u, v\}$. In principle, it would be possible to use similar techniques as in the lower bound of [12]. However since there are only 2 nodes, there also is a simpler argument. For rounds $i = 1, 2, \dots$, let $\mathcal{S}_i$ be the event that there is a successful message transmission between u and v. Also, let $U_i, V_i \in \{0, 1\}$ be random variables that described the listen/transmit behavior of u and v in round i—note that we do not care about their channel selections. We assume that $U_i = 1$ if u transmits in round i and that $U_i = 0$ otherwise. The random variable V_i is defined accordingly for node v. Recall that because u and v pick their *ID*s independently, as long as no node receives a message, they also behave independently and therefore the following hold. For any specific t-round-listen/transmit-pattern we have $p = \{0, 1\}^t$,

$$\mathbf{P}\left(U_{t+1} = 1 \middle| (U_1, \dots, U_t) = p, \bigcap_{i=1}^{t} \overline{\mathcal{S}}_i\right) = \mathbf{P}\left(V_{t+1} = 1 \middle| (V_1, \dots, V_t) = p, \bigcap_{i=1}^{t} \overline{\mathcal{S}}_i\right).$$

We therefore get that

$$\begin{aligned}
q := \mathbf{P}&\left(U_{t+1} = 1 \middle| \bigcap_{i=1}^{t} \overline{\mathcal{S}}_i\right) \\
=&\sum_{p \in \{0,1\}^t} \mathbf{P}\left(U_{t+1} = 1 \middle| (U_1, \dots, U_t) = p, \bigcap_{i=1}^{t} \overline{\mathcal{S}}_i\right) \cdot \mathbf{P}\left((U_1, \dots, U_t) = p \middle| \bigcap_{i=1}^{t} \overline{\mathcal{S}}_i\right) \\
=&\sum_{p \in \{0,1\}^t} \mathbf{P}\left(V_{t+1} = 1 \middle| (V_1, \dots, V_t) = p, \bigcap_{i=1}^{t} \overline{\mathcal{S}}_i\right) \cdot \mathbf{P}\left((V_1, \dots, V_t) = p \middle| \bigcap_{i=1}^{t} \overline{\mathcal{S}}_i\right) \\
= \mathbf{P}&\left(V_{t+1} = 1 \middle| \bigcap_{i=1}^{t} \overline{\mathcal{S}}_i\right).
\end{aligned} \tag{2}$$

Note that with one channel event $\mathcal{S}_i$ occurs if and only if $U_i \neq V_i$—if more than one channel exists, this is a *necessary* requirement. For $t \geq 1$, we can therefore bound

$$\mathbf{P}\left(\mathcal{S}_{t+1} \middle| \bigcap_{i=1}^{t} \overline{\mathcal{S}}_i\right) \leq 2q(1-q) \leq \frac{1}{2}.$$

After t rounds, the success probability can therefore not be better than $1 - 2^{-t}$ and thus for a success probability of $1 - \varepsilon$, we need $\Omega\big(\log \frac{1}{\varepsilon}\big)$ rounds. □

5 Overview of the MIS Algorithm

Algorithm Outline. Our algorithm consists of two main building blocks: the *decay filter* and the *herald filter*. The decay filter is used to reduce the maximum degree of the communication graph to $O(\operatorname{polylog} n)$. The herald filter assumes that the maximum degree is bounded accordingly and establishes an MIS in this setting.

The flow of the algorithm is as follows. Each node, on activation, starts in the decay filter. As time passes, some of the nodes move from the decay filter to the herald filter. Nodes exit the herald filter when either they have joined the MIS and have status *dominating*, or when they have an MIS neighbor and are thus *dominated*. In order to analyze the time complexity of our algorithm, we bound the time each node spends in each of the filters.

We note that nodes do not move backward in this flow. The dominating and dominated statuses are permanent; a node that is in the herald filter does not go back to the decay filter. However, a node u that is in the decay filter might skip the herald filter and directly become *dominated* if u receives a message from a *dominating* neighbor node v. Also, a node that has made progress in one of the filters can be forced to restart at the beginning of its current filter.

A node halts as soon as it discovers that it is dominated. On the other hand, a dominating node v cannot halt: it continues transmitting its status to its neighbors every so often, ensuring that each neighbor w that awakes at a later time becomes *dominated*.

Filter Guarantees. We now present the guarantees of both filters. We later discuss how the filters are implemented and prove the specified guarantees. The first property holds for all components of the algorithm, and acts in parallel with the other filter guarantees. It plays an important role in combining the filters.

(G1) For each node u, if u is awake in round r and it has a *dominating* neighbor v in that round, then w.h.p. node u becomes *dominated* by round $r' = r + O(\log n)$.

Implementation is straightforward: each node u that does not have its final status listens to one of a constant number of channels, w.c.p., every $O(1)$ rounds. Each node that is *dominating* periodically transmits on those channels, w.c.p., every $O(1)$ rounds. If u receives a message from a *dominating* neighbor, then u becomes *dominated*. Since each node can have at most $\alpha(1)$ *MIS* neighbors, applying Chernoff bound gives us guarantee (G1). We show later that both filters satisfy this guarantee.

The guarantees that we get from the decay filter are as follows:

(G2) W.h.p., for each node v and each round r, at most $O(\log n)$ nodes in $N_G^1(v)$ exit the decay filter in round r to enter the herald filter. Each node v that enters the herald filter has spent $\Omega(\log n)$ rounds in the decay filter, long enough to become dominated if v already had a dominating neighbor after waking up.

(G3) W.h.p., for each node v that is in the decay filter in round r, by round $r' = r + O\left(\frac{\log^2 n}{\mathcal{F}} + \log n\right)$, either v is *dominated*, in which case it has a dominating neighbor, or at least one node in $N_G^1(v)$ exits the decay filter and enters the herald filter between rounds r and r'.

The guarantees that we get from herald filter are as follows:

(G4) W.h.p., for each node v that is in the herald filter in round r, by round $r' = r + \tilde{O}(\log n)$, v is *dominating* or *dominated*. In the latter case, v has a dominating neighbor.

(G5) W.h.p., in any round r, the set of dominating nodes is an independent set.

Note that (G2) and (G4) together provide that, w.h.p., the maximum degree in the graph induced by undecided nodes in the herald filter is bounded by some $\Delta_H = O(\text{polylog}\, n)$.

We will describe the algorithm and prove guarantees (G1)–(G5) in the following two sections. Before doing so, we state our main theorem (for a detailed proof we refer to [10]).

Theorem 5.1. *W.h.p., an algorithm satisfying (G1)–(G5) solves the MIS problem in time* $O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$.

Proof. We first show that for each node v that is awake in round r, there is a round $r' = r + O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ such that, by the end of round r', either v is dominating or dominated, or at least one *'new'* node $w \in N_G^2(v)$ has become *dominating*. Here *'new'* means that w was not *dominating* in round r.

We first wait until round r' such that for all nodes $u \in N_G^2(v)$ that are dominating in round r, all neighbors w of u that are awake in round r are dominated by round r'. By guarantee (G1), it holds that $r' = r + O(\log n)$. Next (G3) implies that by round $r'' = r' + O\big(\frac{\log^2 n}{\mathcal{F}} + \log n\big)$, either v is dominating or dominated, or some node $u \in N_G^1(v)$ is in herald filter. From guarantee (G4), we then get that by round $r''' = r'' + \tilde{O}(\log n)$, u is either *dominating* or *dominated*. If u is dominating, it is a 'new' dominating node in $N_G^1(v) \subseteq N_G^2(v)$. If u is dominated, we know that u has a neighbor w that is dominating. We know that $w \in N_G^2(v)$ has not been dominating in round r, as otherwise, was either dominated before or (G2) implies that it does not make it out of decay filter.

It remains to show that this also implies that the algorithm solves the MIS problem in the required time. Property (P2) follows from guarantees (G3) and (G4). Property (P3) follows immediately from (G5). To prove property (P1), consider a node v that is awake in round r. Note that the number of different nodes in $N_G^2(v)$ that can become dominating is at most $\alpha(2)$ due to property (G5). By the above argument, as long as v is not dominating or dominated, w.h.p., we get a new dominating node in $N_G^2(v)$ once every at most $O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds. □

6 Decay Filter

The decay filter is a slightly modified version of the active state of the Active Wake-Up Algorithm in [12]. In essence, the decay filter is a backoff style protocol in which nodes broadcast with exponentially increasing probabilities. The main difference between the decay filter here and the one in [12] is that the graph model of the present paper is more general. Based on a weighted version of Turán's theorem that is proven in Appendix A, it is possible to generalize the analysis from bounded-degree clique partition assumption to general bounded independence graphs.

A pseudo-code is presented in Algorithm 1.

Algorithm description. The decay filter uses $\Theta(\mathcal{F})$ channels, divided into two sets:

(i) the *decay channels* $\mathcal{D}_1, \ldots, \mathcal{D}_{\mathrm{F}}$, where $\mathrm{F} = \Theta(\mathcal{F})$; and

(ii) the *report channels* $\mathcal{R}_1, \ldots, \mathcal{R}_{3\alpha(1)}$.

A node v in the decay filter proceeds as follows. First, v spends $\Theta(\log n)$ rounds listening to one of the report channels, chosen at random in each round. If it hears from an MIS node, it halts and becomes dominated.

Otherwise, node v proceeds through $\log n$ phases. Each phase consists of $\Theta\left(\frac{\log n}{\mathrm{F}}\right) = \Theta\left(\frac{\log n}{\mathcal{F}}\right)$ rounds, except for the last phase, which consists of $\Theta(\log n)$ rounds.

In each round of each phase, each node listens to one of the report channels with probability $\frac{1}{2}$. If node u is not listening to a report channel and it is in phase j, then u chooses uniformly at random one of the decay channels $\mathcal{D}_1, \ldots, \mathcal{D}_{\mathrm{F}}$. Then, with probability $\frac{2^j}{4n}$, u transmits on this selected channel and otherwise u listens to this selected channel. Thus, transmission probabilities are exponentially increasing over the phases, going from $\frac{1}{2n}$ to $\frac{1}{4}$.

If a node u transmits in a round, then u immediately exits the decay filter and enters the herald filter. Moreover, if a node u receives a message on some channel $\mathcal{D}_m$, then u gets knocked out and it restarts the decay filter. If u passes through all the phases without ever transmitting, then u moves to the herald filter. As a side note, notice that if the constants in the asymptotic notation of Algorithm 1 are chosen large enough, with high probability, this last case never happens, i.e., nodes will enter herald filter only through Algorithm 1.

7 Herald Filter

In this section, we present the herald filter. Recall that, to simplify explanations and ease understanding, we assume $\Omega(\log \log n)$ channels to be available.

The herald filter assumes that in the subgraph induced on the nodes in the filter the degree of each node is always bounded by $\Delta_H = O(\mathrm{polylog}\, n)$. Given this assumption, the objective of the filter is to find an MIS.

7.1 Algorithm Outline

Pseudo-code for the herald filter is given by Algorithm 2.

During the algorithm, each node is in one of 7 states: the *active state* $\mathbb{A}$, the *handshake states* $\mathbb{H}'$ and $\mathbb{L}'$, the *red-blue game* states $\mathbb{H}$ and $\mathbb{L}$, the MIS state $\mathbb{M}$ or the exclusion state $\mathbb{E}$. State $\mathbb{A}$ (*active*) indicates the initial state; state $\mathbb{M}$ indicates that the node is in the MIS (permanently); and $\mathbb{E}$ (*eliminated*) indicates nodes that know of a neighboring MIS node. States $\mathbb{L}'$ (*leader candidate*) and $\mathbb{L}$ (*leader*) are temporary states through which a

Algorithm 1 Decay Filter, run @ process v

Channels: $\mathcal{R}_1, \ldots, \mathcal{R}_{3\alpha(1)}$ – report, $\mathcal{D}_1, \ldots, \mathcal{D}_{\mathrm{F}}$ – decay

for $i := 1$ to $\Theta(\log n)$ **do** ▷ initial waiting phase
    listen to channel $\mathcal{R}_k$, k chosen uniformly at random from $\{1, \ldots, 3\alpha(1)\}$
    if $\mathbb{M} \in msg$ **then**
        $state \leftarrow \mathbb{E}$

$phase\text{-}length \leftarrow \Theta(\frac{\log n}{\mathcal{F}})$
for $j := 1$ to $\log n$ **do** ▷ main body
    if $j = \log n$ **then**
        $phase\text{-}length \leftarrow \Theta(\log n)$
    for $k := 1$ to $phase\text{-}length$ **do**
        pick uniformly at random: $l \in \{1, \ldots, 3\alpha(1)\}$, $m \in \{1, \ldots, \mathrm{F}\}$, $q \in [0, 1)$
        switch q **do**
            case $q \in [0, 2^{j-\log n-2})$
                send ID on channel $\mathcal{D}_m$
                exit decay filter and enter herald filter ▷ move to herald filter
            case $q \in [2^{j-\log n-2}, \frac{1}{2})$
                listen on channel $\mathcal{D}_m$
                if $msg \neq \emptyset$ **then**
                    restart decay filter ▷ get knocked out
            case $q \in [\frac{1}{2}, 1)$
                listen on channel $\mathcal{R}_l$
                if $\mathbb{M} \in msg$ **then**
                    $state \leftarrow \mathbb{E}$
exit decay filter and enter herald filter

node v passes to get to state $\mathbb{M}$, while states $\mathbb{H}'$ (*herald candidate*) and $\mathbb{H}$ (*herald*) are accompanying temporary states through which a node u passes to help a neighboring node v to pass through states $\mathbb{L}'$ and $\mathbb{L}$ to get to state $\mathbb{M}$.

In general, a node v can go to state $\mathbb{M}$ (i.e., join the MIS) in two ways: (1) either v does not receive any message for a long time and it joins the MIS assuming it is alone, or (2) v joins the MIS with the help of one of its neighbors u. In the latter case, in order to get to state $\mathbb{M}$, node v goes through states $\mathbb{L}'$ and $\mathbb{L}$, while u goes through states $\mathbb{H}'$ and $\mathbb{H}$ simultaneously. During these states, u helps node v to make sure that no other neighbor of v is trying to join MIS.

Until the state of a node v in herald filter is determined (i.e., until it moves to $\mathbb{M}$ or $\mathbb{E}$), it maintains a counter *lonely*(v) that measures for how long v has not heard from any neighbors; in addition, it maintains a parameter $\gamma(v)$, called the *activity level*, which is always in $\left[\frac{1}{4\Delta_H}, \frac{1}{2}\right]$ and governs the behavior of v in state $\mathbb{A}$. By definition, we assume that for nodes v in states $\mathbb{M}$ and $\mathbb{E}$ and for nodes v that are not presently in the herald filter, we have $\gamma(v) = 0$.

We divide the filter into 4 parts, depending on whether:

(i) the node is in the active state $\mathbb{A}$ (Section 7.2),
(ii) the handshake states $\mathbb{H}'$ and $\mathbb{L}'$ (Section 7.3),
(iii) the red-blue game states $\mathbb{H}$ and $\mathbb{L}$ (Section 7.4), or

Algorithm 2 Herald Filter — run at process v

States: $\mathbb{A}$ – active, $\mathbb{L}/\mathbb{L}'$ – leader (candidate), $\mathbb{H}/\mathbb{H}'$ – herald (candidate)
Channels: $\mathcal{A}_1, \ldots, \mathcal{A}_{n_\mathcal{A}}$ – herald election, $\mathcal{H}$ – handshake, $\mathcal{R}_1, \ldots, \mathcal{R}_{3\alpha(1)}$– report,
$\mathcal{G}$ – red-blue game, $\mathcal{S}_1, \ldots, \mathcal{S}_{n_\mathcal{S}}$ – loneliness support
Input: ε_γ, Δ_H, π_ℓ, α, n, $n_\mathcal{A}$, $n_\mathcal{S}$, τ_{lonely}, $\tau_{red\text{-}blue}$

$count \leftarrow 0$; $lonely \leftarrow 0$; $meet \leftarrow \perp$; $state \leftarrow \mathbb{A}$; $enforce \leftarrow$ **false**; $\gamma \leftarrow \frac{1}{4\Delta_H}$;
while $state \neq \mathbb{E}$ **do**
    $count \leftarrow count + 1$;
    $\gamma \leftarrow \min\left\{\gamma(1+\varepsilon_\gamma), \frac{1}{2}\right\}$
    $lonely \leftarrow lonely + 1$
    uniformly at random pick $q \in [0, 1)$ and $k \in \{1, \ldots, 3\alpha(1)\}$
    pick an $i \in \{1, \ldots, n_\mathcal{A}, \perp\}$ randomly with distribution $\mathbf{P}(i = \perp) = 2^{-n_\mathcal{A}}$ and $\mathbf{P}(i = j) = 2^{-j}$
    switch $state$ **do**
        case $\mathbb{A}$
            run **Active State**
        case $\mathbb{H}'$ or $\mathbb{L}'$
            run the **Handshake**
        case $\mathbb{H}$ or $\mathbb{L}$
            run the **Red-Blue Game**
        case $\mathbb{M}$
            run **MIS state**
endWhile

(iv) the MIS state $\mathbb{M}$ (Section 7.5).

