

1 Information Gathering in Ad-Hoc Radio Networks

2 **Marek Chrobak**

3 Department of Computer Science
4 University of California at Riverside

5 **Kevin P. Costello**

6 Department of Mathematics
7 University of California at Riverside

8 **Leszek Gąsieniec**

9 Department of Computer Science
10 University of Liverpool

11 — Abstract —

12 In the ad-hoc radio network model, nodes communicate with their neighbors via radio signals,
13 without knowing the topology of the underlying digraph. We study the information gathering
14 problem, where each node has a piece of information called a *rumor*, and the objective is to
15 transmit all rumors to the designated target node. For the model without any collision detection
16 we provide an $\tilde{O}(n^{1.5})$ deterministic protocol, significantly improving the trivial bound of $O(n^2)$.
17 We also consider a model with a mild form of collision detection, where a node receives a 1-bit
18 acknowledgement if its transmission was received by at least one out-neighbor. For this model
19 we give an $\tilde{O}(n)$ deterministic protocol for information gathering in acyclic graphs.

20 **2012 ACM Subject Classification** Discrete Mathematics → Combinatorics • Combinatorial Op-
21 timization • Theory of Computation → Design and Analysis of Algorithms → Distributed Al-
22 gorithms • Networks → Ad-Hoc Networks

23 **Keywords and phrases** algorithms, radio networks, information dissemination

24 **Digital Object Identifier** 10.4230/LIPIcs...

25 **Funding** M. Chrobak’s research supported by NSF grant CCF-1536026.

26 **1 Introduction**

27 We address the problem of information gathering in ad-hoc radio networks. A *radio network*
28 is represented by a directed graph (digraph) G , whose nodes represent radio transmitter-
29 s/receivers and directed edges represent their transmission ranges; that is, an edge (u, v) is
30 present in the digraph if and only if node v is within the range of node u . When a node
31 u transmits a message, this message is immediately sent out to all its out-neighbors. How-
32 ever, a message may be prevented from reaching some out-neighbors of u if it collides with
33 messages from other nodes. A collision occurs at a node v if two or more in-neighbors of v
34 transmit at the same time, in which case v will not receive any of their messages, and it will
35 not even know that they transmitted.

36 Radio networks, as defined above, constitute a useful abstract model for studying proto-
37 cols for information dissemination in networks where communication is achieved via broad-
38 cast channels, as opposed to one-to-one links. Such networks do not need to necessarily
39 utilize radio technology; for example, in local area networks based on the ethernet protocol
40 all nodes communicate by broadcasting information through a shared carrier. Different
41 variants of this model have been considered in the literature, depending on the assumptions
42 about the node labels (that is, identifiers), on the knowledge of the underlying topology,



43 and on allowed message size. In this work we assume that nodes are labelled $0, 1, \dots, n - 1$,
 44 where n is the network size. (All our results remain valid if the labels are selected from
 45 the range $0, 1, \dots, O(n)$). We focus on the *ad-hoc model*, where the digraph's topology is
 46 unknown when the computation starts, and a protocol needs to complete its task within a
 47 desired time bound, no matter what the topology is. At the beginning of the computation
 48 each node v is in possession of a unique piece of information, that we refer to as a *rumor*.
 49 Different communication primitives are defined by specifying how these rumors need to be
 50 disseminated across the network. In this paper we do not make any assumptions about the
 51 size of transmitted messages; thus a node can aggregate multiple rumors and transmit them
 52 in one message. In fact, it could as well transmit as one message the complete history of its
 53 past computation.

54 The two most studied information dissemination primitives for this model are broad-
 55 casting and gossiping. In *broadcasting* (or *one-to-all communication*), a single source node s
 56 attempts to deliver its rumor to all nodes in the network. For broadcasting to be mean-
 57 ingsful, we need to assume that all nodes in G are reachable from s . In *gossiping* (or *all-to-all*
 58 *communication*), the objective is to distribute all rumors to all nodes in the network, under
 59 the assumption that G is strongly connected. Both these primitives can be solved in time
 60 $O(n^2)$ by a simple protocol called ROUNDROBIN, where all nodes transmit cyclically one
 61 at a time (see Section 2). Past research on ad-hoc radio networks focussed on designing
 62 protocols that improve this trivial bound.

63 For broadcasting, gradual improvements in the running time have been reported since
 64 early 2000's [6, 21, 2, 3, 12, 13, 11], culminating in the upper bound of $O(n \log D \log \log(D\Delta/n))$
 65 in [10], where D denotes the diameter of G and Δ its maximum in-degree. This is already
 66 almost tight, as the lower bound of $\Omega(n \log D)$ is known [9]. For randomized algorithms, the
 67 gap between lower and upper bounds is also almost completely closed, see [1, 22, 11].

68 In case of gossiping, major open problems remain. The upper bound of $O(n^2)$ was
 69 improved to $\tilde{O}(n^{1.5})$ in [6, 29] and then later to $\tilde{O}(n^{4/3})$ in [18], and no better bound is
 70 currently known¹. No lower bound better than $\Omega(n \log n)$ (that follows from [9]) is known.
 71 In contrast, in the randomized case it is possible to achieve gossiping in time $\tilde{O}(n)$ [11, 23, 7].

72 The reader is referred to survey papers [15, 20, 16, 27, 19] that contain more information
 73 about information dissemination protocols in different variants of radio networks.

74 In this paper we address the problem of *information gathering* (that is, *all-to-one com-*
 75 *munication*). In this problem, similar to gossiping, each node v has its own rumor, and the
 76 objective is to deliver these rumors to a designated target node t . (We assume that t is
 77 reachable from all nodes in G .)

78 The problem of information gathering for trees was introduced in [5], where an $O(n)$ -
 79 time algorithm was presented. Other results in [5] include algorithms for the model without
 80 rumor aggregation or the model with transmission acknowledgements.

81 **Our results.** Our main result, in Section 4, is a deterministic protocol that solves the
 82 information gathering problem in arbitrary ad-hoc networks in time $\tilde{O}(n^{1.5})$. To our know-
 83 ledge this is the first protocol for this problem that achieves running time faster than the
 84 trivial $O(n^2)$ bound. One of our key technical contributions is in solving this problem in
 85 time $\tilde{O}(n^{1.5})$ for acyclic graphs (Section 3). Previous protocols developed for gossiping on
 86 strongly connected graphs rely on feedback (see the discussion below), and are not applic-
 87 able to this problem. This algorithm for acyclic graphs is based on careful application of

¹ We use notation $\tilde{O}(f(n))$ to conceal poly-logarithmic factors; that is, $g(n) = \tilde{O}(f(n))$ iff $g(n) = O(f(n) \log^c n)$ for some constant c . Also, we write $f(n) = \tilde{\Omega}(g(n))$ if and only if $g(n) = \tilde{O}(f(n))$.

88 combinatorial structures called strong selectors, combined with a novel amortization tech-
89 nique to measure progress of the algorithm. To extend this protocol to arbitrary graphs,
90 we integrate it with a gossiping protocol. Roughly, the two sub-protocols run intertwined in
91 parallel, with the sub-protocol for acyclic graphs transferring rumors between strongly con-
92 nected components and the gossiping sub-protocol disseminating them within each strongly
93 connected component. This requires overcoming two challenges. One is that the partition
94 of G into strongly connected components is not actually known, so the combined protocol
95 needs to gradually “learn” the connectivity structure of G while it executes. The second
96 challenge is in synchronizing the computation of the two sub-protocols, since they are based
97 on entirely different principles.

98 In the second part of the paper, in Section 5, we take the “dual” approach to investigate
99 the time complexity of information gathering: rather than optimizing the running time
100 in the general case, we examine what assumptions on the model would permit achieving
101 running time $\tilde{O}(n)$. To this end, we consider a relaxation of our model by allowing a mild
102 form of collision detection. In this new model each node v , after each transmission, receives
103 a 1-bit acknowledgement indicating whether its transmission was received by at least one
104 out-neighbor. With this assumption, we provide an $\tilde{O}(n)$ -time algorithm for information
105 gathering in acyclic radio networks.

106 **Additional context and motivations.** If G is strongly connected then information gath-
107 ering and gossiping are equivalent. (This also naturally applies to connected *undirected*
108 graphs.) Trivially, a gossiping algorithm gathers all rumors in t , solving the information
109 gathering problem. On the other hand, one can solve the gossiping problem by running
110 an information gathering protocol followed by any $\tilde{O}(n)$ -time broadcasting protocol with
111 source node t . However, unlike gossiping, the information gathering problem applies to a
112 wider class of digraphs, namely all digraphs where the target node is reachable from all
113 nodes.

114 This weakening of connectivity assumptions introduces new challenges for information
115 gathering. The crucial one is *lack of feedback*, namely that the nodes in the network do not
116 receive any information about the fate of their transmissions. This should be contrasted
117 with the gossiping problem where the nodes can take advantage of strong connectivity to
118 eventually learn whether their earlier transmissions were successful. In fact, the existing
119 protocols for gossiping critically rely on this feature, as they use it to identify nodes that
120 have collected a large number of rumors, and subsequently broadcast these rumors to the
121 whole network, thus removing them from consideration and reducing congestion.

122 Some evidence that feedback might help to speed up information gathering can be found
123 in [4], where the authors studied the model in which rumor aggregation is not allowed.
124 In this model they developed an $O(n)$ -time protocol for trees, under the assumption that
125 nodes receive immediate acknowledgements of successful transmissions. In contrast, without
126 feedback the best known upper bound (also from [4]) for this problem is $O(n \log \log n)$.

127 Various forms of feedback have been studied in the past in the context of *contention res-*
128 *olution* for multiple-access channels (MAC), where nodes communicate via a single shared
129 channel. (Ethernet is one example.) Depending on more specific characteristics of this
130 shared channel, one can model this problem as the information gathering problem either on
131 a complete graph or a star graph, which is a collection of n nodes connected by directed edges
132 to the target node t . (See [24, 26, 25, 14] for information about contention resolution proto-
133 cols.) For instance, in [5] a tight bound of $\Theta(n \log n)$ was given for randomized information
134 gathering on star graphs (or MACs) even if the nodes have no labels (are indistinguishable)
135 and receive no feedback.

136 As explained earlier, in our model rumor aggregation is allowed. This capability is needed
 137 to beat the $O(n^2)$ upper bound, as without rumor aggregation it is quite easy to show a
 138 lower bound of $\Omega(n^2)$ for both gossiping and information gathering, even for randomized
 139 algorithms and with the topology known [17].

140 Interestingly, the randomized gossiping algorithms in [7, 23] can be adapted to inform-
 141 ation gathering, retaining their $\tilde{O}(n)$ expected running time. Thus randomization can help
 142 not only to overcome collisions, but also lack of feedback².

143 2 Preliminaries

144 **Graph terminology.** Throughout the paper, we assume that the radio network is rep-
 145 resented by a digraph (directed graph) $G = (V, E)$ with a distinguished target node t that
 146 is reachable from all other nodes. By $n = |V|$ we denote the number of nodes in G . If
 147 $(u, v) \in E$ then we refer to u as the *in-neighbor* of v and to v as the *out-neighbor* of u . For
 148 any node v , by $N^-(v) = \{u \in V : (u, v) \in E\}$ we denote the set of its in-neighbors.

149 For brevity, we will refer to strongly connected components of G as *sc-components*. For
 150 each node v , the sc-component containing v will be denoted by $C(v)$. We partition the
 151 set of in-neighbors of v into those that belong to $C(v)$ and those that do not: $N_{\text{scC}}^-(v) =$
 152 $N^-(v) \cap C(v)$ and $N_{\text{ACY}}^-(v) = N^-(v) \setminus C(v)$.

153 The *ancestor set* of v in G , denoted $\text{Anc}[v]$, is the set of all nodes of G from which
 154 v is reachable (via a directed path). Note that $C(v) \subseteq \text{Anc}[v]$. In fact, if A is an sc-
 155 component then all vertices in A have the same ancestor set. It will be thus convenient to
 156 extend this definition to sc-components of G ; if A is an sc-component then its ancestor set
 157 is defined as $\text{Anc}[A] = \text{Anc}[v]$, for some arbitrary $v \in A$. The *proper ancestor set* of A is
 158 $\text{Anc}(A) = \text{Anc}[A] \setminus A$.

159 **Radio networks.** We now give the description of the radio network model that our results
 160 apply to. We first provide the standard definition, as used in the earlier literature. Later
 161 in this section we will show that some restrictions of this model can be relaxed, in order to
 162 simplify the algorithms and proofs.

163 As mentioned in the introduction we assume that each node of G has a unique label from
 164 the set $[n] = \{0, 1, \dots, n - 1\}$. For convenience, we will identify nodes with their labels, so a
 165 “node u ” really means the node with label u . We assume that n is known. All our protocols
 166 still work correctly within the claimed time bounds if the label set is $[N]$ for $N = O(n)$,
 167 provided that N is known upfront (but not necessarily n).

168 The time is divided into discrete *time steps* numbered with non-negative integers. We
 169 assume that all nodes start to execute the protocol simultaneously at time step 0. At each
 170 step each node can be either in the *transmitting state*, when it can only transmit, or in
 171 the *receiving state*, when it can only listen to transmissions from other nodes. Only one
 172 message can be transmitted at each step. This is not an essential restriction because, as
 173 already mentioned, we are not imposing any restrictions on the size or format of messages
 174 transmitted by nodes.