7.2 Active State

Pseudo-code for the active state protocol is given by Algorithm 3.

Consider a node v that is in state $\mathbb{A}$ in round r. In the active state, we use $O(\log \log n)$ channels, divided into three sets:

(i) the *active channels* $\{\mathcal{A}_1, \ldots, \mathcal{A}_{n_\mathcal{A}}\}$,
(ii) the *lonely channels* $\{\mathcal{S}_1, \ldots, \mathcal{S}_{n_\mathcal{S}}\}$, and
(iii) the *report channels* $\{\mathcal{R}_1, \ldots, \mathcal{R}_{3\alpha(1)}\}$ (see Section 6),

where $n_\mathcal{A}, n_\mathcal{S} = O(\log \Delta_H) = O(\log \log n)$. In round r, node v does one of the following three things, with probability $\gamma(v)$ for (a), probability $0.9 - \gamma(v)$ for (b), and probability 0.1 for (c):

(a) Node v picks an active channel using an exponential distribution, choosing channel $\mathcal{A}_j$ with probability 2^{-j}. Then, with a fixed constant probability π_ℓ (chosen in the analysis), v listens to that channel, and with probability $1 - \pi_\ell$, v transmits its *ID* on that channel.
(b) Node v listens to one of the $3\alpha(1)$ report channels chosen uniformly at random.
(c) Node v runs a protocol that we call the *loneliness support block*, which we explain later in this subsection.

In (a), if v transmits on a channel $\mathcal{A}_i$ in state $\mathbb{A}$, then v goes to state $\mathbb{L}'$, attempting to become a leader. On the other hand, if v listens and receives a message from a node u,

Algorithm 3 Active State

if $i = \perp$ **then** $q = 1$
switch q **do**
  case $q \in [0, \pi_\ell \gamma)$
    listen on $\mathcal{A}_i$
    if $msg \neq \emptyset$ **then**
      $ID_{leader} \leftarrow msg.ID$; $state \leftarrow \mathbb{H}'$; $count \leftarrow 0$; $handshake \leftarrow succ$; $lonely \leftarrow 0$
  case $q \in [\pi_\ell \gamma, \gamma)$
    send (ID) on $\mathcal{A}_i$
    $state \leftarrow \mathbb{L}'$; $count \leftarrow 0$; $handshake \leftarrow succ$;
  case $q \in [\gamma, 0.9)$
    listen on $\mathcal{R}_k$
    if $msg.state = \mathbb{M}$ **then**
      $state \leftarrow \mathbb{E}$
  case $q \in [0.9, 1]$ ▷ **loneliness support block**
    pick $j \in \{1, 2, \ldots, n_{\mathcal{S}}\}$ uniformly at random
    with probability 2^{-j} **do**
      send (ID) on channel $\mathcal{S}_j$
    otherwise
      listen to channel $\mathcal{S}_j$
      if $msg \neq \emptyset$ **then**
        $lonely \leftarrow 0$
      if $lonely > \tau_{lonely} = \Theta(\log n \cdot \log\log n)$ **then**
        $state \leftarrow \mathbb{M}$

then v goes to state $\mathbb{H}'$ (while u moves to state $\mathbb{L}'$). In that case node v will try to help u to become a leader and join the MIS. In (b), if v hears an *ID* with status $\mathbb{M}$ on a report channel, then v is dominated by an MIS node and enters state $\mathbb{E}$ (*eliminated*).

Loneliness Support Block. Each node v maintains a counter *lonely*, to keep track of how long it has been in the herald filter without receiving any messages. Whenever v receives a message from a neighbor (anywhere in the herald filter), it resets the *lonely* counter. If *lonely* exceeds a threshold $\tau_{lonely} = \Theta(\log n \log\log n)$, then node v 'assumes' that it is isolated (i.e., that it does not have any neighbor in the herald filter). In this case, v joins the MIS and moves to state $\mathbb{M}$. Node v may in fact *not* be isolated, since a neighbor can show up later. However, we show in Lemma 8.15 that, w.h.p., this is in fact safe.

Every time v executes the loneliness support block, it picks a channel $\mathcal{S}_j$ uniformly at random from the lonely channels. Then v transmits on channel $\mathcal{S}_j$ with probability 2^{-j}; otherwise, it listens to channel $\mathcal{S}_j$. If v receives a message, it resets its *lonely* counter to zero.

Activity Level Adjustment. Now we explain the adjustment of $\gamma(w)$. When w enters the herald filter, $\gamma(w)$ is $\frac{1}{4\Delta_H}$. The value of $\gamma(w)$ is gradually increased by a small constant factor $(1+\varepsilon_\gamma)$ every round, until it reaches the maximal possible value of $\frac{1}{2}$ after $O(\log\log n)$ rounds. The intuitive idea behind this activity level is as follows. Because of nodes waking up asynchronously and the fact that nodes exit the decay filter and enter the herald filter in an asynchronous manner, we need to deal with an undesirable fact: the transmission

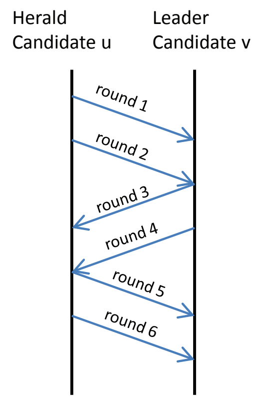

Figure 1: Handshakes between a pair of a leader candidate v and a herald candidate u

of the new nodes that enter the herald filter might affect the MIS election process which is going on among the nodes that entered the herald filter a while before that. With the gradual change in the activity level $\gamma(w)$, we can control this undesired effect and keep it below a tolerable level.

Thus, on first entering the herald filter, a node listens most of the time, but eventually, after some $O(\log\log n)$ steps, it spends a constant fraction of its time using the active channels to try to become a leader or a herald.

7.3 The Handshake

Pseudo-code for the handshake protocol is given by Algorithm 4.

Consider a node h that just moved from state $\mathbb{A}$ to state $\mathbb{H}'$ when it received a message from a node ℓ, that has also just entered state $\mathbb{L}'$. Then, h and ℓ perform a 6-round *handshake* on a designated handshake channel $\mathcal{H}$. If this handshake succeeds, then node h moves to state $\mathbb{H}$ and ℓ moves to state $\mathbb{L}$. Otherwise, both return to state $\mathbb{A}$.

The handshake proceeds as follows: In rounds 1 and 2, h transmits the *ID* of ℓ on $\mathcal{H}$, and ℓ listens. If ℓ receives both of these messages successfully (we show later that it can not receive those messages from a different node h'), then in rounds 3 and 4, ℓ transmits its *ID* on $\mathcal{H}$, and h listens. In addition, ℓ transmits a *meeting channel*, i.e., a randomly chosen report channel, which is used later in the red-blue game (see Section 7.4). Finally, assuming that these messages are received successfully by h, then in rounds 5 and 6, h again transmits the *ID* of ℓ on $\mathcal{H}$ and ℓ listens. If in any of these rounds, either of these nodes does not receive the message that it was supposed to receive, then it considers the handshake failed and returns to state $\mathbb{A}$.

Each of the 3 transmissions in the handshake is repeated twice in order to synchronize properly with the red-blue game and the nodes in the MIS. Nodes in these later states broadcast in every other round. By requiring two consecutive successful rounds of the

Algorithm 4 The Handshake

```
switch state do
    case ℍ′
        switch count do
            case 1, 2, 5, 6
                Send ID_leader on 𝓗
            case 3, 4
                Listen on 𝓗
                if msg = ∅ then
                    handshake ← fail
                else
                    meet ← msg.[2]
        if handshake = fail then
            count ← 0, state ← 𝔸
        if count = 6 then
            count ← 0, state ← ℍ; game ← succ
    case 𝕃′
        switch count do
            case 1, 2, 5, 6
                Listen on 𝓗
                if msg = ∅ then
                    handshake ← fail
            case 3, 4
                meet ← k
                Send (ID, meet) on 𝓗
        if handshake = fail then
            count ← 0, state ← 𝔸
        if count = 6 then
            count ← 0, state ← 𝕃; game ← succ
```

handshake, we can be sure that there is no concurrent red-blue game or neighboring MIS node.

Note that it is possible for ℓ to consider the handshake failed due to not receiving a message in round 5 or 6, while h assumes that the handshake was performed successfully. This situation is detected by h in the first 6 rounds of the red-blue game.

It is easy to see that one of the necessary conditions for a handshake between some $v \in \mathbb{L}'$ and $u \in \mathbb{H}'$ to be successful is that h must be the *only* herald candidate trying to perform a handshake with ℓ at that time. Hence, the nodes that enter states $\mathbb{H}$ and $\mathbb{L}$ can be viewed as *leader-herald pairs*.

7.4 The Red-Blue Game

Ideally, we would like the leaders to form an independent set (and to also be independent of nodes in $\mathbb{M}$). This would allow us to send the leaders directly to the MIS. However, this is not always the case as multiple leaders can be adjacent. The goal of the red-blue game is to detect such bad leaders (i.e., adjacent leaders) and *knock them out*, back to state $\mathbb{A}$, along with their heralds.

For this purpose, we use a simple algorithm which we call the *red-blue game*. The red-blue game uses a designated channel $\mathcal{G}$, along with the handshake channel $\mathcal{H}$ and the report channels.

Pseudo-code for the red-blue game protocol is given by Algorithm 5.

A single red-blue game is a 6-round protocol that is executed by a leader-herald pair (ℓ, h). During each game, it is possible that the pair is *knocked out*, meaning that both nodes go back to state $\mathbb{A}$. If the pair finishes $\Theta(\log n)$ red-blue games without getting knocked out, then ℓ assumes that it does not have an adjacent leader and joins the MIS.

The 6 rounds of a red-blue game are as follows: In rounds $1, 3$ and 5 of the game, both ℓ and h transmit on the handshake channel $\mathcal{H}$. These transmissions block channel $\mathcal{H}$ so that adjacent nodes cannot perform a successful handshake and thus, no new adjacent leader-herald pair can be created until either ℓ joins the MIS or the pair is knocked out.

The main rounds of the game are rounds 2 and 4. In both rounds, h broadcasts ℓ's *ID* on channel $\mathcal{G}$. At the beginning of the 6-round protocol, ℓ picks a random color in the set $\{red, blue\}$. In round 2, if ℓ chose *red*, then it transmits its *ID* on channel $\mathcal{G}$, and if it chose *blue*, it listens to $\mathcal{G}$. In round 4, the behavior is reversed: ℓ listens if it chose *red* and it transmits if it chose *blue*.

Each time ℓ is listening to channel $\mathcal{G}$, by default, it should receive the message of h. If ℓ does not receive that message, it means that another node is also transmitting on channel $\mathcal{G}$—either a leader, another herald or an MIS node. If this happens, ℓ gets knocked out.

In round 6, ℓ transmits on the meeting channel $\mathcal{R}_{meet}$, while h listens on it (if this is the first red-blue game, then $\mathcal{R}_{meet}$ was chosen randomly from the report channels during the handshake phase and transmitted from ℓ to h). The content of the sent message is whether the red-blue game succeeded (i.e., whether ℓ detected any collisions) and the meeting channel for the next red-blue game chosen uniformly at random among the report channels. If h does not receive a message from ℓ indicating that the game succeeded, then h gets knocked out. (Notice that h may not receive such a message due to a collision, in which case ℓ gets knocked out in the next red-blue game when it fails to receive a message from h.) Note that the nodes that are knocked out go back to state $\mathbb{A}$ only after they have finished the 6 rounds of their red-blue game.

The objective of the even rounds is that if two leaders are adjacent and act synchronously (round-wise), then with probability at least $\frac{1}{2}$, *both* leaders get knocked out. This is because if both leaders choose different colors *red* and *blue*, then they fail to receive the message from their respective heralds in rounds 2 and 4. Thus, if a leader-herald pair passes the red-blue game $O(\log n)$ times, then, w.h.p., there is no synchronized neighboring leader.

In the analysis, we show that because of the handshake rules, there are only very few configurations for two leader-herald pairs to be adjacent. Basically either the two leaders or the two heralds neighbor each other and operate synchronously, or if the leader of one and the herald of another pair are neighboring, then their starts of the red-blue games are shifted by exactly 2 rounds. When combined with the properties of the red-blue game, this ensures that only one leader moves on to the MIS.

7.5 The MIS State

Pseudo-code for the MIS state protocol is given by Algorithm 6.

Algorithm 5 The Red-Blue Game

```
switch state do
    case ℍ
        switch count do
            case 1, 3, 5  mod 6                          ▷ block 𝓗
                Send (state, ID_leader) on 𝓗
            case 2  mod 6                                ▷ help leader with game
                Send (state, ID_leader) on 𝓖
            case 4  mod 6                                ▷ help leader with game
                Send (state, ID_leader) on 𝓖
            case 6  mod 6
                Listen on 𝓡_meet                          ▷ from previous game
                if msg ≠ (ID_leader, succ, ∗) then
                    count ← 0, state ← 𝔸
                else
                    meet ← msg.[3]
                if count > τ_red-blue then
                    state ← 𝔼
    case 𝕃
        switch count do
            case 1, 3, 5  mod 6                          ▷ block 𝓗
                if count (mod 6) = 1 then
                    pick randomly color ∈ {red, blue}
                Send (state, ID) on 𝓗
            case 2  mod 6                                ▷ red-blue game
                if color = blue then
                    Listen on 𝓖;
                    if msg = ∅ or ID ∉ msg then
                        game ← fail
                else Send (ID) on 𝓖
            case 4  mod 6                                ▷ red-blue game
                if color = red then
                    Listen on 𝓖;
                    if msg = ∅ or ID ∉ msg then
                        game ← fail
                else Send (ID) on 𝓖
            case 6  mod 6                                ▷ Send game & new 𝓡_meet
                Send (ID_leader, game, k) on 𝓡_meet
                meet ← k
                if game = fail then
                    count ← 0, state ← 𝔸
                if count > τ_red-blue then
                    state ← 𝕄
```

Nodes in the MIS state need to continue broadcasting to prevent neighboring nodes from joining the MIS. This is accomplished by broadcasting with constant probability on $\mathcal{H}$, $\mathcal{G}$ and the report channels. More specifically, each node v that is in state $\mathbb{M}$ (i.e., that has joined the MIS) performs one of the following two steps:

(i) If v did not broadcast its *ID* on channel $\mathcal{H}$ in the previous round, then it does so in the current round.

(ii) If v did broadcast on channel $\mathcal{H}$ in the previous round, then with probability:

   a. $\frac{1}{2}$ it broadcasts its *ID* and status on channel $\mathcal{H}$,
   b. $\frac{1}{4}$ it broadcasts its *ID* and status on channel $\mathcal{G}$,
   c. $\frac{1}{4}$ it broadcasts its *ID* and status on channel $\mathcal{R}_k$, with k chosen uniformly at random in $\{1, \ldots, 3\alpha(1)\}$.

Case (a) blocks any ongoing handshakes. Case (b) knocks back neighboring leaders to state $\mathbb{A}$, preventing the red-blue game from succeeding. Case (c) knocks back neighboring

Algorithm 6 MIS state

if *enforce* **then**
  send (*state*, *ID*) on $\mathcal{H}$
  enforce $\leftarrow$ **false**
else
  switch q **do**
    case $q \in \left[0, \frac{1}{2}\right)$
      send (*state*, *ID*) on $\mathcal{H}$
      enforce $\leftarrow$ **false**
    case $q \in \left[\frac{1}{2}, \frac{3}{4}\right)$
      send (*state*, *ID*) on $\mathcal{G}$
      enforce $\leftarrow$ **true**
    case $q \in \left[\frac{3}{4}, 1\right)$
      send (*state*, *ID*) on $\mathcal{R}_k$
      enforce $\leftarrow$ **true**

heralds to state $\mathbb{A}$, and also eliminates neighboring nodes in state $\mathbb{A}$, causing them to move to $\mathbb{E}$. These ongoing broadcasts ensure that we satisfy guarantee (G1) introduced in Section 5.