175 If a node u transmits a message at a time τ , this message is sent to all out-neighbors
 176 of u in the same step. If v is one of these out-neighbors then v will receive this message

² We should point out, however, a subtle difference. In the randomized case of information gathering, while gathering all rumors in the target node will indeed complete in expected time $\tilde{O}(n)$, the nodes in the network have no way to determine whether the process has completed, except for simply just waiting for $O(n^2)$ steps. In case of gossiping, the expected running time of $\tilde{O}(n)$ includes broadcasting the confirmation of the process' completion to the whole network.

177 at time τ provided that v is in the receiving state and that u is the only in-neighbor of v
 178 that transmits at time τ . If there are two or more in-neighbors of v that transmit at time τ
 179 then a *collision occurs*, and v does not receive any information (independently of its current
 180 state). In other words, collisions are indistinguishable from absence of transmissions. There
 181 is no feedback mechanism available in this model, that is a sender of a message does not
 182 receive any information as to whether its transmission was successful or not. (We will relax
 183 this restriction later in Section 5.)

184 **Selectors.** A *strong* (n, k) -*selector* is a sequence of label sets $S = (S_0, S_1, \dots, S_{\ell-1})$ (that
 185 is, $S_i \subseteq [n]$ for each i) that “singles out” each label from each subset with at most k labels,
 186 in the following sense: for each $X \subseteq [n]$ with $|X| \leq k$ and each $x \in X$ there is an index
 187 i such that $S_i \cap X = \{x\}$. It is known [8] that there exist strong (n, k) -selectors of size
 188 $\ell = O(k^2 \log n)$.

189 Such selectors are often used in protocols for ad-hoc radio networks for reducing collisions.
 190 In a protocol based on a strong (n, k) -selector S each time step is associated with some set
 191 S_i and only the nodes in S_i are allowed to transmit at this time step. To illustrate more
 192 concretely how such selectors are used in our work, consider a node v of in-degree at most
 193 k , and suppose that its in-neighbors are initially dormant, staying in the receiving state,
 194 and they activate at certain (possibly different) times. Each in-neighbor w of v , after it
 195 activates, transmits its message at a time step τ if and only if $w \in S_{\tau \bmod \ell}$. (Observe
 196 that, importantly, at any time τ all nodes use the *same* transmission set $S_{\tau \bmod \ell}$.) Let X
 197 be the set of in-neighbors of v and suppose that some $w \in X$ activates at time η . Then
 198 the definition of strong (n, k) -selectors implies that there will be $\tau \in [\eta, \eta + \ell)$ such that
 199 $S_{\tau \bmod \ell} \cap X = \{w\}$. In other words, w will be the only in-neighbor of v that transmits at
 200 time step τ . So if v stays dormant until time $\eta + \ell$ then v will receive w 's message after at
 201 most $\ell = O(k^2 \log n)$ time steps since the activation time of w . (In fact, it is sufficient if v
 202 happens to be in the receiving state at time τ .) We stress that this is true independently of
 203 the label assignment and of activation times of the nodes other than w .

204 For all $j = 0, 1, \dots, \frac{1}{2} \log n$, by 2^j -SELECT = $(S_0^j, S_1^j, \dots, S_{\ell_j-1}^j)$ we will denote a strong
 205 $(n, 2^j)$ -selector of size $\ell_j = O(4^j \log n)$. Without loss of generality we can assume that $\ell_{j+1} =$
 206 $4\ell_j$ for all $j \leq \frac{1}{2} \log n - 1$. We remark that in Section 5, where we consider transmissions with
 207 acknowledgements, it will be desirable to have many (but not necessarily all) of a collection
 208 of competing in-neighbors of a node transmit successfully. For this purpose we will there
 209 introduce a different type of selectors.

210 Another basic protocol that is often used is called ROUNDROBIN. In this protocol all
 211 nodes transmit cyclically one by one; that is each node w transmits in a step τ if and only if
 212 $w = \tau \bmod n$. In ROUNDROBIN there are no collisions, so, in the setting above, node w will
 213 successfully transmit its message to v in at most n time steps. Observe that a protocol based
 214 on a strong (n, k) -selector can be faster than ROUNDROBIN only when $k = O(\sqrt{n/\log n})$.

215 *Note:* To avoid clutter, in the definitions above, as well as later throughout the paper,
 216 we omit the notation for rounding and assume that in all formulas representing integer
 217 quantities (the number of nodes, steps, etc.) their values are appropriately rounded. This
 218 will not affect asymptotic running time estimates.


219 **Simplifying assumptions.** To streamline the presentation of our protocols, in the paper
 220 we use a relaxed communication model with two additional features:

221 (MFC) We assume that some number κ of radio *frequency channels*, numbered $0, 1, \dots, \kappa - 1$,
 222 is available for communication. So a node may receive and transmit up to κ messages at
 223 each step, one per channel. There is no interference between channels; that is, a message

224 sent on one channel cannot collide with a message sent on a different channel. For each
 225 individual frequency the standard collision rule applies: if two in-neighbors of a node
 226 transmit on this frequency channel at the same time step, then a collision occurs and
 227 this node does not receive any information on this channel at this step.
 228 (SRT) Further, for each frequency f , a node can receive and transmit at frequency f in
 229 a single step. The restriction is that the messages transmitted at all frequencies in any
 230 step do not depend on the messages received in this step.

231 All protocols in this paper are presented in terms of this relaxed model. Below we
 232 explain, however, how to convert these protocols to the standard radio network model (as
 233 defined earlier), increasing the running time only by factor $O(\kappa)$; that is, a protocol that
 234 uses features (MFC) and (SRT), and for which we give an upper bound $O(T)$ on the running
 235 time, can be converted into a protocol in the standard model whose running time is $O(\kappa T)$.
 236 Since $\kappa = O(\log n)$ in our protocols, their $\tilde{O}(\cdot)$ -complexity is not affected.

237 *Simulating multiple frequencies.* We first explain how we can convert *any* protocol \mathcal{A} that
 238 uses κ frequencies and runs in time $O(T)$ into a protocol \mathcal{A}' that uses only one frequency
 239 and runs in time $O(\kappa T)$. This can be done by straightforward time multiplexing. In more
 240 detail: \mathcal{A}' organizes all time steps $0, 1, 2, \dots$ into *rounds*. Each round $r = 0, 1, 2, \dots$ consists
 241 of κ consecutive steps $r\kappa, r\kappa + 1, \dots, r\kappa + \kappa - 1$. Each step s of \mathcal{A} is simulated by round s of
 242 \mathcal{A}' . For each frequency f , the message transmitted at frequency f by \mathcal{A} is transmitted by
 243 \mathcal{A}' in step $s\kappa + f$, that is the f th step of round s . At the end of round s , \mathcal{A}' will know all
 244 messages received in this round, so it will know what messages would \mathcal{A} receive in step s ,
 245 and therefore it knows the state of \mathcal{A} and can determine the transmissions of \mathcal{A} in the next
 246 step.



■ **Figure 1** Partition of \mathcal{A}' 's time steps into rounds, for $\kappa = 10$ frequencies.

247 *Simulating simultaneous receiving/transmitting.* By the argument above, we can assume
 248 that we have only one frequency channel. We first give a generic argument that applies to
 249 arbitrary protocols. (The construction in this paragraph is not strictly needed for our results,
 250 but we include it, as it has some independent interest and provides useful context.) For an
 251 arbitrary protocol we claim that we can disallow simultaneous receiving and transmitting
 252 at the cost of only adding a logarithmic factor to the running time. To see this, suppose
 253 that \mathcal{B} is some transmission protocol where nodes can transmit and listen at the same time.
 254 (Recall that the transmission of \mathcal{B} at any step does not depend on the information it receives
 255 in the same step.) We use a strong $(n, 2)$ -selector $2\text{-SELECT} = (S_i^1)_i$ of size $\ell_1 = O(\log n)$.
 256 We replace each step τ of \mathcal{B} by a time segment I_τ of length ℓ_1 . For any node w and any
 257 $i = 0, 1, \dots, \ell_1 - 1$, if $w \in S_i^1$ then at the i th step of segment I_τ node w transmits whatever
 258 message it would transmit in \mathcal{B} at time τ ; otherwise w is in the receiving state. By definition,
 259 in this new protocol \mathcal{B}' nodes do not transmit and receive at the same time. Further, for
 260 any edge (u, v) , let i be such that $S_i^1 \cap \{u, v\} = \{u\}$. If u transmitted successfully to v in
 261 step τ of \mathcal{B} , then in \mathcal{B}' in the i -th step of segment I_τ u will be in the transmitting state and
 262 v will be in the receiving state, guaranteeing that u 's message will reach v .

263 We now claim that for the type of protocols presented in the paper this additional factor
 264 of $\log n$ is not necessary; that is, allowing simultaneous reception and transmission does
 265 not affect their asymptotic running time at all. Our protocols ARBGATHER in Section 4
 266 and ACYGATHERACK in Section 5 are *selector-based*, namely the computation of each node
 267 is divided into time intervals, where in each interval the node transmits either according to
 268 ROUNDROBIN or according to some strong selector, in the manner described earlier in this
 269 section. In case of ROUNDROBIN, the simultaneous reception and transmission capability is
 270 (trivially) not needed. For intervals where a strong selector is used, the argument how this
 271 capability can be removed was given in [4]. Roughly, the idea is that whenever a protocol
 272 uses a strong (n, k) -selector, this selector can be replaced by a strong $(n, k + 1)$ -selector.
 273 The size of this $(n, k + 1)$ -selector is $O(k^2 \log n)$, so for $k > 1$ it is asymptotically the same
 274 as for a strong (n, k) -selector. (And the contribution to our running time bounds from
 275 the strong $(n, 1)$ -selectors is comparatively negligible.) This guarantees that, during each
 276 complete cycle (of length $O(k^2 \log n)$) of this selector, for any node v with k in-neighbors
 277 and any v 's in-neighbor u there will be a step when v is in the receiving state and u is the
 278 only in-neighbor in the transmitting state.

279 We will make yet another assumption in the paper, this one concerning the initial know-
 280 ledge that the nodes possess about the digraph. In the standard definition given earlier in
 281 this section, the nodes know the size n of the digraph but do not know its topology. Without
 282 any loss of generality, we will relax the latter restriction as follows:

283 (INN) We assume that when the computation starts each node v knows the labels of its
 284 in-neighbors, that is the set $N^-(v)$.

285 The reason that we can assume (INN) is that any protocol in our paper, prior to starting
 286 its execution, can execute one cycle of ROUNDROBIN, where each node transmits only its
 287 own label. This costs only time $O(n)$, so the asymptotic running time of the protocol is not
 288 affected. (In the paper we only consider protocols whose worst-case running time is $\Omega(n)$.)

289 **3** $\tilde{O}(n^{1.5})$ -Time Protocol for Acyclic Digraphs

290 We first consider ad-hoc radio networks whose underlying digraph G is acyclic and has one
 291 designated target node t that is reachable from all other nodes in G . We give a determ-
 292 inistic information gathering protocol that gathers all rumors in the target node t in time
 293 $O(n^{1.5} \log^3 n)$, independently of the topology of G .

294 As explained in Section 2, we make Assumptions (MFC), (SRT) and (INN), namely
 295 that the protocol has multiple frequency channels available, that on each frequency it can
 296 simultaneously receive and transmit messages at each step, and that each node knows its
 297 in-neighbors.

298 The protocol is based on strong selectors 2^j -SELECT and on ROUNDROBIN, following the
 299 principles outlined in Section 2. Thus, whenever we say that a node w *transmits according*
 300 *to* 2^j -SELECT during some time interval $[\eta, \eta']$, we mean that for any time step $\tau \in [\eta, \eta']$
 301 node w transmits if and only if $w \in S_{\tau \bmod \ell_j}^j$. (Recall that ℓ_j is the size of 2^j -SELECT.) The
 302 concept of *transmitting according to* ROUNDROBIN is defined analogously.

303 Let $\theta = \frac{1}{2}(\log n - \log \log n) + 2$. In the algorithm below we use a sequence of $\theta + 1$
 304 values $\beta_0, \beta_1, \dots, \beta_\theta$, defined as follows: $\beta_0 = 0$, $\beta_j = \sum_{g < j} \ell_g$ for $j = 1, \dots, \theta - 1$, and
 305 $\beta_\theta = \sum_{g < \theta} \ell_g + n$.

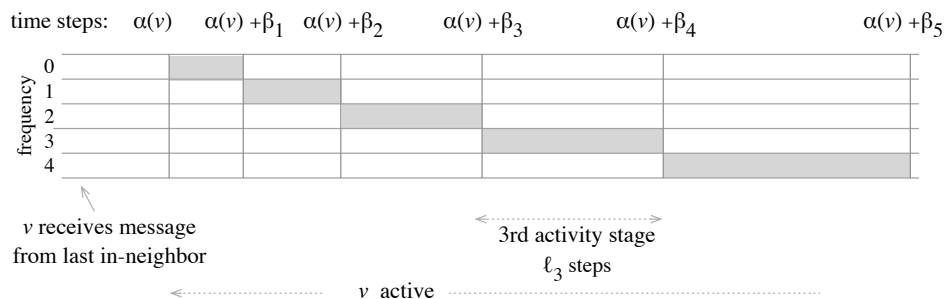
306 **Protocol ACYGATHER.** We start with an overview of the algorithm. The algorithm
 307 transmits on θ frequencies numbered $0, 1, \dots, \theta - 1$. Each node v uses information received

308 from its in-neighbors to determine its *activation time*, denoted by $\alpha(v)$. Node v will be *active*
 309 during its *activity period* $[\alpha(v), \alpha(v) + \beta_\theta]$; before step $\alpha(v)$ it is called *dormant* and after step
 310 $\alpha(v) + \beta_\theta - 1$ it is called *expired*. Dormant and expired nodes do not transmit; active nodes
 311 may or may not transmit at any given step. While active, v will also periodically compute
 312 and transmit a value called its *recommended wake-up time*, denoted rws_v . (The out-neighbors
 313 of v use these values to determine their own activation times.) Each message transmitted
 314 by v will contain the following information: (i) all rumors collected by v , including its own,
 315 (ii) the label of v , and (iii) the current value of rws_v .