Note that channel $\mathcal{H}$ is blocked at least once every two rounds. Thus, after v has been in state $\mathbb{M}$ for 6 rounds, no new neighbors of v can switch to state $\mathbb{L}$. On the other hand, note that in every period of two rounds, with constant probability, v transmits once on channel $\mathcal{G}$. The transmissions on channel $\mathcal{G}$ knock back adjacent leaders that might have been created when (or immediately after) v switched to state $\mathbb{M}$ due to the *lonely* counter. Finally, the transmissions on the report channels let the neighboring nodes of v know that they are dominated by v, causing them to halt. Note that those transmissions can also knock back neighboring heralds to state $\mathbb{A}$.

8 MIS Algorithm Analysis

We first cover in Section 8.1 the analysis of the decay filter, proving that guarantees (G2) and (G3) hold. Subsections 8.2 to 8.6 cover the complex analysis of the herald filter.

8.1 Decay Filter Analysis

The analysis is an adaptation and generalization of the analysis of the active state in the Active Wake-Up Algorithm of [12]. In particular, while the analysis in [12] requires that the network graph can be partitioned into cliques such that the graph induced by the cliques has bounded degree, we generalize the analysis to the significantly broader class of bounded independence graphs. The core argument of the analysis is based on Lemma A.2 in Appendix A and on the following Lemma 8.1 which was proven in [11].

Lemma 8.1. *Assume there are k bins and n balls with non-negative weights $w_1, \dots, w_n \leq \frac{1}{4}$, as well as a parameter $q \in (0, 1]$. Assume that $\sum_{i=1}^{n} w_i = c \cdot \frac{k}{q}$ for some constant $c \geq 1$.*

Each ball is independently selected with probability q and each selected ball is thrown into a uniformly random bin. Then, w.v.h.p.(k), there are at least $\frac{k}{4}$ bins in which the total weight of all balls is between $\frac{c}{3}$ and $2c$.

In the following, for a given round r, let $p_u(r)$ be the transmission probability of node u (i.e., $p_u(r) = \frac{2^j}{4n}$ if u is in phase j of decay filter and $p_u(r) = 0$ if u is not in one of the $\log n$ phases of the main body part) and let $P_u(r) := \sum_{v \in N^1(u)} p_v(r)$ be the sum of transmission probabilities in the neighborhood of u (in G), also called the *probability mass of* u by us. As a key property, we first show that in the neighborhood of each node u, the total probability sum $P_u(r)$ remains bounded for all rounds r, w.h.p.

Lemma 8.2. *W.h.p., for all rounds $r \geq 1$ and for all nodes $u \in V$, we have $P_u(r) = O(\mathrm{F})$.*

Proof. We need to show that w.h.p., for each node u and round r, $P_u(r)$ does not exceed $\eta\mathrm{F}$ for a sufficiently large constant $\eta > 0$. For contradiction, assume that $P_u(r) > \eta\mathrm{F}$ for some node u and round r and that round r is the first time for which $P_v > \eta\mathrm{F}$ for any node v. Let $T = \Theta\left(\frac{\log(n)}{\mathrm{F}}\right)$ be the length of one of the $\log n$ phases. As nodes only double their transmission probabilities every T rounds, and since new nodes start the decay filter with probability $\frac{1}{4n}$, for any of the T rounds r' preceding round r, we have $P_u(r') > \frac{\eta\mathrm{F}}{2} - 1 \geq \frac{\eta\mathrm{F}}{3}$ (for sufficiently large η).

Consider one such round $r' \in [r - T, r - 1]$ for which $P_u(r') > \frac{\eta\mathrm{F}}{3}$ and where by minimality of r, also $P_v(r') \leq \eta\mathrm{F}$ for all nodes v. Let $P_u^C(r')$ be the sum of the transmission probabilities of the nodes in $N^1(u)$ that choose a specific channel C among the F channels in the considered round r'. Since each node in the main-body of decay filter picks channel C with probability $\frac{1}{2\mathrm{F}}$ (with probability $1/2$ it listens to one of the report channels), we have $\frac{\eta}{6} \leq \mathbb{E}\left[P_u^C(r')\right] \leq \frac{\eta}{2}$. Assuming that η is large enough, Lemma 8.1 therefore implies that w.v.h.p.(F), on at least $\frac{\mathrm{F}}{4}$ channels C, $\frac{\eta}{18} \leq P_u^C(r') \leq \eta$.

For node u, let $P_{u,2}(r') := \bigcup_{v \in N^2(u)} p_v(r')$ be the sum of the transmission probabilities of all nodes in the 2-neighborhood of u and analogously let $P_{u,2}^C(r')$ be the probability sum of the nodes in $N^2(u)$ that pick channel C. From the bounded independence property of G, we know that the 2-neighborhood of u can be covered by the union of $O(1)$ 1-neighborhoods. Consequently, we get $P_{u,2}(r') = O(\mathrm{F})$ because $P_v(r') \leq \eta\mathrm{F}$ for all nodes v. Hence, choosing a large enough constant κ, the number of channels C for which $P_{u,2}^C(r') > \kappa\mathrm{F}$ is less than $\frac{\mathrm{F}}{8}$. Consequently, w.v.h.p.(F), there are at lest $\frac{\mathrm{F}}{8}$ channels C on which $\frac{\eta}{18} \leq P_u^C(r') \leq \eta$ and $P_{u,2}^C(r') \leq \kappa\mathrm{F}$.

Consider one such channel C and let $N^C(u, r') \subseteq N^1(u)$ be the subset of nodes in $N^1(u)$ that choose channel C in round r', i.e., $N^C(u)$ is the set of nodes contributing to $P_u^C(r')$. Further, let $S_C \subseteq N^C(u, r')$ be the nodes $v \in N_C(u, r')$ for which $P_u^C(r') = \Omega(1)$. Note that by the bounded independence property of G, the subgraph induced by $P_u^C(r')$ has bounded independence number. Therefore, applying Lemma A.2 to the graph induced by $N^C(u, r')$ implies that the nodes in S_C contribute a constant fraction of the probability mass in $P_u^C(r')$. Hence, since we chose a channel C for which $P_u^C(r') = \Theta(1)$, w.c.p. exactly one node $v \in S_C$ transmits in round r'. Also because $P^C_{u,2}(r') = O(1)$, w.c.p. no other node in $N^2(u)$ transmits on channel C in round r'. Hence, node v reaches all its neighbors on channel C and because $v \in S_C$ this knocks out a constant fraction of the probability mass contributing to $P^C_u(r')$.

We have therefore shown that on channel C and similarly on any other of the at least $\frac{\mathrm{F}}{8}$ "good" channels, w.c.p. at least a constant fraction of the probability mass contributing to $P^C_u(r')$ and thus a $\Theta\big(\frac{1}{\mathrm{F}}\big)$-fraction of the total probability mass contributing to $P_u(r')$ is eliminated (i.e., the respective nodes are knocked out). As soon as nodes have picked the channels to operate on, what happens on different channels is independent. Therefore overall on the at least $\frac{\mathrm{F}}{8}$ "good" channels, a constant fraction of the probability mass contributing to $P_u(r')$ is eliminated w.v.h.p.(F). Therefore if choosing the constant factor in the $T = \Theta\big(\frac{\log(n)}{\mathrm{F}}\big)$ rounds of a phase large enough, during the T rounds preceding round r, w.h.p., a arbitrarily large enough constant fraction of the probability mass contributing to $P_u(r)$ is eliminated. This contradicts the assumption that $P_u(r)$ exceeds $\eta\mathrm{F}$. Finally note that for each wake-up pattern of nodes, the number of rounds in which some node can be in decay filter is clearly upper bounded by a polynomial in n (each time a node is reset to the beginning of decay filter, some other node moves on to herald filter). We can therefore apply a union bound over all nodes u and rounds r to get the claim of the lemma for arbitrary u and r. □

We can now move on to proving the two decay filter guarantees (G2) and (G3).

Lemma 8.3. *(G2): With high probability, for each node v and each round r, at most $O(\log n)$ nodes in $N^1_G(v)$ come out of the decay filter in round r to enter the herald filter. Each node that enters herald filter has spent $\Omega(\log n)$ rounds in decay filter.*

Proof. The second part of the claim follows immediately from the waiting part of the decay filter. The proof of the first part is based on the previous lemma that bounds the total probability mass in each neighborhood. By Lemma 8.2, the sum of transmission probabilities in each neighborhood $N^1_G(v)$ is always at most $O(\mathcal{F}) = O(\log n)$. Therefore, the number of nodes that transmit and exit decay filter (line 14) is $O(\log n)$. As nodes decide independently whether to transmit, by a standard Chernoff argument, this bound also holds w.h.p. Nodes that exit decay filter in line 23 have a constant transmission probability and thus also as a consequence of Lemma 8.2, the number of such nodes is bounded by $O(\log n)$ in each round. □

Lemma 8.4. *(G3): W.h.p., for each node u that is in decay filter in round r, by round $r' = r + O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$, either u is* dominated*, in which case it has a dominating neighbor, or at least one node in $N^1(u)$ gets out of decay filter and enters herald filter.*

Proof. First of all, if node u gets out of decay filter via lines 4 or 22, u is dominated and thus the first part of the claim of the lemma is satisfied. Otherwise, if u never receives a message in line 18, u exits decay filter after $O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds in line 14 or in line 23, implying the second part of the claim. Finally, if u hears a message from node v in line 18, node v exits decay filter and enters herald filter and therefore also in this case the second part of the claim is satisfied. $\square$

8.2 The Analysis of the Active State

We first present some facts about the transitions of nodes from state $\mathbb{A}$ to states $\mathbb{L}'$ and $\mathbb{H}'$. For $k = O(1)$ we show that for the k-neighborhood of some node u, the probability that no node in $N^k(u)$ is being elected as a herald candidate (switching from state $\mathbb{A}$ to $\mathbb{H}'$) is constant, and, by adjusting π_ℓ, arbitrarily close to 1. We then give some conditions under which the creation of a single herald candidate happens with constant probability.

Definition 8.5. *(**Activity Sum**) 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.*

Lemma 8.6. *Fix a constant positive integer k. For any round t and node u, with probability $1 - O(\pi_\ell \alpha(k))$, no node $v \in N^k(u)$ switches from state $\mathbb{A}$ to state $\mathbb{H}'$ in round t.*

Proof. For the whole proof we only use the graph $G_\mathbb{A}$ induced by nodes in state $\mathbb{A}$ in round t. We also solely focus on nodes v that do have at least one active neighbor in $G_\mathbb{A}$, as isolated nodes cannot become herald candidates. We will use the notation $N_\mathbb{A}(v)$ and $N^d_\mathbb{A}(v)$ to refer to $N_{G_\mathbb{A}}(v)$ and $N^d_{G_\mathbb{A}}(v)$, respectively.

To become a herald candidate, a node v in state $\mathbb{A}$ must receive a message from one of its neighbors on one of the channels $\mathcal{A}_1, \ldots, \mathcal{A}_{n_\mathcal{A}}$. This is only possible if in round t, v chooses to listen on a channel $\mathcal{A}_i$ and exactly one of v's neighbors in $G_\mathbb{A}$ broadcasts on channel $\mathcal{A}_i$.

Hence, consider an arbitrary channel $\mathcal{A}_i$ from the herald election channels $\mathcal{A}_1, \ldots, \mathcal{A}_{n_\mathcal{A}}$. Let $p_v(i) = 2^{-i} \cdot \gamma(v)$ be the probability that an active node v chooses to broadcast or listen on channel $\mathcal{A}_i$. In addition, we define $P_v(i) := 2^{-i}\Gamma(v) = \sum_{w\in N^1_\mathbb{A}(v)} p_w(i)$. Let $B_i^{v,w}$ be the event that v listens on channel $\mathcal{A}_i$, while exactly one of its neighbors $w \in N_\mathbb{A}(v)$ transmits on channel $\mathcal{A}_i$ and all other neighbors $w' \in N_\mathbb{A}(v)$ are either not on channel $\mathcal{A}_i$ or they choose to listen as well.

$$\begin{aligned}
\mathbf{P}(B_i^{v,w}) &= \pi_\ell p_v(i) \cdot (1-\pi_\ell) p_w(i) \cdot \prod_{w'\in N_\mathbb{A}(v)\setminus\{w\}} \big(1 - p_{w'}(i)(1-\pi_\ell)\big) \\
&\le \pi_\ell p_v(i) p_w(i) \cdot \prod_{w'\in\{v,w\}} \frac{1}{1-(1-\pi_\ell)p_{w'}(i)} \cdot \prod_{w'\in N^1_\mathbb{A}(v)} \big(1 - p_{w'}(i)(1-\pi_\ell)\big) \\
&\le \pi_\ell p_v(i) p_w(i) \cdot 4 \cdot e^{-\frac{1}{2}P_v(i)}.
\end{aligned}$$

In the last inequality, we use that $p_{w'}(i) \leq \frac{1}{2}$ and that $\pi_\ell \leq \frac{1}{2}$. Define B_i^v to be the event that v listens on $\mathcal{A}_i$ and exactly one of its neighbors transmits on that channel. Since $B_i^v = \bigcup_{w \in N_{\mathbb{A}}(v)} B_i^{v,w}$ and the events $B_i^{v,w}$ are disjoint for different w, we have

$$\mathbf{P}(B_i^v) = \sum_{w \in N_{\mathbb{A}}(v)} \mathbf{P}(B_i^{v,w}) \leq \pi_\ell p_v(i) P_v(i) \cdot 4e^{-\frac{1}{2}P_v(i)} =: C_i^v.$$

For any $x > 0$ and constant c, $cx^2e^{-x} = O(1)$, which by using $x = P_v(i)$ implies that $C_i^v = O\big(\pi_\ell \frac{p_v(i)}{P_v(i)}\big) = O\big(\pi_\ell \frac{\gamma(v)}{\Gamma(v)}\big)$ for any fixed i. Next we show that $\sum_{i=1}^{n_\mathcal{A}} C_i^v = O\big(\pi_\ell \frac{\gamma(v)}{\Gamma(v)}\big)$, too.

$$\frac{C_{i+1}^v}{C_i^v} = \frac{p_v(i+1)}{p_v(i)} \frac{P_v(i+1)}{P_v(i)} e^{-\frac{1}{2}P_v(i+1) + \frac{1}{2}P_v(i)} = \frac{1}{4} e^{\frac{1}{4}\Gamma(v)2^{-i}} < \frac{1}{2} \qquad \forall i \geq \log \Gamma(v) \tag{3}$$

$$\frac{C_i^v}{C_{i+1}^v} = \frac{p_v(i)}{p_v(i+1)} \frac{P_v(i)}{P_v(i+1)} e^{-\frac{1}{2}P_v(i) + \frac{1}{2}P_v(i+1)} = 4e^{-\frac{1}{4}\Gamma(v)2^{-i}} < \frac{1}{2} \qquad \forall i \leq \log \Gamma(v) - 4 \tag{4}$$

We can therefore deduce the upper bounds

$$\sum_{i \geq \log \Gamma(v)} C_i^v \leq 2C_{\lceil \log \Gamma(v) \rceil}^v \text{ and } \sum_{i \leq \log \Gamma(v)} C_i^v \leq 2C_{\max\{1, \lfloor \log \Gamma(v) - 4 \rfloor\}}^v,$$

proving the claim that $\sum_{i=1}^{n_\mathcal{A}} C_i^v = O\big(\pi_\ell \frac{\gamma(v)}{\Gamma(v)}\big)$.

Using Lemma A.1, choosing $G' := G_\mathbb{A}[N_\mathbb{A}^k(u)]$, $w(v) := \gamma(v)$ and $W(v) := \Gamma(v)$, we get that $\sum_{v \in G'} \frac{\gamma(v)}{\Gamma(v)} \leq \alpha(G') \leq \alpha(k)$. (Note that the independence number of a graph is larger than or equal to the independence number of any induced subgraph.)