316 We are now ready to detail the steps of the algorithm, and in particular to describe how
 317 exactly the values of $\alpha(v)$ and rws_v are computed by a node v .

318 First, we explain how v determines its activation time $\alpha(v)$. If v is a source node (that
 319 is, its in-degree is 0), then $\alpha(v) = 0$. Otherwise $\alpha(v)$ is determined by the messages received
 320 by v , as follows. For each in-neighbor u of v , we denote by $rws_{u \rightarrow v}^1$ the *first* rws_u value
 321 received by v from u . (This may not be the first rws_u value *transmitted* by u , since earlier
 322 transmissions of u might have collided at v .) Node v waits until it receives messages from
 323 all its in-neighbors, and, as soon as this happens, if u is the last in-neighbor of v that
 324 successfully transmitted to v , then v sets $\alpha(v) = rws_{u \rightarrow v}^1$.

325 Next, we define the transmission sequence of v . The activity period $[\alpha(v), \alpha(v) + \beta_\theta]$ of
 326 v is divided into θ *activity stages*, where, for $j = 0, 1, \dots, \theta - 1$, the j th activity stage consists
 327 of the time interval $[\alpha(v) + \beta_j, \alpha(v) + \beta_{j+1}]$. (See Figure 2.) During its j th activity stage,
 328 for $j \leq \theta - 2$, node v transmits according to selector 2^j -SELECT using frequency j . During
 329 the $(\theta - 1)$ th activity stage, the protocol transmits using ROUNDROBIN on frequency $\theta - 1$.
 330 At all other times v does not transmit. The recommended wake-up time value included in
 331 v 's messages during its j th activity stage is $rws_v = \alpha(v) + \beta_{j+1}$. (Note that it changes from
 332 one activity stage to next.)



■ **Figure 2** Illustration of activity stages. (The picture is not up to scale. In reality the length of activity stages increases at rate 4.) Shaded regions show frequencies used in different activity stages.


333 **Correctness.** We first show that the algorithm is correct, in the sense that each rumor will
 334 eventually reach the target node t . This should be intuitively clear, because once a node
 335 becomes active, it is guaranteed to successfully transmit its message to its all out-neighbors
 336 using the ROUNDROBIN protocol during its last activity stage.

337 Formally, we can prove correctness by induction. For a given node v , let $\delta(v)$ denote the
 338 length of the *longest* directed path from some source node to v . (This path cannot repeat
 339 vertices due to our assumption that G is acyclic.) Let $\delta^* = \max_{v \in G} \delta(v)$. The definition of
 340 $\delta()$ implies directly the following observation:

341 ► **Observation 1.** The values of $\delta()$ satisfy the following properties:

- 342 (i) $\delta(v) = 0$ if and only if v is a source node of G .
 343 (ii) $\delta(v) = \delta^*$ if and only if $v = t$.
 344 (iii) If $u \in N^-(v)$ then $\delta(u) < \delta(v)$.
 345 (iv) If $\delta(v) > 0$ then there is $u \in N^-(v)$ with $\delta(u) = \delta(v) - 1$.

346 It is convenient to visualize $\delta()$ in terms of partitioning of G into layers $B_0, B_1, \dots, B_{\delta^*}$,
 347 where B_i denotes the set of nodes with $\delta(v) = i$. This partitioning is illustrated in Figure 3.
 348 The key properties are that B_0 consists of the source nodes, $B_{\delta^*} = \{t\}$, and all edges go
 349 from lower- to higher-indexed layers. (This representation will be useful in Section 5.)



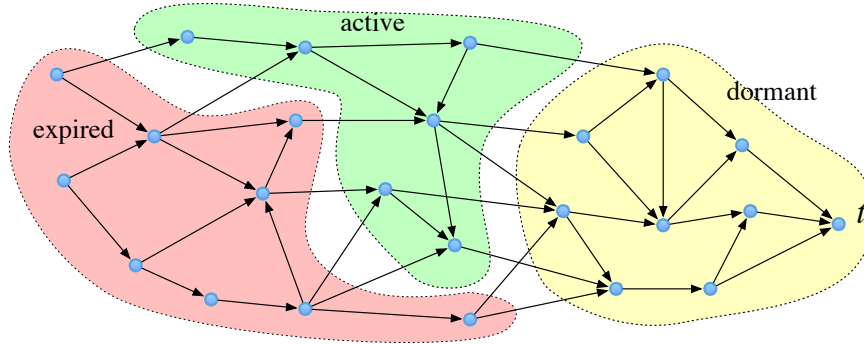
■ **Figure 3** Partiton of G into layers $B_0, B_1, \dots, B_{\delta^*}$. In this example $\delta^* = 5$.

350 We claim that for each $i = 0, 1, \dots, \delta^*$, and for each node $v \in B_i$, we have $\alpha(v) \leq i\beta_\theta$,
 351 and that at time $\alpha(v)$ all rumors from $Anc[v]$ are already in v . This is sufficient, because
 352 this guarantees that after $\delta^*\beta_\theta$ steps all rumors will be collected in the target node t . The
 353 claim is trivially true for $i = 0$, because all nodes in B_0 are source nodes. For $i > 0$, suppose
 354 that the claim holds for indices $i' = 0, 1, \dots, i - 1$. This means that for each $w \in N^-(v)$,
 355 as $w \in \bigcup_{i'=0}^{i-1} B_{i'}$, we have $\alpha(w) \leq (i - 1)\beta_\theta$ and at time $\alpha(w)$ all rumors from $Anc[w]$
 356 are in w . In its last activity stage $[\alpha(w) + \beta_{\theta-1}, \alpha(w) + \beta_\theta)$ node w will transmit using
 357 ROUNDROBIN, so v will receive a message from w no later than at time $\alpha(w) + \beta_\theta - 1$. This
 358 messages includes all rumors in $Anc[w]$ and w 's current recommended activation time value
 359 rws_w , whose maximum value is $\alpha(w) + \beta_\theta \leq (i - 1)\beta_\theta + \beta_\theta \leq i\beta_\theta$. Then the definition of
 360 $\alpha(v)$ implies that $\alpha(v) \leq i\beta_\theta$. At time $\alpha(v)$ node v will have received messages from all its
 361 in-neighbors, so at that time it will collect rumors from $\bigcup_{w \in N^-(v)} Anc[w] \cup \{v\} = Anc[v]$,
 362 completing the proof of the inductive step, and the claim.