Let B^v be the event that v moves from state $\mathbb{A}$ to $\mathbb{H}'$ and $B = \bigcup_{v \in N^k(u)} = \bigcup_{v \in N_\mathbb{A}^k(u)}$. Then,

$$\mathbf{P}(B) \leq \sum_{v \in N_\mathbb{A}^k(u)} \mathbf{P}(B^v) \leq \sum_{v \in N_\mathbb{A}^k(u)} \sum_{i=1}^{n_\mathcal{A}} C_i^v = \sum_{v \in N_\mathbb{A}^k(u)} O\left(\pi_\ell \frac{\gamma(v)}{\Gamma(v)}\right) = O\left(\pi_\ell \alpha(k)\right).$$

Choosing a sufficiently small π_ℓ concludes the proof. □

Definition 8.7. *(**Fatness**) We call a node u (or respectively its neighborhood $N(u)$) η-*fat *for some value $\eta > 0$, if it holds that $\Gamma(u) \geq \eta \cdot \max_{v \in N(u)}\{\Gamma(v)\}$.*

Lemma 8.8. *Let t be a round in which for a node u in state $\mathbb{A}$ in the herald filter it holds that there is no herald, leader, or herald candidate in $N^2(u)$. Furthermore, all neighbors of MIS nodes in $N^2(u)$ are in state $\mathbb{E}$, $\Gamma(u) \geq 1$, and either*

(a) *$\Gamma(u) < 3\alpha(1)$, u is $\frac{1}{3\alpha(1)}$-fat, and $\gamma(u) = \frac{1}{2}$, or*

(b) *u is $\frac{1}{2}$-fat and $\Gamma(u) \geq 3\alpha(1)$.*

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

Proof. We make use of the notations $X(u,t)$ and $N^d_{\mathbb{A}}$ as explained in Section 2. For the remainder of the proof, assume that in rounds $t, \ldots, t+7$, no node in $N^2(u)$ joins the MIS, as otherwise, the claim of the lemma is trivially satisfied. In order to prove the lemma, we first show that either in round t or in round $t+1$, w.c.p., a herald candidate is created in $N^1(u)$. Formally, we define the event H_u as follows. In round t', event H_u occurs iff there are two neighboring nodes $v, w \in N^1_{\mathbb{A}}(u,t')$ such that

- v and w both operate on a channel $\lambda \in \{\mathcal{A}_1, \ldots, \mathcal{A}_{n_\mathcal{A}}\}$
- no other neighbor of v and w chooses channel λ, and
- no other node in $N^2_{\mathbb{A}}(u,t')$ receives a message on channel λ.

Clearly, if event H_u holds either in round t or $t+1$, the nodes v, w have a probability of $2\pi_\ell(1-\pi_\ell)$ of becoming a herald-leader candidate pair and no other herald candidate is created on channel λ in that round. Combined with appropriate applications of Lemma 8.6, this suffices to prove the claim of the lemma.

In the following, for a node $v \in N^1_{\mathbb{A}}(u,t')$, let $\Gamma_{\mathbb{A}}(v,t') := \sum_{w \in N^1(v,t') \cap N^1_{\mathbb{A}}(u,t')} \gamma(w,t')$ be the total activity value of all active nodes in round t' in the 1-neighborhood of v restricted to the 1-neighborhood of u. To estimate the probability that H_u occurs in a round $t' \in \{t, t+1\}$, we first show that in one of the two rounds $t' \in \{t, t+1\}$, with probability at least $\frac{1}{4}$ it holds that u is in state $\mathbb{A}$ and $\Gamma_{\mathbb{A}}(u,t') := \sum_{v \in N^1_{\mathbb{A}}(u,t')} \gamma(v,t') \geq \frac{2}{3} \cdot \Gamma(u,t)$. Assume that the claim is not true for $t' = t$. As the lemma statement is based on the assumption that u is in state $\mathbb{A}$ in round t, this implies that $\Gamma_{\mathbb{A}}(u,t) < \frac{2}{3}(\Gamma(u,t) - \gamma(u,t))$. Also by the assumptions of the lemma, in round t, no nodes in $N(u)$ are in states $\mathbb{H}'$, $\mathbb{H}$, or $\mathbb{L}$. As nodes w in states $\mathbb{M}$ and $\mathbb{E}$ have $\gamma(w) = 0$ and thus do not contribute to $\Gamma(u)$, we therefore have $\Gamma_{\mathbb{L}'}(u,t) := \sum_{v \in N^1_{\mathbb{L}'}(u,t)} \gamma(v,t) = \Gamma(u,t) - \Gamma_{\mathbb{A}}(u,t)$. Because by assumption, there are no nodes in state $\mathbb{H}'$ in round t, all nodes that are in state $\mathbb{L}'$ in round t switch back to state $\mathbb{A}$ for the next round. As by assumption, no nodes switch to states $\mathbb{M}$ or $\mathbb{E}$, and node v that is in state $\mathbb{A}$ in round t can only move out of $\mathbb{A}$ if it decides to operate on one of the channels $\mathcal{A}_1, \ldots, \mathcal{A}_{n_\mathcal{A}}$. This happens with probability at most $\gamma(v,t) \leq \frac{1}{2}$. Therefore, with probability at least $\frac{1}{2}$, at least half of the total activity value of the nodes in $N^1_{\mathbb{A}}(u,t)$ remains in state $\mathbb{A}$ for round $t+1$. And (independently) with probability at least $\frac{1}{2}$, also node u remains in state $\mathbb{A}$ for round $t+1$. Thus, with probability at least $\frac{1}{4}$, u is in state $\mathbb{A}$ in round $t+1$ and at least half of the total activity contributing to $\Gamma_{\mathbb{A}}(u,t)$ also contributes to $\Gamma_{\mathbb{A}}(u,t+1)$. Therefore, with probability at least $\frac{1}{4}$,

$$\Gamma_{\mathbb{A}}(u,t+1) \geq \big(\Gamma(u,t) - \Gamma_{\mathbb{A}}(u,t)\big) + \frac{1}{2} \cdot \Gamma_{\mathbb{A}}(u,t) = \Gamma(u,t) - \frac{1}{2} \cdot \Gamma_{\mathbb{A}}(u,t) \geq \frac{2}{3}\Gamma_{\mathbb{A}}(u,t).$$

The last inequality follows because we assumed that $\Gamma_{\mathbb{A}}(u,t) < \frac{2}{3}\Gamma(u,t)$. We also use the assumption that no nodes switch to states $\mathbb{M}$ or $\mathbb{E}$ and therefore the activity of nodes in $N^1(u)$ can only grow from round t to round $t+1$. We therefore in the following assume that $t' \in \{t, t+1\}$ such that $\Gamma_{\mathbb{A}}(u,t') \geq \frac{2}{3} \cdot \Gamma(u,t)$ and u is in state $\mathbb{A}$ in round t'.

To show that in round t', event H_u occurs, we distinguish the two cases given in the lemma statement. We start with the simpler case (a), where in round t, $1 \leq \Gamma(u) < 3\alpha(1)$ and $\gamma(u) = \frac{1}{2}$. Because no node in $N^1(u)$ switches to states $\mathbb{M}$ or $\mathbb{E}$ in round t, activity levels can only grow as they can grow at most by a factor of $1+\epsilon_\gamma$ in a single round, we know that $\Gamma(u,t) \leq \Gamma(u,t') \leq (1+\epsilon_\gamma)\Gamma(u,t)$. We know that $\Gamma_{\mathbb{A}}(u,t') \geq \frac{2}{3}\Gamma(u,t) \geq \frac{2}{3}$. Consequently, u is in state $\mathbb{A}$, it has activity level $\gamma(u,t') = \frac{1}{2}$, and the total activity level $\Gamma_{\mathbb{A}}(u,t') - \gamma(u,t')$ of all *neighbors* is between $\frac{1}{6}$ and $(1+\epsilon_\gamma)3\alpha(1) = O(1)$. Therefore, w.c.p., u and exactly one of its neighbors v operate on channel $\lambda = \mathcal{A}_1$. (Recall that a node w in state $\mathbb{A}$ chooses channel $\mathcal{A}_1$ with probability $\frac{\gamma(w)}{2}$.) Because we assume that u is $\frac{1}{3\alpha(1)}$-fat at time t, $\Gamma(v,t)$ is also bounded and therefore, w.c.p., no other neighbor of v picks channel $\mathcal{A}_1$. Hence, the only thing that is missing to show that event H_u occurs with constant probability is to prove that no other node in $N^2(u)$ hears a message on channel $\mathcal{A}_1$ in round t'. This follows from the following claim by choosing $S = N^2_{\mathbb{A}}(u) \cap (N(u) \cup N(v))$.

Claim 8.9. *Consider a round r, a set $S \subset N^2_{\mathbb{A}}(u,r)$, and let $\partial S \subseteq S$ be the nodes in S that are adjacent to some node in $N^2_{\mathbb{A}}(u,r) \setminus S$. For a channel $\lambda = \mathcal{A}_i \in \{\mathcal{A}_1, \ldots, \mathcal{A}_{n_\mathcal{A}}\}$, conditioned on the event that nodes in ∂S do not choose to operate on channel λ, the probability that a node in $N^2_{\mathbb{A}}(u,r) \setminus S$ receives a message on channel λ in round r is $1 - O(\pi_\ell)$.*

Proof of Claim 8.9. In the following, we use the notation $\mathcal{N} := N^2_{\mathbb{A}}(u,r) \setminus S$. Further, let A be the event that nodes in δS do not operate on channel λ and for a node $x \in \mathcal{N}$, let B_x be the event that node x receives a message on channel λ. Event B_x occurs iff x listens on channel λ and exactly one of its neighbors broadcasts on channel λ. The probability for a node $x \in \mathcal{N}$ to pick channel λ is $\gamma(x) \cdot 2^{-i}$. We therefore have

$$\begin{aligned}\mathbf{P}(B_x|A) &= \frac{\pi_\ell \gamma(x,r)}{2^i} \sum_{z \in N_{\mathbb{A}}(x) \setminus S} \frac{(1-\pi_\ell)\gamma(z,r)}{2^i} \cdot \prod_{y \in N_{\mathbb{A}}(x) \setminus (S \cup \{z\})} \left(1 - \frac{(1-\pi_\ell)\gamma(y,r)}{2^i}\right) \\ &\leq \frac{\pi_\ell \gamma(x,r)}{2^i} \frac{\Gamma(x,r)}{2^i} \cdot e^{-\Gamma(x,r)2^{-i}} = O\left(\frac{\gamma(x,r)}{\Gamma(x,r)} \cdot \pi_\ell\right).\end{aligned}$$

Let X be the number of nodes $x \in \mathcal{N}$ that receive a message on channel λ in round r. For the expectation of X, we then get

$$\mathbb{E}[X|A] = O(\pi_\ell) \cdot \sum_{x \in \mathcal{N}} \frac{\gamma(x,r)}{\Gamma(x,r)} = O(\pi_\ell).$$

The second equation follows from Lemma A.2 because the graph induced by $\mathcal{N}$ has independence at most $\alpha(2)$. Applying the Markov inequality, we get $\mathbf{P}(X \geq 1|A) \leq \mathbb{E}[X|A] = O(\pi_\ell)$, which concludes the proof of Claim 8.9. $\square$

We can now continue with the proof of Lemma 8.8. We have shown that in case (a), the event H_u occurs with constant probability. Let us therefore switch to case (b), where $N(u,t)$ is $\frac{1}{2}$-fat and $\Gamma(u,t) \geq 3\alpha(1)$. For the following argumentation, we define

$$\hat{N}^1_{\mathbb{A}}(u,t') := \left\{ v \in N^1_{\mathbb{A}}(u,t') : \Gamma_{\mathbb{A}}(v,t') \geq \frac{\Gamma_{\mathbb{A}}(u,t')}{2\alpha(1)} \right\} \quad \text{and} \quad \hat{\Gamma}_{\mathbb{A}}(u,t') := \sum_{v \in \hat{N}^1_{\mathbb{A}}(u,t')} \gamma(v,t').$$

To analyze the probability of the event H_u, consider two neighboring nodes $v, w \in N^1_{\mathbb{A}}(u,t')$. We define $L_{v,w}$ to be the event that in round t' both v and w decide to operate on channel $\lambda := \lceil \log_2 \Gamma(u,t) \rceil$ and no other node in $N(v) \cup N(w)$ chooses the same channel λ. Further $H_{v,w}$ is the event that $L_{v,w}$ occurs and in addition, no node in $N^2(u) \setminus (N(v) \cup N(w))$ receives a message on channel λ in round t'. Claim 8.9 implies that $\mathbf{P}(H_{v,w}|L_{v,w}) = 1 - O(\pi_\ell)$. Further, note that $H_u = \bigcup_{v,w \in N^1_{\mathbb{A}}(u), \{v,w\} \in E} H_{v,w}$, and we have $H_{v,w} = H_{w,v}$ and $H_{v,w} \cap H_{v',w'} = \emptyset$ for $\{v,w\} \neq \{v',w'\}$. It therefore holds that

$$\mathbf{P}(H_u) = \frac{1 - O(\pi_\ell)}{2} \cdot \sum_{\substack{\{v,w\} \in E, \\ (v,w) \in (N^1_{\mathbb{A}}(u,t'))^2}} \mathbf{P}(L_{v,w}). \tag{5}$$

The probability for a node $v \in \mathbb{A}$ to choose channel λ is $\gamma(v) \cdot 2^{-\lceil \log \Gamma(u,t) \rceil} \in \left[\frac{1}{2\Gamma(u,t)}, \frac{1}{\Gamma(u,t)}\right]$. We can therefore bound the probability that $L_{v,w}$ occurs in round t' as

$$\begin{aligned}
\mathbf{P}(L_{v,w}) \quad &\geq \quad \frac{1}{4} \cdot \frac{\gamma(v,t')\gamma(w,t')}{\Gamma(u,t)^2} \cdot \prod_{x \in N(v) \cup N(w)} \left(1 - \frac{\gamma(x,t')}{\Gamma(u,t)}\right) \\
&\geq \quad \frac{\gamma(v,t')\gamma(w,t')}{4\Gamma(u,t)^2} \cdot 4^{-\sum_{x \in N(v) \cup N(w)} \frac{\gamma(x,t')}{\Gamma(u,t)}} \\
&\geq \quad \frac{\gamma(v,t')\gamma(w,t')}{4\Gamma(u,t)^2} \cdot 4^{-(1+\epsilon_\gamma)\frac{\Gamma(v,t)+\Gamma(w,t)}{\Gamma(u,t)}} \\
&\geq \quad \frac{\gamma(v,t')\gamma(w,t')}{\Gamma(u,t)} \cdot 4^{-1-(1+\epsilon_\gamma)4\frac{\Gamma(u,t)}{\Gamma(u,t)}} \;=\; \frac{1}{4^{5+4\epsilon_\gamma}} \cdot \frac{\gamma(v,t')\gamma(w,t')}{\Gamma(u,t)}.
\end{aligned}$$

The last inequality follows because in round t, node u is $\frac{1}{2}$-fat. In the following, we restrict our attention to the events $L_{v,w}$ for $v \in \hat{N}^1_{\mathbb{A}}(u,t')$ as these are the only ones for which we obtain a significant lower bound on the probability that they occur. For $v \in \hat{N}^1_{\mathbb{A}}(u,t')$, let $K_v := \bigcup_{w \in N(v) \cap N^1_{\mathbb{A}}(u,t')} L_{v,w}$ be the event that $L_{v,w}$ occurs for some neighbor w of v. For a node $v \in \hat{N}^1_{\mathbb{A}}(u,t')$, we then have

$$\begin{aligned}
\mathbf{P}(K_v) &= \sum_{w\in N(v)\cap N^1_{\mathbb{A}}(u,t')} \mathbf{P}(L_{v,w}) \\
&\geq \frac{1}{4^{5+4\epsilon_\gamma}} \cdot \frac{\gamma(v,t')}{\Gamma(u,t)^2} \sum_{w\in N(v)\cap N^1_{\mathbb{A}}(u,t')} \gamma(w,t') \\
&= \frac{1}{4^{5+4\epsilon_\gamma}} \cdot \frac{\gamma(v,t')\big(\Gamma_{\mathbb{A}}(v,t') - \gamma(v,t')\big)}{\Gamma(u,t)^2} \\
&\geq \frac{1}{4^{5+4\epsilon_\gamma}} \cdot \frac{\gamma(v,t')\big(\Gamma_{\mathbb{A}}(u,t') - \alpha(1)\big)}{2\alpha(1)\Gamma(u,t)^2} \qquad (6) \\
&\geq \frac{1}{4^{5+4\epsilon_\gamma}} \cdot \frac{\gamma(v,t')}{6\alpha(1)\Gamma(u,t)}.
\end{aligned}$$

Inequality (6) follows because $\gamma(v,t') \leq \frac{1}{2}$ and since $v \in \hat{N}^1_{\mathbb{A}}(v)$ and thus $\Gamma_{\mathbb{A}}(v,t') \geq \frac{\Gamma_{\mathbb{A}}(u,t')}{2\alpha(1)}$. The last inequality follows from $\Gamma_{\mathbb{A}}(u,t') \geq \frac{2}{3} \cdot \Gamma(u,t) \geq \frac{2}{3} \cdot 3\alpha(1) = 2\alpha(1)$. Using (5), we can now bound the probability of event H_u in round t' as

$$\mathbf{P}(H_u) \geq \frac{1-O(\pi_\ell)}{2} \sum_{v\in\hat{N}^1_{\mathbb{A}}(u,t')} \mathbf{P}(K_v) \geq \frac{1-O(\pi_\ell)}{3\alpha(1)4^{6+4\epsilon_\gamma}} \cdot \sum_{v\in\hat{N}^1_{\mathbb{A}}(u,t')} \frac{\gamma(v,t')}{\Gamma(u,t)} = \frac{1-O(\pi_\ell)}{3\alpha(1)4^{6+4\epsilon_\gamma}} \cdot \frac{\hat{\Gamma}_{\mathbb{A}}(u,t')}{\Gamma(u,t)}. \qquad (7)$$

Applying Lemma A.2 to the graph induced by the nodes in $N^1_{\mathbb{A}}(u,t')$, the activity sum of nodes in $\hat{N}^1_{\mathbb{A}}(u,t')$ can be lower bounded as

$$\hat{\Gamma}_{\mathbb{A}}(u,t') \geq \frac{\Gamma_{\mathbb{A}}(u,t')}{2\alpha(1)} \geq \frac{\Gamma(u,t)}{3\alpha(1)}.$$

Together with (7), this proves that also in case (b), the event H_u occurs with constant probability in a round $t' \in \{t, t+1\}$. Note also that in both cases (a) and (b), for π_ℓ sufficiently small, the probability that H_u occurs can be lower bounded by a constant q that is independent of the probability π_ℓ.