363 **Running time.** Next, we analyze the running time of Protocol ACYGATHER. Figure 4
 364 provides a snapshot of the computation of Protocol ACYGATHER that should be helpful in
 365 understanding our analysis. It shows three types of nodes: expired, active and dormant.
 366 The in-neighbors of dormant nodes can be of any three types, but each dormant node has
 367 at least one in-neighbor that is either dormant or active. The in-neighbors of active nodes
 368 are either expired or active. All in-neighbors of expired nodes are expired. These properties
 369 follow from the algorithm, because the activation times of all non-source nodes v satisfy

$$370 \quad \max_{u \in N^-(v)} \alpha(u) < \alpha(v) \leq \max_{u \in N^-(v)} \alpha(u) + \beta_\theta,$$

371 and because the activation periods of all nodes have the same length (which implies that an
 372 expiration time of a node is strictly after the expiration times of all its in-neighbors).



■ **Figure 4** Bird's-eye view of the computation of Protocol ACYGATHER.

373 To establish our upper bound, we choose in the graph G a *critical path*

374
$$P = (v_0, v_1, \dots, v_p = t),$$

375 defined as follows: for each $a = p - 1, p - 2, \dots, 0$, v_a is the in-neighbor of v_{a+1} who was last
 376 in succeeding to transmit its message to v_{a+1} (that is, $\alpha(v_{a+1}) = rws_{v_a \rightarrow v_{a+1}}^1$), and v_0 is a
 377 source node. The argument about correctness of protocol ACYGATHER, presented above,
 378 implies that P is well defined. (Note that, since we define this path in the backward order,
 379 the indexing of the nodes v_a can be determined only after we determine the whole path.)
 380 The definition of P implies that the overall running time is upper-bounded by the time for
 381 the rumor of v_0 to reach t along P .

382 If at a step τ a node v is in its j -th activity stage (that is, $\tau \in [\alpha(v) + \beta_j, \alpha(v) + \beta_{j+1})$)
 383 then we refer to j as v 's *stage index in step* τ . We extend this (artificially) to dormant and
 384 expired nodes as follows: if v is dormant then its stage index is -1 , and if v is expired then
 385 its stage index is θ . As time progresses, an active node will move from one activity stage
 386 to next, which results in an increment of its stage index. Within any time interval multiple
 387 nodes may have their stage indices incremented.

388 Note that the stage index of each individual node is incremented $\theta + 1 = O(\log n)$
 389 times, so the total number of these increments in the whole computation is $O(n \log n)$. Our
 390 estimate on the running time is obtained by “charging” the delay between activation times
 391 of consecutive nodes on P to stage-index increments in the graph. The intuition is that
 392 if collisions cause a long delay when v_a attempts to send its message to v_{a+1} , then v_{a+1}
 393 must have many in-neighbors that transmit in this time period. But then these in-neighbors
 394 will have their stage indices incremented during this time, and since the overall number of
 395 stage-index increments is $O(n \log n)$, we cannot have too many such long delays.


396 We now formalize this intuition. Consider some node $v_a \neq t$ on P . (See Figure 5.) Our
 397 argument is based on the following key lemma.

398 ► **Lemma 1.** *There are the total of $\Omega(\frac{1}{\sqrt{n \log n}}(\alpha(v_{a+1}) - \alpha(v_a)))$ stage-index increments in*
 399 *the time interval $[\alpha(v_a), \alpha(v_{a+1}))$.*

400 **Proof.** Suppose that the first transmission of v_a that is successfully received by v_{a+1} occurs
 401 during v_a 's h -th activity stage.

402 ► **Claim 1.** For $a < p$ and $h < \theta - 1$ we have $\alpha(v_{a+1}) - \alpha(v_a) = O(4^h \log n)$.

403 This claim follows from the definition of P , as $\alpha(v_{a+1}) = rws_{v_a \rightarrow v_{a+1}}^1 = \alpha(v_a) + \beta_{h+1}$,
 404 and $\beta_{h+1} = \sum_{g < h+1} \ell_g = O(4^h \log n)$.



■ **Figure 5** Illustration of the time analysis for acyclic graphs.

405 We now continue the proof of the lemma. If $h = 0$ then there is at least one stage
 406 increment in $[\alpha(v_a), \alpha(v_{a+1}))$ (namely the increment of the stage index of v_a from -1 to 0
 407 at time $\alpha(v_a)$) and $\alpha(v_{a+1}) - \alpha(v_a) = \ell_0 = O(\log n)$, so the lemma holds trivially.

408 Thus for the rest of the proof we can assume that $1 \leq h \leq \theta - 1$. By the choice of h , v_a
 409 has not succeeded in its $(h - 1)$ th activity stage $[\alpha(v_a) + \beta_{h-1}, \alpha(v_a) + \beta_h)$. Let U be the set
 410 of in-neighbors of v_{a+1} (including v_a) whose $(h - 1)$ th activity stage overlapped that of v_a .

411 ► **Claim 2.** $|U| > 2^{h-1}$.

412 To justify Claim 2, we argue by contradiction. Suppose that $|U| \leq 2^{h-1}$. During this
 413 activity stage v_a transmitted according to 2^{h-1} -SELECT using only frequency $h - 1$. Further,
 414 by the definition of the protocol, at each step of this stage the in-neighbors of v_{a+1} with stage
 415 index other than $h - 1$ did not use frequency $h - 1$ for transmissions. So the transmissions
 416 from v_a to v_{a+1} in this stage can only conflict with transmissions from $U \setminus \{v_a\}$ to v_{a+1} . The
 417 definition of strong selectors and the assumption that $|U| \leq 2^{h-1}$ imply that then v_a would
 418 have successfully transmitted to v_{a+1} during its $(h - 1)$ th activity stage, contradicting the
 419 definition of h . Thus Claim 2 is indeed true.

420 ► **Claim 3.** The total number of stage-index increments in time interval $[\alpha(v_a), \alpha(v_{a+1}))$ is
 421 at least 2^{h-1} .

422 The proof of Claim 3 is quite simple: The $(h - 1)$ th activity stage lasts ℓ_{h-1} steps,
 423 so for each node in U its $(h - 1)$ th activity stage ends before time $\alpha(v_a) + \beta_h + \ell_{h-1} <$
 424 $\alpha(v_a) + \beta_{h+1} = \alpha(v_{a+1})$. By the definition of U , it also cannot end before $\alpha(v_a)$. So, each
 425 time the $(h - 1)$ th activity stage of a node in U ends, it contributes 1 to the total number
 426 of stage-index increments in time interval $[\alpha(v_a), \alpha(v_{a+1}))$. This implies Claim 3.