To complete the proof, assume that in round t', event H_u occurs with probability q and if it occurs, nodes v and w are the two nodes in $N^1(u)$ participating on channel λ (channel $\mathcal{A}_1$ in case (a)). Let M be the event that no herald is created on a channel $\mathcal{A}_i \neq \lambda$ in round t'. Clearly, the probability that M occurs is lower bounded by the probability that no herald is created on any channel in round t'. By Lemma 8.6, we therefore have $\mathbf{P}(M) = 1 - O(\pi_\ell)$. For the probability that events H_u and M both occur, we then get

$$\mathbf{P}(M \cap H_u) = 1 - \mathbf{P}(\overline{M} \cup \overline{H}_u) \geq 1 - \mathbf{P}(\overline{M}) - \mathbf{P}(\overline{H}_u) = q - O(\pi_\ell).$$

Recall that probability q is a constant independent of π_ℓ. Conditioned on the event that $M \cap H_u$ occurs, the probability that one of the two nodes v, w listens on channel λ and the other one broadcasts on the channel is $2\pi_\ell(1-\pi_\ell)$. In that case one of the two nodes becomes a herald candidate and the other one its leader candidate. Also, $M \cap H_u$ implies that in round t' no other herald candidates are created in $N^2(u)$. Let t'' be the round in $\{t, t+1\} \setminus t'$. If in addition in round t'' and in the remaining rounds $t+2, \ldots, t+7$ no herald candidate is created in $N^2(u)$, nodes v and w make it through the handshake and become an isolated leader-herald pair as claimed by the lemma. By Lemma 8.6, this happens with probability $1 - O(\pi_\ell)$, which by choosing π_ℓ sufficiently small concludes the proof. □

8.3 The Analysis of the Handshake

In the following lemma, we study the circumstances under which two adjacent leader-herald pairs can coexist.

Lemma 8.10. *In round r consider two leader-herald pairs (l_1, h_1) and (l_2, h_2) and suppose that the pairs started their most recent handshakes in rounds r_1 and r_2, $r_1 \leq r_2$, respectively. Say that edge e is* crossing *if one of its endpoints is in $\{l_1, h_1\}$ and its other endpoint is in $\{l_2, h_2\}$. Then, either no crossing edge exists or exactly one of the following conditions is satisfied: (1) $r_1 = r_2$ and crossing edges are $\{l_1, l_2\}$ and/or $\{h_1, h_2\}$, (2) $r_2 = r_1 + 2$ and the only crossing edge is $\{l_1, h_2\}$.*

Proof. Since both pairs are in $\mathbb{L} \times \mathbb{H}$ in round r, they successfully finished their respective handshakes at the end of rounds $r_1 + 5$ and $r_2 + 5$, respectively. Suppose that there is a crossing edge. First, assume that $r_2 \geq r_1 + 4$. Then, since nodes (l_1, h_1) block the channel $\mathcal{H}$ every other round starting with round $r_1 + 6 \leq r_2 + 2$, the handshake of pair (l_2, h_2) cannot succeed, a contradiction.

Suppose $r_2 = r_1 + 3$. In round $r_1 + 3$, leader l_1 and herald h_2 transmit. So there is no edge between l_1 and l_2, or between h_1 and h_2, because if there was, a collision would happen at l_2 resp. h_1, causing the handshake to fail. In round $r_1 + 4$, heralds h_1 and h_2 transmit. By the same logic there is no edge between h_1 and l_2, or between h_2 and l_1, a contradiction to our assumption of an existing crossing edge.

Suppose $r_2 = r_1 + 1$. In round $r_1 + 1$, heralds h_1 and h_2 transmit. So there is no edge between h_1 and l_2, or between h_2 and l_1. In round $r_1 + 2$, leader l_1 and herald h_2 transmit. So there is no edge between l_1 and l_2, or between h_1 and h_2.

Suppose $r_2 = r_1 + 2$. In round $r_1 + 2$, leader l_1 and herald h_2 transmit. So there is no edge between l_1 and l_2, or between h_1 and h_2. In round $r_1 + 6$, herald h_1 transmits on channel $\mathcal{H}$ and leader l_2 is receiving a message from h_2. So there is no edge between h_1 and l_2. Thus, the only existing crossing edge is (l_1, h_2), between the older leader and the younger herald.

Finally, let $r_1 = r_2$. If l_1 neighbors h_2 or l_2 neighbors h_1, then in both cases the handshakes of one of both pairs (l_1, h_1) or (l_2, h_2) fails in round $r_1 = r_2$, contradicting our assumption of a successful handshake for both pairs. Due to the synchronous actions a crossing edge between two leaders or two heralds can exist. ☐

8.4 The Analysis of the Red-Blue Game

We next study the exact guarantees of when and how pairs get knocked out in the red-blue games.

Definition 8.11. *(**Maturity**) We say that candidate v is* mature *in round t, if v is in round* 5 *or* 6 *of its respective handshake.*

Definition 8.12. *(**Good Pair, Bad Pair**) Consider a leader-herald pair (l, h) in round t. We say pair (l, h) is a* good pair *if in round t none of the neighbors of l (other than h) is in state $\mathbb{L}$, $\mathbb{H}$ or is a mature candidate. Otherwise we say that (l, h) is a* bad pair*.*

Lemma 8.13. *If a pair (l, h) is good in round t and they started their first red-blue game in round r, then, w.h.p., either*

- *the related leader l joins the MIS by the end of round $r + \tau_{red-blue} = r + O(\log n)$, or*
- *a node $v \in N(l) \cup N(h)$ joins the MIS before round $r + \tau_{red-blue} = r + O(\log n)$ by increasing its lonely counter above τ_{lonely}.*

Proof. Consider a herald-leader pair (l, h) that is good in round t.

The herald h only listens in round 6 of the red-blue game, but the only nodes sending on channels $\mathcal{R}_k$ are leaders or MIS nodes. So, in round $t' \geq t$, h either gets knocked out by its own leader l, which only happens if l records a failure itself and which we cover below. Or h gets knocked out by a different leader which is also in round 6 of its red-blue game, which, by Lemma 8.10, cannot happen, because at time t, h is already a herald and thus made it through the handshake successfully. Or h neighbors an MIS node v in round t'. If v moved to state $\mathbb{M}$ from state $\mathbb{L}$, then v started its first red-blue game in round $r - 2$ by Lemma 8.10, i.e., it survived $\Theta(\log n)$ games while being next to h—w.h.p., this cannot happen. Thus, for h to be knocked out by $v \neq l$, v needs to join the MIS via the *lonely* counter.

The leader l only listens in either round 2 or round 4 of a particular red-blue game on channel $\mathcal{G}$. It thus gets knocked out in round $t' \geq t$ only if another node $v \neq h$ sends in that particular round on $\mathcal{G}$. Let v be another leader. By Lemma 8.10 that leader started its handshake in round $r - 6$, too, but then l neighbors that leader also in round t, which is a contradiction to the definition to (l, h) being a good pair in round t. Thus, let v be a herald in round t'. By Lemma 8.10 v started its red-blue game 2 rounds after l, but then v was either a herald or a mature candidate by time t, also a contradiction to the definition of a good pair. So v has to be an MIS node. If v went through state $\mathbb{L}$ then by the time t' node l also is an MIS node due to Lemma 8.10, so both nodes survived $\Theta(\log n)$ games with none getting knocked out, w.h.p. that does not happen. Thus v joined the MIS via the *lonely* counter. ☐

Lemma 8.14. *Consider a node v and suppose that in an arbitrary round t, there is a leader or herald of a bad pair in $N^3(v)$. Then, with constant probability, in round $t+12$, no node in $N^3(v)$ is in state $\mathbb{H}'$ and all leaders and heralds are part of a good pair.*

Proof. Consider all the bad leader-herald pairs (bad at time t) and all the pairs $(l', h') \in \mathbb{L}' \times \mathbb{H}'$ at time t with both ends in $N^4(v)$, which have finished their handshake successfully by the end of round $t+5$. Let us call each such pair *interesting.* We know that with a constant probability, in rounds $t, \ldots, t+12$, no node in $N^3(v)$ other than those in interesting or good pairs is in a state in $\{\mathbb{L}, \mathbb{H}, \mathbb{H}'\}$. We assume this as given. For each interesting pair, also consider their 6 round red-blue game that is fully contained in round interval $[t+1, t+11]$. We always consider the first element in our pair notation the leader or leader candidate of the corresponding interesting pair.

For an interesting pair (l, h) we say it *conflicts*, if l neighbors another leader or herald node in $[t, t+5]$. We show that with constant probability all such pairs get knocked out by round $t+12$. An interesting pair, which is not conflicting satisfies the conditions of a good pair by round $t+5$, which does not conflict with the statement of the lemma. Thus we also assume all interesting pairs to be conflicting.

Consider an interesting pair (l_1, h_1) and suppose that they finish their related red-blue game in round $r_1 \in [t+6, t+11]$. We say (l_1, h_1) is *isolated* if there does not exist another pair (l_2, h_2) (which *has* to be interesting, too) such that l_1 and l_2 are adjacent.

Given this definition, the proof has two parts. We first show that with constant probability, no conflicting interesting isolated pair remains *not knocked out.* Then in the second part, we show that with at least a constant probability, no non-isolated conflicting interesting pair remains not-knocked out. These two show that no interesting conflicting pairs can remain *not knocked out* which would complete the proof.

The total number of isolated interesting pairs is bounded by $\alpha(3)$. Hence, the probability that the leader of each isolated interesting pair is red in its red-blue game fully contained in rounds $[t+1, t+11]$ is lower bounded by $2^{-\alpha(3)} = \Omega(1)$. In the rest of the proof, we assume that the leader of each isolated interesting pair is red (their choices are independent). Consider two neighboring interesting isolated pairs (l_1, h_1) and (l_2, h_2) and suppose these two pairs started their red-blue games in rounds respectively $r_1, r_2 \in [t+1, t+6]$ and $r_1 \leq r_2$. By definition of isolated pairs and from Lemma 8.10, we get that either (1) $r_1 = r_2$ and the only crossing edge is between h_1 and h_2, or (2) $r_2 = r_1 + 2$ and the only crossing edge is between l_1 and h_2. In the first case, these two pairs do not conflict with each other. In the second case, since l_1 is red in the red-blue game starting in round r_1, in round $r_1 + 3$, leader l_1 does not receive the message of its herald h_1 on channel $\mathcal{G}$. This is because, in that round, h_2 is transmitting on channel $\mathcal{G}$. Hence, pair (l_1, h_1) gets knocked out. This completes the study of interesting isolated pairs since all of them are conflicting, i.e., their leader (resp. possible leader candidate at time t) neighbors another herald and the color choice causes them to be knocked out.

In order to complete the proof, we show that with at least a constant probability, each non-isolated interesting pair is knocked out. Consider two non-isolated interesting pairs (l_1, h_1) and (l_2, h_2) such that l_1 and l_2 neighbor each other. Due to Lemma 8.10 both pairs finish their red-blue game in the same round.

Let L_r be the set of non-isolated leaders of interesting pairs that finish their red-blue game in round $r \in [t+1, t+6]$, and let H_r be the graph induced on leaders L_r. Note that for each $r \in [t+1, t+6]$, H_r is an induced subgraph of $N^4(v)$ and thus, it has maximum independent size at most $\alpha(4)$.

Let S_r be a maximal independent set in H_r. Then, let T_r be a maximal independent set in $L_r \setminus S_r$. Also, define T'_r as follows: for each non-isolated leader $l \in S_r$ that does not have a neighbor in T_r, add an arbitrary neighbor of l to T'_r. It is clear that $|S_r| \leq \alpha(4)$, $|T_r| \leq \alpha(4)$, and $|T'_r| \leq |S_r| \leq \alpha(4)$. Thus, with a constant probability, all leaders in S_r choose red color while all leaders in sets T_r and T'_r choose blue color, in their red-blue game ending in round r. For each leader $l \in L_r$, if $l \in S_r$, then $T_r \cup T'_r$ contains a neighbor of l, and if $l \notin S_r$, then l has one neighbor in S_r and one neighbor in T_r. Thus, if all leaders in S_r choose red color while all leaders in sets T_r and T'_r choose blue color, then each non-isolated leader $l \in L_r$ has at least one adjacent non-isolated leader in H_r that has a color different than that of l. If this happens, each non-isolated leader $l \in L_r$ is knocked out. Thus, with constant probability, each non-isolated leader $l \in L_r$ is knocked out. Hence, noting that adding adjacent leaders can only increase the probability of getting knocked out (note that pairs continue their game till the end of the game even if they are knocked out), we conclude that, with constant probability, for each $r \in [t+1, t+6]$, each non-isolated leader $l \in L_r$ is knocked out. This completes the proof. □

8.5 The Analysis of the MIS State

Here we present the main safety guarantee of our MIS algorithm.

Lemma 8.15. *W.h.p., the nodes in state* $\mathbb{M}$ *always form an independent set. Moreover, if a node* v *enters state* $\mathbb{M}$ *in round* t *and a node* $w \in N(v)$ *is awake in round* t*, then in round* $t' = t + O(\log n)$*, w.h.p.,* w *is in state* $\mathbb{E}$*.*

Proof. We start by proving the first part of the lemma. Fix two arbitrary adjacent nodes u and v and let $\mathcal{E}_{u,v}$ be the event that nodes u and v are the first pair of adjacent nodes (first in time) that are in the $\mathbb{M}$ simultaneously. Here, if two or more adjacent node pairs are in $\mathbb{M}$ simultaneously (i.e., in the case of a tie between pairs), then break the tie based on the ids the nodes (comparing the smaller ids of the two pairs first and then the larger ids of the two pairs). We show that for each pair of adjacent nodes u and v, w.h.p., event $\mathcal{E}_{u,v}$ does not happen. A union bound over all pairs then shows that, w.h.p., no two adjacent nodes are in state $\mathbb{M}$ simultaneously, proving the independence claim.

Considering the fact that nodes can enter $\mathbb{M}$ through two paths—(P1) through being a leader, or (P2) because of having *lonely* $> \tau_{lonely}$—we divide the analysis into three cases (A), (B) and (C) as follows.

(A) Nodes u and v both take path (P1), i.e., they both enter $\mathbb{M}$ through becoming a leader. Using Lemma 8.10, we can infer that adjacent nodes u and v can enter (and coexist in) state $\mathbb{M}$ by taking path (P1) only if u and v enter state $\mathbb{M}$ simultaneously. In this case, nodes u and v perform $\Theta(\log n)$ red-blue games synchronously. In each of these red-blue games, with probability 0.5, nodes u and v choose different colors and thus, both of them get knocked out. Hence, during $\Theta(\log n)$ red-blue games, w.h.p., the two nodes u and v get knocked out.

(B) Nodes u and v both take path (P2), i.e., they both enter $\mathbb{M}$ through having $\mathit{lonely} > \tau_{\mathit{lonely}}$. Without loss of generality, suppose that v enters $\mathbb{M}$ no later than u. Then, consider $\tau_{\mathit{lonely}} = \Theta(\log n \log\log n)$ rounds before the time that u enters $\mathbb{M}$. If v was in the MIS for $\Omega(\log n)$ rounds before u joins, w.h.p., v would eliminate u before that point via broadcasting on report channels. Thus assume that they join $\mathbb{M}$ within $O(\log n)$ rounds, implying that they spent $\Theta(\tau_{\mathit{lonely}})$ rounds next to each other in states $\mathbb{A}$ and $\mathbb{L}'$. Note that in each two consecutive rounds of this interval, with at least a constant probability, there is a round that both u and v are in state $\mathbb{A}$. Let r be one of these rounds and let k be the number of neighbors of u in state $\mathbb{A}$. Note that $k \geq 1$ because v is in among them. Moreover, $k \in O(\log^3 n)$. Thus, in round r, u has a probability $\Omega(\frac{1}{\log\log n})$ for listening to channel $\mathcal{S}_j$ where $j = \lceil \frac{1}{k} \rceil$. If u listens to this channel, the probability that exactly one of its neighbors in state $\mathbb{A}$ transmits on channel $\mathcal{S}_j$ is $\frac{k}{2^{-j}}(1 - \frac{1}{2^{-j}})^{k-1} = \Theta(1)$. That is, in round r, with probability at least a positive constant, u receives a message from v or some other node. Hence, throughout the $\Theta(\tau_{\mathit{lonely}}) = \Theta(\log n \log\log n)$ rounds of this period, with high probability, u receives a message and thus, it does not enters $\mathbb{M}$.

(C) One of the nodes u and v takes path (P1) while the other one takes path (P2). That is, without loss of generality, v enters $\mathbb{M}$ through becoming a leader of a good pair and u enters $\mathbb{M}$ through having $\mathit{lonely} > \tau_{\mathit{lonely}}$. Let r_v be the round in which v starts its first red-blue game and let r'_v be the round in which v enters $\mathbb{M}$. Moreover, let r_u be the round in which u enters herald filter and let r'_u be the round in which u enters state $\mathbb{M}$. Note that $r'_v - r_v \geq \tau_{\mathit{red\text{-}blue}} = \Theta(\log n)$ and $r'_u - r_u \geq \tau_{\mathit{lonely}} = \Theta(\log n \log\log n)$. We emphasize that this is a deterministic guarantee.

First consider the case $r'_v \geq r'_u + \frac{1}{2}\tau_{\mathit{red\text{-}blue}}$. In that case u had $\frac{1}{2}\tau_{\mathit{red\text{-}blue}}$ rounds to disrupt v's red-blue games. Consider a single 6-round red-blue game of v, starting at a round $r \geq r'_u$. u's action in round r might be deterministic, but w.c.p. in round $r+1$ it sends on $\mathcal{G}$. Similarly, w.c.p., v chooses blue for that red-blue game, in which case v gets knocked out. This applies for each red-blue game after round r'_u and thus, using a Chernoff bound, w.h.p., after $O(\log n)$ games, v is knocked out.

Next consider $r'_v < r'_u + \frac{1}{2}\tau_{\mathit{red\text{-}blue}}$. Then u neighbors v for at least $\frac{1}{2}\tau_{\mathit{red\text{-}blue}}$ rounds while v is in state $\mathbb{L}$, before v moves to state $\mathbb{M}$. Since u and v are the first nodes that violate independence, u has no more than $\alpha(1)$ neighbors in state $\mathbb{M}$ during that interval. If v broadcasts on one of the report channels in round r of the interval, then u might have more than $\alpha(1)$ leaders nearby, in which case at least one of them is part of a bad pair.

By Lemma 8.14 in round $r+12$ w.c.p. all leaders in $N(u)$ are part of good pairs and no node in $N^1(u)$ went to state $\mathbb{H}'$ in rounds $r, \ldots, r+12$ due to Lemma 8.6—in the case that Lemma 8.14 provides us with a new MIS node instead, notice that this can happen at most a constant number of times. Since u moves to state $\mathbb{M}$ via its *lonely* counter, w.c.p., u does not broadcast in round $r+11$, i.e., it is in state $\mathbb{A}$ in round $r+12$. Thus, w.c.p., it hears from v in that round. Applying a Chernoff bound again concludes the claim and thus the first part of the lemma.

For the second part of the lemma, we now assume the first part. However, we remark that we proved the statement of the second part directly in some of the cases of the first part. Suppose node v is in state $\mathbb{M}$ in round t and consider an arbitrary neighbor w of v that is awake in round t. Note that in rounds $\geq t$, v blocks channel $\mathcal{H}$ once every two rounds. Moreover, once every two rounds in rounds $\geq t$, v blocks channel $\mathcal{G}$ with a constant probability. Finally, once every two rounds in rounds $\geq t$, v transmits on a channel $\mathcal{R}_k$. We know that in rounds $\geq t+6$, w is in one of states $\mathbb{L}', \mathbb{H}', \mathbb{H}, \mathbb{E}$. Hence, in every 6 consecutive rounds after $t+6$, w receives a message from v with constant probability. Therefore, with high probability, by round $t+\Theta(\log n)$, w has received a message from v and has thus quit the algorithm, by entering state $\mathbb{E}$. □

8.6 Putting the Pieces Together

In this section, we wrap everything up to show that guarantees (G1), (G4) and (G5) hold. Together with the guarantees (G2) and (G3) handled in Section 6 we finalize the proof of Theorem 5.1.

Lemma 8.15 immediately proves (G1) for all nodes currently in the herald filter. All nodes in the decay filter, both at the beginning as well as during the main body part, listen to the report channels w.c.p. in every round. Thus, either they learn of a neighboring MIS node within $O(\log n)$ rounds, or they move forward to the herald filter in that time bound. In the latter case, after that transition, Lemma 8.15 takes care of those nodes.

Lemma 8.15 also immediately gives us (G5). The only thing that remains to be shown is the progress guaranteed by (G4). To do so, we use the following lemma.

Lemma 8.16. *Consider two neighboring nodes u and v such that both are in the herald filter at time t and assume that no node in $N(u) \cup N(v)$ has joined the MIS by time t. Then, w.h.p., some node u' in the $O(\log\log n)$-neighborhood of u joins the MIS between times t and $t+O(\log n)$.*

Proof. W.l.o.g., assume that at time t, the activity of u and v have already reached the maximum level $\gamma(u) = \gamma(v) = \frac{1}{2}$. Otherwise, there is a new MIS node in the $O(\log\log n)$-

neighborhood of u by time $t' = t + O(\log\log n)$ when the activity level of both u and v reach $\frac{1}{2}$ or we can start the argument at time t'.

Clearly, $\Gamma(u) \geq 1$. W.h.p., within $O(\log n)$ rounds all MIS nodes in $N^3(u)$ inform their neighbors about their state, eliminating them.

First assume that $\Gamma(u) < 3\alpha(1)$ for $O(\log n)$ rounds and that u is $\frac{1}{3\alpha(1)}$-fat. In that case we can apply Lemma 8.8 to get that w.c.p. a node u' in $N(u)$ moves to state $\mathbb{L}$ alone, i.e., forming a good pair with some node v'. If this is the case, then by Lemma 8.13, w.c.p., we will get an MIS node in $N^2(u)$ within $O(\log n)$ rounds, which finishes the claim. If u' however is part of a bad pair, then w.c.p. within 12 rounds all members of bad pairs in $N^2(u)$ go back to state $\mathbb{A}$. By Chernoff at most $O(\log n)$ bad pairs are created before a good pair is, also finishing our claim. Thus consider the case in which $\Gamma(u) < 3\alpha(1)$ for $O(\log n)$ rounds, but u is not sufficiently fat. By the definition of fatness and the fact that $\Gamma(u) \geq 1$, however, this implies that one of u's neighbors v has an activity sum of $\Gamma(v) \geq 3\alpha(1)$. If v itself is not $\frac{1}{2}$-fat, then, due to guarantee (G2) from decay filter, within distance $\delta = O(\log\log n)$ there is a node w that is both $\frac{1}{2}$-fat and has $\Gamma(w) \geq 2^\delta \cdot 3\alpha$. In which case we can again apply Lemma 8.8. Standard Chernoff gives us the desired result.

Thus let $\Gamma(u)$ rise above $3\alpha(1)$ within $O(\log n)$ rounds. Note that the only way for $\Gamma(u)$ to drop is to have a new MIS node within $N^2(u)$ arising, so we assume that this is not the case. With analogous reasoning to the previous case in which u was not sufficiently fat, we can apply to u the same chain of reasoning that we applied to v before. But that covers all possible cases, concluding our proof.

□

Lemma 8.17. *(G4): W.h.p., for each node v that is in the herald filter in round r, by round $r' = r + \tilde{O}(\log n)$, v is* dominating *or* dominated*. In the latter case, v has a dominating neighbor.*

Proof. Consider a node u in the herald filter in round r. If u does not neighbor any other node in the herald filter, within $O(\log n \log\log n)$ rounds it decides to move on to state $\mathbb{M}$ or it stops being alone. Since the former completes the proof we assume the latter and let v be one of u's neighbors in the herald filter. If v is an MIS node, then by Lemma 8.15 u is dominated within $O(\log n)$ rounds, completing the proof. Thus we assume that $N(u) \cap \mathbb{M} = \emptyset$.

Case 1: v already neighbors a node in the MIS. But then within $\tau_{lonely} = O(\log n)$ rounds v becomes eliminated and either u is alone again or it neighbors another node v' in the herald filter.

Case 2: $(N(u) \cup N(v)) \cap \mathbb{M} \overset{(*)}{=} \emptyset$. The conditions for Lemma 8.16 are satisfied and we get that a new MIS node arises in $N^d(u)$ with $d = O(\log\log n)$ within $O(\log n)$ rounds. W.h.p., every new MIS node complies with Lemma 8.15 and thus this can happen at most $\alpha(d)$ times before the conditions of Lemma 8.16 can no longer be satisfied. Since α is a polynomial, the total amount of time needed until condition $(*)$ is violated *or* it has no neighbor v anymore in the herald filter, is $O(\alpha(d)\log n) = O\big(\log n\,(\log\log n)^{\deg(\alpha)}\big) = \tilde{O}(\log n)$.

In both cases, after $\tilde{O}(\log n)$ rounds u either

- joins the MIS itself or
- one of its neighbors does and it dominates u or
- a node in distance 2 joins the MIS and eliminates all its neighbors.

The first two possibilities complete our proof again. In the latter, u is either alone or it neighbors a different node v' such that one of the above two cases holds and we repeat. In total, however, u can drop back to being alone or neighboring a different node v' at most $\alpha(2)$ times, and every time it has to wait at most $\tau_{lonely} + O(\alpha(d)\log n) = \tilde{O}(\log n)$ rounds. This completes our proof.

□

9 Connected Dominating Set and its Applications

In this section, we use our MIS algorithm solution as a building block in solving other problems efficiently in the multichannel environment. Our main technical result is a new algorithm that uses the MIS solution as a subroutine to build a constant-degree connected dominating set (CDS) in $O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds. We then use this structure as an overlay to derive solutions to broadcast and leader election that run in $O\big(D + \frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds, and to k-message multi-message broadcast that runs in $O\big(D+k+\frac{\log^2 n}{\mathcal{F}}\big)+\tilde{O}(\log n)$ and $O\big(D+k\log n+\frac{\log^2 n}{\mathcal{F}}\big)+\tilde{O}(\log n)$ rounds for unrestricted and restricted message sizes, respectively.

9.1 Connected Dominating Set

First, we show how to construct a constant-degree CDS in $O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds. At a high-level, our solution builds an MIS and then connects every pair of MIS nodes that are within 3 hops using a constant-length path. As argued in [8], the resulting structure is a constant-degree CDS. [3]

Our algorithm faces three key challenges:

(1) MIS nodes must keep revisiting the MIS algorithm to prevent newly activated nodes from joining the MIS;

(2) since nodes might end the MIS stage at different rounds, they might also start the CDS stage at different rounds, causing synchronization issues (the CDS stage cycles through fixed-length phases);

[3]Technically, the argument in [8] leverages a slightly stronger constraint on the graph structure than the bounded independence we assume here. But it is straightforward to see that the same constant-degree property holds for the latter property.

(3) MIS nodes must discover MIS nodes within distance 3 within $\tilde{O}(\log n)$ rounds.

We overcome the first challenge by letting MIS nodes act w.c.p. on the report channels $\mathcal{R}_1, \ldots, \mathcal{R}_{3\alpha(1)}$ every constant number of rounds.

For the second challenge we let MIS nodes participate only with a constant probability. Non-MIS nodes also only participate if they have exactly one participating MIS node in their neighborhood with whom they are synchronized. This will eventually lead to the event that only a single MIS node and some of its neighbors are active, which is enough for the algorithm to progress.

The third challenge is overcome by letting non-MIS nodes disseminate information about their dominating nodes with a strategy that considers the possibility of having many neighboring nodes that 'serve' the same nodes in the ongoing cycle.

Overview. We provide a detailed description of the algorithm below before proceeding with the analysis. Pseudo-code for the algorithm is given by Algorithm 7.

We assume the algorithm can make use of its own collection of $\Theta(\mathcal{F})$ channels, which are disjoint from those used by the MIS algorithm. Those channels are $\mathcal{C}_1, \ldots, \mathcal{C}_{n_\mathcal{C}}$ and the special channel $\mathcal{C}$. For simplicity, we assume here that $n_\mathcal{C} = \log n$, and explain how to adapt the protocol in case $n_\mathcal{C}$ is smaller.

Transition from the MIS to CDS Stage. Our CDS algorithm begins with nodes running our MIS algorithm. We call this the *MIS stage*. Eventually, nodes move on to the CDS-specific routines described below. We call this the *CDS stage*. We must be careful about this transition to ensure that a node that advances to the CDS stage does not neglect its obligations to the MIS algorithm. In more detail, when a node u is elected to the MIS we have it remain in state $\mathbb{M}$ for an additional $c_1 \log n$ rounds, for some constant $c_1 \geq 1$ we fix below. Only after these rounds does it move on to the CDS stage. At this point, we still require u to return every constant number of rounds to the report channels monitored in the first $\Theta(\log n)$ rounds of the decay filter, and broadcast with constant probability. Notice, if we choose a sufficiently large constant c_1, by the time u enters the CDS stage, every awake neighbor of u either knows about u or is early enough in the initial listening stage of the decay filter that it will subsequently receive a message from u on one of the report channels before it can advance. It is here that we set c_1 to get sufficiently high probability of this property holding. On the other hand, a node that learns that it is *not* in the MIS, can immediately advance to the CDS stage.