427 Now, to complete the proof of the lemma we have two cases. If $h < \theta - 1$ then, by
 428 Claim 3, the number of stage-index increments in time interval $[\alpha(v_a), \alpha(v_{a+1}))$ is at least

$$429 \quad 2^{h-1} = \frac{1}{2} \cdot 2^{-h} \cdot 4^h = \Omega\left(\frac{1}{\sqrt{n} \log n} (\alpha(v_{a+1}) - \alpha(v_a))\right),$$

430 because $2^{-h} \geq 2^{-\frac{1}{2} \log n} = 1/\sqrt{n}$ (as $h \leq \frac{1}{2} \log n$) and $4^h = \Omega((\alpha(v_{a+1}) - \alpha(v_a))/\log n)$, by
 431 Claim 1.

432 On the other hand, if $h = \theta - 1$ then $\alpha(v_{a+1}) - \alpha(v_a) = n$. From the definition of θ
 433 we have $h \geq \frac{1}{2} \log n - \log \log n + 1$ so, applying Claim 3, we obtain that the number of
 434 stage-index increments in time interval $[\alpha(v_a), \alpha(v_{a+1}))$ is at least $2^{h-1} \geq 2^{\frac{1}{2} \log n - \log \log n} =$
 435 $\sqrt{n}/\log n = \frac{1}{\sqrt{n} \log n} (\alpha(v_{a+1}) - \alpha(v_a))$. ◀

436 We now show how we can use Lemma 1 to establish an $O(n^{1.5} \log^3 n)$ upper bound on
 437 the running time of Protocol ACYGATHER.

438 ► **Theorem 2.** *Let G be an acyclic directed graph with n vertices and a designated target node*
 439 *reachable from all other nodes. Algorithm ACYGATHER completes information gathering on*
 440 *G in time $O(n^{1.5} \log^3 n)$.*

441 **Proof.** Let T be the running time of Protocol ACYGATHER in our relaxed model of com-
 442 munication, namely with assumptions (MFC), (SRT), and (INN). Since $\alpha(v_0) = 0$ and
 443 $T \leq \alpha(v_p)$, we can bound this running time as $T \leq \sum_{a=0}^{p-1} (\alpha(v_{a+1}) - \alpha(v_a))$. Then
 444 Lemma 1 implies that the total number of stage index increments during the computation
 445 is $\Omega(T/\sqrt{n} \log n)$. Since this number is also $O(n \log n)$, it gives us that $T = O(n^{1.5} \log^2 n)$.

446 The number of frequencies is $\kappa = O(\log n)$. Thus, as explained in Section 2, we can
 447 eliminate all three assumptions (MFC), (SRT), and (INN), increasing the running time only
 448 by a factor $O(\log n)$. Such modified Protocol ACYGATHER will run in time $O(n^{1.5} \log^3 n)$
 449 in the standard model of ad-hoc radio networks. ◀

450 *Note:* With more careful analysis the logarithmic factors in Theorem 2 can be slightly
 451 improved (at least to $O(n^{1.5} \log^{2.5} n)$). We leave this as an exercise for the reader.

452 **4 $\tilde{O}(n^{1.5})$ -Time Protocol for Arbitrary Digraphs**

453 We now extend our information gathering protocol ACYGATHER from Section 3 to arbitrary
 454 digraphs, retaining running time $O(n^{1.5} \log^3 n)$. Throughout this section G will denote an
 455 n -vertex digraph with a designated target node t that is reachable from all other nodes in
 456 G .

457 The main obstacle we need to overcome is that protocol ACYGATHER critically depends
 458 on G being acyclic. For instance, in that protocol each node waits until it receives messages
 459 from all its in-neighbors. If cycles are present in G , this leads to a deadlock, where each node
 460 in a cycle waits for its predecessor. On the other hand, the known gossiping protocols [6,
 461 29, 18] do not work correctly if the graph is not strongly connected, because they rely on
 462 broadcasting to periodically flush out some rumors from the system, and on leader election
 463 to synchronize computation.

464 The idea behind our solution is to integrate protocol ACYGATHER with the gossiping pro-
 465 tocol from [18], using ACYGATHER to transmit information between different sc-components
 466 of G and using gossiping to disseminate information within sc-components. The idea is nat-
 467 ural but it faces some technical challenges. One challenge is that the sc-components are
 468 actually not known. In fact, a node v doesn't even know the size of $C(v)$, but it needs to
 469 provide this size to the gossiping protocol. To get around this issue, the algorithm runs in
 470 parallel $\log n$ copies of a gossiping protocol for sizes that are powers of 2. Another challenge
 471 is that transmissions from outside of an sc-component may interfere with the execution of the
 472 gossiping protocol in this sc-component. This will be addressed by executing the gossiping
 473 algorithm repeatedly until it succeeds. In order to verify whether the gossiping succeeded,
 474 its nodes will distribute and collect additional information, not just rumors.

475 **Protocol SCCGOSSIP for gossiping.** We will refer to the gossiping algorithm from [18]
 476 as SCCGOSSIP. The following property of SCCGOSSIP is crucial for our algorithm:

477 (sc) If the input digraph is strongly connected and has at most \bar{n} vertices, with the
 478 node labels from the set $[N] = \{0, 1, \dots, N - 1\}$, then algorithm SCCGOSSIP completes
 479 gossiping in time $O(\bar{n}^{4/3} \log^{10/3} N)$.

480 The bound on the running time of SCCGOSSIP given in [18] (see their Theorem 2) assumes
 481 that N is bounded polynomially in \bar{n} and it does not explicitly separate the dependence on
 482 \bar{n} and N . In Appendix A we explain how the bound in Property (scc) follows from the
 483 analysis in [18].

484 As explained earlier, one idea of our algorithm is to execute SCCGOSSIP on its sc-
 485 components. The details of this will be provided shortly. For now, we only make an
 486 observation that captures one basic principle of this process. Let A be an sc-component
 487 of size n_A and let $j = \lceil \log n_A \rceil$, so that $2^{j-1} < n_A \leq 2^j$. Let SCCGOSSIP_j denote SCCGO-
 488 SIPP specialized for strongly connected digraphs of size $\bar{n} = 2^j$ and label set $[N] = [n]$, and
 489 let $T_{\text{SCC}}(j)$ be the running time of SCCGOSSIP_j on such digraphs. Suppose also that all
 490 nodes in $\text{Anc}(A)$ are idle and that the nodes in A execute SCCGOSSIP_j , all starting at the
 491 same time and ignoring any information collected so far in A . Since the nodes in $\text{Anc}(A)$
 492 are idle, there will be no interference from outside A . Then the execution of SCCGOSSIP_j
 493 on A will be identical to its execution on the sub-digraph of G induced by A ; that is, as if
 494 the rest of G did not exist. Therefore, using Property (scc), this execution of SCCGOSSIP_j
 495 will complete correctly in time $T_{\text{SCC}}(j) = O(n_A^{4/3} \log^{10/3} n)$. In our analysis we will use the
 496 estimate $T_{\text{SCC}}(j) = \tilde{O}(n_A^{4/3})$, hiding the polylogarithmic factor, as the degree of this factor
 497 will not affect the overall asymptotic running time of our algorithm.