Handling Asynchronous Starts. Because we assume asynchronous starts, nodes will potentially finish the MIS stage at different rounds. We must ensure the CDS stage handles this properly. In more detail, notice that the CDS stage described below consists of four constant-round phases—*announce, dense search, sparse search*, and *report*—which it cycles through again and again. We call each cycle through all four phases an *instance*. Due to

Algorithm 7 Non-MIS node v participating in instance i of MIS node u with $\mathcal{F} = \Omega(\log n)$

path, contains paths that includes v
discovered, contains discovered MIS nodes
knowledge, an array such that $knowledge[u]$ contains the MIS nodes u knows about at this point

▷ Announce phase (1 round)
Set $bcast \leftarrow 1$ with probability $\frac{1}{2}$ otherwise $bcast \leftarrow 0$
if $path \neq \perp$ **and** $bcast = 1$ **then**
    BROADCAST($\langle non - MIS, path \rangle$ on channel $\mathcal{C}$)
else
    LISTEN(on channel $\mathcal{C}$)
    if received message $m = \langle MIS, u, S \rangle$ **then**
        $knowledge[u] \leftarrow knowledge[u] \cup S$
        if S contains a path that includes v, add this to *path* and **join CDS**
▷ Sparse Search phase (1 round)
Set $bcast \leftarrow 1$ with probability $2^{-(i \textbf{ mod } \lceil \log \log n \rceil)}$ otherwise $bcast \leftarrow 0$
if $bcast = 1$ **then**
    BROADCAST($\langle u \rangle$ on channel $\mathcal{C}$)
else
    move to next round
▷ Dense Search phase (1 round)
Set $bcast \leftarrow 1$ with probability $\frac{1}{2}$ otherwise $bcast \leftarrow 0$
if $bcast = 1$ **then**
    Choose channel $\mathcal{C}_c$ with probability 2^{-c}
    BROADCAST($\langle u \rangle$ on channel $\mathcal{C}_c$)
else
    wait until next round
▷ Report phase (3 rounds)
if *discovered* contains nodes R not in $knowledge[u]$ **then**
    with probability $\frac{1}{2}$ **do**
        BROADCAST(R on channel $\mathcal{C}$) ▷ – round 1 of phase
    $herald \leftarrow 0$ ▷ – round 2 of phase
    choose channel $\mathcal{C}_c$ with probability 2^{-c}
    with probability $\frac{1}{2}$ **do**
        BROADCAST(R on channel $\mathcal{C}_c$)
    otherwise
        LISTEN(on channel $\mathcal{C}_c$)
        if received a message **then**
            $herald \leftarrow 1$
    if $herald = 1$ **then** ▷ – round 3 of phase
        BROADCAST(R on channel $\mathcal{C}$)
    else
        move to next round
else
    wait 3 rounds

bounded independence (which bounds the number of MIS nodes in a region), a non-MIS node entering the CDS stage will hear from its MIS neighbor(s) already in the stage within $O(\log n)$ rounds, and a new MIS node entering the stage will notify all its neighbors of its presence also within $O(\log n)$ rounds.

Our remaining concern is the possibility that the phases executed by one MIS node might be out of synchronization with those of a nearby MIS node. We will address this situation below by having non-MIS nodes participate in a given instance of an MIS neighbor only if that neighbor flipped heads for the instance while all nearby overlapping instances flipped tails (we formalize this below). A non-MIS node that is not participating in an instance in a given round, generates a random bit. If the bit is 0, it listens on channel $\mathcal{C}$, otherwise it chooses a channel from $\mathcal{C}_1, \ldots, \mathcal{C}_{n_\mathcal{C}}$ with uniform randomness and listens. We will show this provides sufficient synchronization for the search phases to function.

CDS Stage: MIS Nodes. Each MIS node v in the CDS stage divides rounds into instances each consisting of four phases.

An MIS node v begins by generating a string of random bits with uniform independent probability. We call this its *coin sequence*. This coin sequence will be shared with its neighbors (i.e., non-MIS nodes that u dominates) who will use it to decide whether or not to participate in a given instance. For simplicity, we assume v fits sufficiently many bits in a single message to satisfy the demands of this stage (which consumes 1 bit per instance). In practice, v can simply update the bits in its message on a regular basis.

In more detail, during the first phase of each instance, if v has a 1 in the corresponding bit of its coin sequence, it will broadcast, w.c.p., an *announcement* message on channel $\mathcal{C}$ that contains: its coin sequence, the *ID*s of all MIS nodes within 3 hops that it knows about, and a 2 or 3 hop path to each such nearby node. During the other phases of the instance, v listens on channel $\mathcal{C}$. If v learns about a new nearby MIS node during any of these rounds, and v's *ID* is smaller, it chooses a path to this node (as demonstrated below, it will learn about at least one path when it learns about the new nodes) and adds it to the information it includes in its announcement messages.

CDS Stage: Isolation Non-MIS nodes have a more complicated task as they must decide for which, if any, of their MIS neighbors' instances to participate. Note that a given non-MIS node u has knowledge of the full coin sequence of each MIS neighbor it knows about. Node u considers a given instance of a given MIS neighbor v to be *locally isolated* if and only if v has a 1 in the corresponding bit for its coin sequence, and all overlapping instances that u knows about have a 0 in the corresponding bit. Node u participates only in locally isolated instances. Similarly, we say an instance for v is *globally isolated* if no overlapping instance from an MIS node within 3 hops of v has a 1 bit in the corresponding bit of its coin sequence.

CDS Stage: Non-MIS Node Participation When a non-MIS node u decides to participate in a given instance i of an MIS neighbor v, it proceeds through the following four phases. These phases are described below and in the pseudo-code presented in Algorithm 7.

During the ***announce phase***, if u is on a path selected by an MIS node (i.e., u is in the CDS), then it broadcasts this information w.c.p. on channel $\mathcal{C}$. Otherwise, it listens on channel $\mathcal{C}$. If it hears a message from MIS node v, then this message will include:

(a) all the MIS nodes that are within 3-hops of v that v has discovered; and

(b) paths to each of these nearby MIS nodes.

If u learns from this message that has been selected on one such path, it joins the CDS, stores the path, and it will add it to the list of paths it announces in subsequent announce phase rounds. Otherwise, it simply updates its local snapshot of v's knowledge of nearby MIS nodes. (It maintains these snapshots so that in later phases, if it learns about a nearby MIS node, it knows whether or not v has already heard about this nearby node.)

During the ***sparse search phase***, neighbors of v will attempt to announce v to nearby nodes. To do so, u will broadcast v's identity on channel $\mathcal{C}$ with probability $2^{-(i \mod \log\log n)}$, where i is the instance number.

The ***dense search phase*** has neighbors of v try another tactic for announcing v. Their behavior decides on the size of $\mathcal{F}$. In this phase, if $n_{\mathcal{C}} = \log n$, u chooses a channel $\mathcal{C}_c$ with probability 2^{-c} (i.e., as in the herald protocol from the MIS algorithm), and then broadcasts w.c.p. the *ID* of v. If $n_{\mathcal{C}}$ is smaller, then nodes cycle through $\lceil \frac{\log n}{\mathcal{F}} \rceil$ subsets of these probabilities, selecting probabilities based on the instance number.

The ***report phase*** allows a neighbor u of MIS node v to try to report information about nearby MIS nodes u has recently heard of and that are not included in the most recent announcement messages u has received from v. This phase requires 3 rounds. If u has new information to report to MIS neighbor v, then u will spend the next 3 rounds attempting to report to v. Notice, when u learns about a new MIS node, it has also learned a path to this MIS node (i.e., if u hears from some neighbor w of a new MIS node, then the path is u to w to the new node; if it hears directly from the new MIS node, then the path consists only of u itself), Assume u attempts to report to v. During the first round, u broadcasts w.c.p. its new knowledge on channel $\mathcal{C}$. (This handles the case where only a small number of v's neighbors have information to report.) In the next round, as in the dense search phase, it chooses channel $\mathcal{C}_c$ with probability 2^{-c} and then broadcasts w.c.p. its new information. If u decides to listen in the second round and it receives a message it then acts as a herald for this message in the third round, rebroadcasting it w.c.p. on channel $\mathcal{C}$. This case handles larger collections of nodes to report to v. Again, if there are fewer channels, then we can cycle through the possible subsets of probabilities.

CDS Stage: Non-MIS Nodes Not Participating. As mentioned, in rounds in which a non-MIS node u is not participating in an instance, it either listens on channel $\mathcal{C}$, or chooses a channel from $\mathcal{C}_1, \ldots, \mathcal{C}_{n_{\mathcal{C}}}$ with uniform probability. If u hears a search phase message during these rounds, it will add the corresponding MIS node to its set of nodes it has discovered, and label it with a path. If u hears an announce message from a non-MIS node, it only cares if that node is announcing a path that includes u, at which point u would add this path to its list of paths it is on, and joins the CDS, if it has not already.

Analysis

We now analyze the above algorithm. We begin by proving a property regarding global isolation that follows in a straightforward manner from our bounded independence assumption (which tells us there are $O(1)$ MIS nodes in any constant-hop neighborhood).

Lemma 9.1. *Fix an instance i of MIS node v. W.c.p., instance i is globally isolated.*

We leverage this property to prove the efficiency of the announce phase.

Lemma 9.2. *Fix some round r and node v such that v is either an MIS node or v knows it is in a CDS path. It follows that, w.h.p., every neighbor of v will receive at least one announce phase message from v in the interval $[r, r + \Theta(\log n)]$.*

Proof. Node v will succeed w.c.p. every time it participates in an instance. To see why, notice that there are at most a constant number of other MIS nodes and nodes on a CDS path within 2 hops of v (by our bounded independence assumption). Node v participates w.c.p. every constant number of rounds (by Lemma 9.1). Therefore, $\Theta(\log n)$ rounds are enough for it to succeed with high probability. □

To analyze the search and report phases, the following notation will prove useful. Let M be the set of MIS nodes and $C = V \setminus M$ be the covered (i.e., non-MIS) nodes. For a given covered node $c_i \in C$, let $S(c_i) \subseteq M$ be the one or more MIS nodes that neighbor c_i. Conversely, for MIS node $m_i \in M$, let $W(m_i) \subseteq C$ be the non-MIS nodes neighboring m_i. For any two nodes u, v, let $d(u, v)$ be the shortest distance between u and v in hops. Notice, if an instance of m_i is globally isolated, no node within range of $\{m_i\} \cup W(m_i)$ is participating.

The following lemma is an implication of Lemma 5.2 of our study of leader election in single hop multichannel networks [11] (the original lemma actually proves something stronger). It will prove useful in analyzing the dense search phase.

Lemma 9.3 (Follows from [11])**.** *Assume each node in a non-empty set A of nodes independently chooses a channel from among $\log n$ channels using an exponential probability distribution. W.c.p., there is at least one channel with at least one and no more than constant number of nodes.*

We now show that the sparse and dense search phase combine to ensure that covered nodes efficiently discover nearby supervisors.

Lemma 9.4. *Let $s_i, s_j \in M, i \neq j$, be two MIS nodes such that $d(s_i, s_j) \leq 3$. Within $O\left(\frac{\log^2 n}{\mathcal{F}}\right) + \tilde{O}(\log n)$ rounds at least one of the following two conditions hold:*
(1) some $w_i \in W(s_i)$ has s_j in its discovered set; or
(2) some $w_j \in W(s_j)$ has s_i in its discovered set.

Proof. We know $d(s_i, s_j) > 1$, by the independence property of the MIS. If $d(s_i, s_j) = 2$, then by definition $W(s_i) \cap W(s_j) \neq \emptyset$, so there is a covered node that *begins* the algorithm with knowledge of both s_i and s_j—satisfying the lemma.

Assume, therefore, that $d(s_i, s_j) = 3$. Every 3-hop path between these MIS nodes must go from s_i to a node in $W(s_i)$ to a node in $W(s_j)$ to s_j. Let $B_i \subseteq W(s_i)$ be the nodes covered by s_i that neighbor nodes in $W(s_j)$, and let $B_j \subseteq W(s_j)$ be the nodes covered by s_j that neighbor nodes in $W(s_i)$ (here the B indicates that these nodes *bridge* the worker sets). In the two search phases, workers (i.e., nodes covered by MIS nodes) attempt to propagate the *ID*s of their MIS nodes to their neighbors. To satisfy the lemma, therefore, it is sufficient for a message from B_i to reach B_j, or vice versa.

We consider two cases. In the first case, some node $u \in B_j$ has less than $\log n$ neighbors in B_i. Let N_u be the set of these neighbors. We turn our attention to the sparse search phase. In particular, consider the next instance of s_i in which during the sparse search phase the probability $\hat{p} \in \left[\frac{1}{k}, \frac{2}{k}\right]$, for $k = |N_u|$, is used. W.c.p. this instance is globally isolated. If it is globally isolated, the nodes in $|N_u|$ broadcast with probability $\hat{p}$ during this phase. By the definition of global isolation, u is not participating during this round, so it will receive on channel $\mathcal{C}$ with probability $\frac{1}{2}$. It follows that u receives a message from B_i with probability p_{solo}, bounded as:

$$p_{solo} = \frac{1}{2} \sum^{k} \hat{p} \prod^{k-1} (1 - \hat{p}) \geq \frac{1}{2} k \cdot \frac{1}{k} \left(1 - \frac{2}{k}\right)^{k} \geq \frac{1}{2 \cdot 4^2} = \Theta(1).$$

Therefore, w.c.p., information moves from B_i to B_j every $\log \log n$ phases (and thus, every $O(\log \log n)$ rounds). It is straightforward to therefore conclude that $O(\log n \log \log n)$ rounds are enough to ensure this occurs with high probability.

The second case assumes that *every* node in B_j has at least $\log n$ neighbors B_i. Here we turn our attention to the dense search phase. Consider an instance of s_j that is globally isolated. Consider the dense search phase of this instance. If $n_{\mathcal{C}} = \log n$), it follows from applying Lemma 9.3 with $A = B_j$, that w.c.p. there is one channel $\mathcal{C}_c$ on which a single node u from B_j broadcasts. Let N_u be u's neighbors in B_i. We know that the nodes in N_u are not participating in this round. We also know that $|N_u| \geq \log n$. In expectation, half of these nodes decide to choose a channel from $\mathcal{C}_1, \ldots, \mathcal{C}_{n_{\mathcal{C}}}$ at random. Therefore, w.c.p., we have at least one node on $\mathcal{C}_c$.

We have shown, therefore, that in this case, information passes from B_j to B_i w.c.p. every constant number of rounds. On the other hand, if $n_{\mathcal{C}} < \log n$, we have to wait up to $\left\lceil \frac{\log n}{n_{\mathcal{C}}} \right\rceil$ instances in expectation until we get to an instance where the dense search phase includes a well-matched channel selection probability for the size of B_j.

In either case, $O\left(\frac{\log^2 n}{\mathcal{F}}\right) + \tilde{O}(\log n)$ rounds are enough to ensure this happens with high probability. $\square$

We now turn our attention to the report phase. Let $u \in M$ be an MIS node. Assume u has at least 1 neighbor with information to report to u. Our goal is to prove that at least one such node succeeds in delivering this information w.c.p. within a constant number of rounds. Because the report phase uses the same algorithm strategy as the herald filter from the MIS algorithm, we can leverage our existing analysis of that filter to simplify the below.

Lemma 9.5. *Fix some MIS node u and round r such that at least one neighbor of u has information to report to u in round r. W.h.p., u will receive a report during a report phase round during the interval $[r, r + \Theta(\log n)]$.*

Proof. Consider the next instance of u. W.c.p., this instance is globally isolated. Assume this is the case. We now turn our attention to the report phase of the instance. Let R be the non-empty set of nodes competing to report to u. If $|R| = 1$, then, w.c.p., the single $v \in R$ will broadcast in the first round of this phase. Because the instance is globally isolated, it will have no contention, and therefore succeed with constant probability.

The second case is when $|R| > 1$. Here, the nodes in R will participate in the basic logic of the herald filter, in that they first choose a channel $\mathcal{C}_c$ with probability $\frac{1}{2^c}$, and then broadcast with constant probability. Because of global isolation, we can consider R to be a *fat* region in the terminology of the MIS algorithm. We can apply the same style of argument as in Lemmas 8.6 and 8.8, to prove that, w.c.p., there will be a single channel with a single broadcaster, and this channel will include a constant number of receivers. These receivers will then go on to successfully report to u w.c.p. in the next round.

It follows that $O(\log n)$ rounds are sufficient to ensure u receives the new information with high probability. □

We are now ready to pull together the pieces and prove our main theorem. In the following, when we say we construct a CDS within t rounds, we mean within t rounds of the last node being activated. A slightly more general statement of this theorem would note that we build a CDS on connected subgraph G' within t rounds of the last node in G' being activated.

Theorem 9.6. *W.h.p., a constant-degree connected dominating set can be constructed in $O\left(\frac{\log^2 n}{\mathcal{F}}\right) + \tilde{O}(\log n)$ rounds.*

Proof. Once all relevant nodes are activated, all nodes have moved on to the CDS stage within $O\left(\frac{\log^2 n}{\mathcal{F}}\right) + \tilde{O}(\log n)$ rounds. Within an additional $O(\log n)$ rounds, all non-MIS nodes have heard from all neighboring MIS nodes (by Lemma 9.2). At this point, fix any pair of MIS nodes v and w that are two hops away. They have at least one common neighbor u. By Lemma 9.5, w.h.p., u will report the existence of v to w (and vice versa) within $O(\log n)$ additional rounds. If v and w are instead three hops away, then Lemma 9.4 combined with Lemma 9.5 provides that, w.h.p., they will learn about each other within $O\left(\frac{\log^2 n}{\mathcal{F}}\right) + \tilde{O}(\log n)$ rounds. Therefore, within $O\left(\frac{\log^2 n}{\mathcal{F}}\right) + \tilde{O}(\log n)$ rounds, all MIS nodes within 3 hops will have learned of each other. A final application of Lemma 9.2 tells us that $O(\log n)$ additional rounds ensures that the connecting paths chosen by these MIS nodes have been propagated to the nodes on the paths. At this point, every MIS node, and a short path between every nearby pair of MIS nodes, has joined the MIS. By our earlier argument (also made in [8]), the result is a constant-degree connected dominating set. If we choose sufficiently large constants for the high probability results above, then the polynomially many union bounds required to combine them leaves us with a probability that is still sufficiently high. $\square$

9.2 Other Problems

In this section, we use our MIS algorithm solution as a building block in solving other problems efficiently in the multichannel environment. Our main technical result is a new algorithm that uses the MIS solution as a subroutine to build a constant-degree connected dominating set (CDS) in $O\big(\frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds. We then use this structure as an overlay to derive solutions to broadcast and leader election that run in $O\big(D + \frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds, and to k-message multi-message broadcast that runs in $O\big(D+k+\frac{\log^2 n}{\mathcal{F}}\big)+\tilde{O}(\log n)$ and $O\big(D+k\log n+\frac{\log^2 n}{\mathcal{F}}\big)+\tilde{O}(\log n)$ rounds for unrestricted and restricted message sizes, respectively.

The problems below use a CDS as an overlay network. To best match the typical assumptions for these problems, we will assume synchronous starts—i.e., the CDS algorithm starts and ends at the same rounds for all nodes. Our algorithms all work without this assumption as well, requiring in this case only that the theorem statements be rewritten to guarantee their running time holds *after the first round in which a complete CDS is constructed.*

Broadcast. First, build a constant-degree CDS using the above algorithm. Then, the source delivers the message to its CDS neighbors. On receiving the message, a CDS node re-broadcasts it with constant probability in each round. Because the CDS nodes have constant degree, a standard Chernoff analysis shows that, w.h.p., the message will reach every CDS node in $O(D + \log n)$ rounds (and therefore every node within $O(\log n)$ more rounds). Combined with the running time of the CDS algorithm, the total running time is $O\big(D + \frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds, nearly reaching the $\Omega\big(D + \frac{\log^2 n}{\mathcal{F}}\big)$ centralized lower bound for the multi-channel setting [17]. Formally:

Theorem 9.7. *W.h.p., the problem of global broadcast can be solved in* $O\big(D + \frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ *rounds.*

Multi-Message Broadcast. The k-message multi-message broadcast problem assumes k messages must be propagated to the entire network. As before, first build a constant-degree CDS. We then use the same logic as the report phase of the CDS algorithm to propagate the k messages from their sources to nearby CDS nodes. This routine uses

$\log n$ channels and, w.h.p., can deliver all k messages to nearby nodes in $O(k + \log n)$ rounds. Once the messages are in the CDS, how we propagate depends on our assumption on message size. For unrestricted message size, we can run the above simple broadcast algorithm, simply combining all messages a node has received into a single message, in each round. This requires $O(D + \log n)$ rounds to propagate all k once we have our CDS. If we assume restricted message size (i.e., $O(\text{polylog}\, n)$ bits), we can use the algorithm and analysis of [22]. As established in [22], this will require $O(D + k \log n)$ rounds (formally, F_{prog} in the relevant theorem is $O(1)$ while F_{ack} is in $O(\log n)$). From this we conclude:

Theorem 9.8. *W.h.p., it is possible to solve k-message multi-message broadcast in time $O\big(D + k + \frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ with unrestricted messages sizes, and in $O\big(D + k \log n + \frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds with restricted message sizes.*

Leader Election. To elect a leader, run the broadcast algorithm with *all* nodes initiating broadcast with a message containing their own ID, and having each node update its broadcast message in each round to include the largest *ID* it has received so far. Using a standard Chernoff analysis, we can show that the largest ID will propagate to all nodes within $O(D + \log n)$ rounds. Formally:

Theorem 9.9. *W.h.p., leader election can be solved in $O\big(D + \frac{\log^2 n}{\mathcal{F}}\big) + \tilde{O}(\log n)$ rounds.*

References

[1] I. 802.11. Wireless LAN MAC and Physical Layer Specifications, June 1999.

[2] N. Alon, L. Babai, and A. Itai. A Fast and Simple Randomized Parallel Algorithm for the Maximal Independent Set Problem. *Journal of Algorithms*, 1986.

[3] N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. A Lower Bound for Radio Broadcast. *Journal of Computer and System Sciences*, 43(2):290–298, 1991.

[4] N. Alon and J. Spencer. *The Probabilistic Method.* Wiley International, 3rd edition, 2008.

[5] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the Time-Complexity of Broadcast in Multi-Hop Radio Networks: An Exponential Gap Between Determinism and Randomization. *J. Computer and System Sciences*, 45(1):104–126, 1992.

[6] L. Barenboim, M. Elkin, S. Pettie, and J. Schneider. The Locality of Distributed Symmetry Breaking. In *Proc. Symp. on Foundations of Computer Science (FOCS)*, pages 321–330, 2012.

[7] Bluetooth Consortium. *Bluetooth Specification Version 2.1*, July 2007.

[8] K. Censor-Hillel, S. Gilbert, F. Kuhn, N. Lynch, and C. Newport. Structuring Unreliable Radio Networks. In *Proc. 30th Symp. on Principles of Distributed Computing (PODC)*, pages 79–88, 2011.

[9] I. Chlamtac and S. Kutten. On Broadcasting in Radio Networks–Problem Analysis and Protocol Design. *IEEE Transactions on Communications*, 33(12):1240–1246, 1985.

[10] S. Daum, M. Ghaffari, S. Gilbert, F. Kuhn, and C. Newport. Maximal independent sets in multichannel radio networks. Technical Report 275, University of Freiburg, Department of Computer Science, 2013.

[11] S. Daum, S. Gilbert, F. Kuhn, and C. Newport. Leader Election in Shared Spectrum Networks. In *Proc. Symp. on Principles of Dist. Computing (PODC)*, pages 215–224, 2012.

[12] S. Daum, F. Kuhn, and C. Newport. Efficient Symmetry Breaking in Multi-Channel Radio Networks. In *Proc. 26th International Symposium on Distributed Computing (DISC)*, pages 238–252, 2012.

[13] S. Dolev, S. Gilbert, R. Guerraoui, D. R. Kowalski, C. Newport, F. Kuhn, and N. Lynch. Reliable Distributed Computing on Unreliable Radio Channels. In *the Proc. of the 2009 MobiHoc S^3 Workshop*, 2009.

[14] S. Dolev, S. Gilbert, R. Guerraoui, F. Kuhn, and C. Newport. The Wireless Synchronization Problem. In *Proceedings of the Principles of Distributed Computing*, pages 190–199, 2009.

[15] S. Dolev, S. Gilbert, R. Guerraoui, and C. Newport. Gossiping in a Multi-Channel Radio Network: An Oblivious Approach to Coping with Malicious Interference. In *Proc. Symp. on Dist. Comp. (DISC)*, 2007.

[16] S. Dolev, S. Gilbert, R. Guerraoui, and C. Newport. Secure Communication Over Radio Channels. In *Proc. Symp. on Principles of Dist. Computing (PODC)*, 2008.

[17] S. Dolev, S. Gilbert, M. Khabbazian, and C. Newport. Leveraging Channel Diversity to Gain Efficiency and Robustness for Wireless Broadcast. In *Proc. Symp. on Dist. Comp. (DISC)*, 2011.

[18] A. Ephremides, J. E. Wieselthier, and D. J. Baker. A Design Concept for Reliable Mobile Radio Networks With Frequency Hopping Signaling. *Proc. of the IEEE*, 75(56–73), 1987.

[19] M. Farach-Colton, R. J. Fernandes, and M. A. Mosteiro. Lower Bounds for Clear Transmissions in Radio Networks. In *Proc. Latin American Symposium on Theoretical Informatics*, 2006.

[20] S. Gilbert, R. Guerraoui, D. Kowalski, and C. Newport. Interference-Resilient Information Exchange. In *Proc. Conf. on Computer Comm. (INFOCOM)*, 2009.

[21] T. Jurdzinski and G. Stachowiak. Probabilistic Algorithms for the Wakeup Problem in Single-Hop Radio Networks. In *Proc. Int. Symp. on Algorithms and Computation*, pages 535–549, 2002.

[22] F. Kuhn, N. Lynch, and C. Newport. The Abstract MAC Layer. In *Proc. Symp. on Dist. Comp. (DISC)*, 2009.

[23] F. Kuhn, T. Moscibroda, T. Nieberg, and R. Wattenhofer. Fast Deterministic Distributed Maximal Independent Set Computation on Growth-Bounded Graphs. In *Proc. 19th International Symposium on Distributed Computing (DISC)*, 2005.

[24] M. Luby. A Simple Parallel Algorithm for the Maximal Independent Set Problem. *SIAM Journal on Computing*, 15(4):1036–1053, 1986.

[25] T. Moscibroda and R. Wattenhofer. Maximal Independent Sets In Radio Networks. In *Proc. 24th Symp. on Principles of Distributed Computing (PODC)*, pages 148–157, 2005.

[26] C. Newport. *Distributed Computation on Unreliable Radio Channels.* PhD thesis, MIT, 2009.

[27] S. Schmid and R. Wattenhofer. Algorithmic Models for Sensor Networks. In *Proc. 14th Int. Workshop on Parallel and Distributed Real-Time Systmes*, pages 1–11, 2006.

[28] J. Schneider and R. Wattenhofer. A Log-Star Distributed Maximal Independent Set Algorithm for Growth-Bounded Graphs. In *Proc. Symp. on Principles of Dist. Computing (PODC)*, pages 35–44, 2008.

[29] M. Strasser, C. Pöpper, and S. Capkun. Efficient Uncoordinated FHSS Anti-jamming Communication. In *Proc. Symp. on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, 2009.

[30] M. Strasser, C. Pöpper, S. Capkun, and M. Cagalj. Jamming-resistant Key Establishment using Uncoordinated Frequency Hopping. In *the Proceedings of the IEEE Symposium on Security and Privacy*, 2008.

A Properties of Graphs with Bounded Independence

For the analysis of our algorithms in the bounded independence graph model, we require the followin general statements about the independence number of graphs. As we are not aware of any reference where they appear, we give full proofs. The first lemma can be seen as a weighted version of Turán's theorem.

Lemma A.1. *Let $G = (V, E)$ be a graph and assume that every node $u \in V$ has a positive edge weight $w_u > 0$. Define $W := \sum_{v \in V} w_v$ and for each $u \in V$, $W_u := \sum_{v \in N_G^+(u)} w_v$. It then holds that*

$$\sum_{v \in V} \frac{w_v}{W_v} \leq \alpha(G) \quad \textit{and} \tag{8}$$

$$\sum_{v \in V} w_v \cdot W_v \geq \frac{W^2}{\alpha(G)}, \tag{9}$$

where $\alpha(G)$ is the independence number of G.

Proof. We adapt the well-known probabilistic proof of Turán's theorem [4]. Given a global order on the node set V, we can construct an independent set S of G as follows. We start with an empty set and consider the vertices in the given order. When considering a node $v \in V$, we add v to S iff no neighbor of v is already added to S. We randomly order the vertices as follows. The first node in the order is chosen randomly with probability proportional to the weights w_v, i.e., node v is chosen with probability $\frac{w_v}{W}$. All subsequent vertices in the random order are again chosen randomly from the remaining set of vertices with probability proportional to w_v. Let S be the independent set obtained by this random process.

A sufficient condition for node v to be in S is that v appears in the constructed random order before any of its neighbors. The probability for this to happen is exactly $\frac{w_v}{W_v}$. The expected size of S can therefore be estimated as follows:

$$\mathbb{E}[|S|] \overset{(*)}{\geq} \sum_{v \in V} \frac{w_v}{W_v} = W \cdot \sum_{v \in V} \frac{\frac{w_v}{W}}{W_v} \overset{(**)}{\geq} W \cdot \frac{1}{\sum_{v \in V} \frac{w_v}{W} \cdot W_v} = \frac{W^2}{\sum_{v \in V} w_v \cdot W_v}. \tag{10}$$

As there exists an independent set of size at least $\mathbb{E}[|S|]$, we clearly have $\alpha(G) \geq \mathbb{E}[|S|]$ and thus (*) proves (8). The inequality (**) follows from the convexity of the function $\frac{1}{x}$. That, together with $\alpha(G) \geq \mathbb{E}[|S|]$, proves (9). $\square$

Note that by choosing $w_v = 1$ for all $v \in V$ and if we denote the average degree of G by $\bar{d}$, the statement of Lemma A.1 simplifies to

$$\sum_{v \in V} \frac{1}{\deg(v)} \leq \alpha(G) \quad \text{and} \quad \bar{d} + 1 \geq \frac{n}{\alpha(G)},$$

which is equivalent to Turán's theorem.

As a consequence of Lemma A.1, we obtain the following lemma.

Lemma A.2. *Let $G = (V, E)$ be a graph and assume that every node $u \in V$ has a positive edge weight $w_u > 0$. Define $W := \sum_{v \in V} w_v$ and for each $u \in V$, $W_u := \sum_{v \in N_G^+(u)} w_v$. Let $V_{heavy} \subseteq V$ be the set of nodes v for which $W_v \geq \frac{W}{2\alpha(G)}$. The total weight of node in V_{heavy} is at least*

$$\sum_{v \in V_{heavy}} w_v > \frac{W}{2\alpha(G)}.$$

Proof. We define $V_{light} := V \setminus V_{heavy}$. Note that therefore for all $v \in V_{light}$, $W_v < \frac{W}{2\alpha(G)}$. Summing up $w_v W_v$ for nodes in V_{light} thus yields

$$\sum_{v \in V_{light}} w_v \cdot W_v < \sum_{v \in V_{light}} w_v \cdot \frac{W}{2\alpha(G)} \leq \frac{W^2}{2\alpha(G)}. \tag{11}$$

For nodes in V_{heavy}, we then get

$$\begin{aligned}
\sum_{v \in V_{heavy}} w_v \quad &\geq \quad \frac{1}{W} \cdot \sum_{v \in V_{heavy}} w_v W_v \\
&= \quad \frac{1}{W} \cdot \left(\sum_{v \in V} w_v W_v - \sum_{v \in V_{light}} w_v W_v \right) \\
&\overset{\text{Eq. (11)}}{>} \quad \frac{1}{W} \cdot \left(\sum_{v \in V} w_v W_v - \frac{W^2}{2\alpha(G)} \right) \\
&\overset{\text{Lem. A.1}}{\geq} \quad \frac{W}{2\alpha(G)},
\end{aligned}$$

which concludes the proof. □

By choosing $w_v = 1$ for all nodes, we directly obtain the following corollary.

Corollary A.3. *More than $\frac{n}{2\alpha(G)}$ nodes of an n-node graph G with independence number $\alpha(G)$ have degree at least $\frac{n}{2\alpha(G)} + 1$.* □

B Low Number of Channels

In this section we shortly describe how to adapt the algorithm to work with $o(\log \log n)$ channels. The decay filter stays unchanged (low number of channels imply long phases), the only problems arise in the herald filter, as we need $\Theta(\log \log n)$ herald election channels and loneliness support channels. Let $m := \max\{n_{\mathcal{A}}, n_{\mathcal{S}}\} = \Theta(\log\log n)$. We will use $\frac{\mathcal{F}}{4}$ channels for each of those two groups of channels and the remaining $\frac{\mathcal{F}}{2}$ channels to cover the other channels needed. We define $c := 2^{\frac{4m}{\mathcal{F}}}$ and use c as the base for the exponential distributions, i.e., for herald election nodes choose channel $\mathcal{A}_i$ with probability c^{-i} and in the loneliness support block nodes broadcast with probability c^{-i} if they choose channel $\mathcal{S}_i$. The following lemmas have to be adjusted. Lemma 8.8 no longer provides a pair with probability $\Omega(\pi_\ell)$, but with probability $\Omega\big(\frac{\mathcal{F}}{m} \cdot \pi_\ell\big)$. Lemma 8.16 now guarantees the creation of an MIS node within $O\big(\frac{m}{\mathcal{F}} \cdot \log n\big)$ rounds instead of $O(\log n)$ rounds. The chain of reasoning in all proofs remains, only the calculations have to be adjusted.