498 **Algorithm ARBGATHER.** Our protocol can be thought of as running two parallel sub-
 499 routines, the SCC-subroutine and the ACY-subroutine, that use two disjoint sets of frequen-
 500 cies. There will be θ ACY-frequencies indexed $0, 1, \dots, \theta - 1$, where $\theta = \frac{1}{2}(\log n - \log \log n) + 2$,
 501 as in Section 3. These will be used by the ACY-subroutine to simulate protocol ACYGATHER.
 502 We will also have θ' SCC-frequencies indexed $0, 1, \dots, \theta' - 1$, used by the SCC-
 503 subroutine to simulate protocol SCCGOSSIP. Due to using different frequencies, there will
 504 be no signal interference between these two subroutines. In the SCC-subroutine, each
 505 SCC-frequency j will be used to simulate SCCGOSSIP_j . Thus these different instances of
 506 SCCGOSSIP_j will also do not interfere with each other.

507 Before providing the detailed descriptions of these subroutines, we give an overview of
 508 the algorithm, expanding on the intuitions described earlier in this section. A good way
 509 to visualize the computation of Algorithm ARBGATHER is to think of the ACY-subroutine
 510 as a “master process” that transmits rumors between sc-components (that form an acyclic
 511 graph), while the SCC-subroutine is run as a “background” iterative process in all nodes.
 512 Consider some sc-component A . The nodes in A will repeatedly run in parallel θ' copies
 513 of protocol SCCGOSSIP , one for each SCC-frequency, all attempting to determine A and to
 514 collect rumors from $\text{Anc}[A]$. While a node v in A executes these protocols SCCGOSSIP , its
 515 in-neighbors in $\text{Anc}(A)$ (that is, outside of A) may be transmitting to v . Messages sent from
 516 $\text{Anc}(A)$ to v on ACY-frequencies contribute to progress, as they contain rumors from $\text{Anc}(A)$,
 517 that v needs to collect, and do not interfere with v 's executions of SCCGOSSIP . Messages
 518 sent from $\text{Anc}(A)$ to v on SCC-frequencies are problematic, because they can collide with
 519 transmissions of SCCGOSSIP 's coming from in-neighbors of v in A . (And also, they do not
 520 provide useful information.)

521 We focus on one particular value of j , namely $j = \lceil \log n_A \rceil$, as for this j SCCGOSSIP_j
 522 is most likely to succeed quickly on A . Some number of repetitions of SCCGOSSIP_j in A
 523 may fail, either because of collisions with transmissions on SCC-frequencies from $\text{Anc}(A)$,
 524 or because of some nodes in A not being yet ready to participate in the gossiping in A , if
 525 they have not yet received messages from all of their in-neighbors in $\text{Anc}(A)$. Eventually
 526 though, all the nodes in $\text{Anc}(A)$ will complete their own executions of SCCGOSSIP_j , stop the
 527 simulation of SCCGOSSIP altogether, and switch to the ACY-subroutine. As a result, these

528 nodes will no longer interfere with the executions of SCCGOSSIP_j in A , and their execution
 529 of the ACY -subroutine will successfully transmit their messages to the nodes in A . At this
 530 point a repetition of SCCGOSSIP_j in A will succeed, in the sense that each node $v \in A$ will
 531 determine A , it will receive all rumors from $\text{Anc}[A]$, and it will also learn that *other nodes* in
 532 A have successfully completed this process. This will happen for all nodes in A during the
 533 same repetition of SCCGOSSIP_j , allowing the algorithm to “synchronize” the computation
 534 of all nodes in A so that they all can simultaneously stop executing the SCC -subroutine and
 535 switch to executing the ACY -subroutine.

536 Next, we give formal descriptions of both subroutines.

537 *The SCC-subroutine.* For each $j = 0, 1, \dots, \theta' - 1$, SCC -frequency j will be used to simulate
 538 protocol SCCGOSSIP_j . For $s = 0, 1, \dots$ let $\psi_{j,s} = 2sT_{\text{SCC}}(j)$. For each SCC -frequency j , the
 539 computation of v on frequency j is partitioned into j -frames, where the s -th j -frame, for any
 540 $s \geq 0$, is $[\psi_{j,s}, \psi_{j,s+1})$ — a time interval sufficient for two complete simulations (described
 541 below) of SCCGOSSIP_j on a digraph with at most 2^j nodes. For each j , these simulations
 542 start at time 0 and stop as soon as v determines that for at least one SCC -frequency j' the
 543 simulation of $\text{SCCGOSSIP}_{j'}$ successfully completed in $C(v)$. (More precisely, they stop at
 544 time $\alpha_{\text{ACY}}(v)$ that will be defined shortly.)

545 Consider a j -frame r , for some $r \geq 0$. In this j -frame v executes two runs of SCCGOSSIP_j .
 546 The first one is called the *exploration run* and it is used to distribute node labels; that is,
 547 any node v uses its own label v as the “rumor” for the purpose of gossiping. The second
 548 run of SCCGOSSIP_j is called the *dissemination run*. In this run v ’s “rumor” is the 5-tuple

$$549 \quad [v, \tilde{C}(v), N^-(v), \tilde{N}_{\text{ACY}}^-(v), R(v)],$$

550 where

- 551 ■ $\tilde{C}(v)$ is the set of labels received by v during the exploration run of j -frame r , including
 552 v itself.
- 553 ■ $\tilde{N}_{\text{ACY}}^-(v) \subseteq N^-(v)$ is the set of in-neighbors of v that have successfully transmitted a
 554 message to v on some ACY -frequency before time $\psi_{j,r}$ (the beginning of r -th j -frame),
 555 and
- 556 ■ $R(v)$ is the set of all (original) rumors received on ACY -frequencies before time $\psi_{j,r}$, plus
 557 the rumor of v .

558 Let $\tilde{C}'(v)$ be the set of node labels received by v in this dissemination run of SCCGOSSIP_j ,
 559 including v itself. Then, immediately after the dissemination run, v performs three tests:

- 560 *Test 1:* Is it true that $\tilde{C}(v) = \tilde{C}'(v)$?
- 561 *Test 2:* Is it true that $\tilde{C}(v) = \tilde{C}(u)$ for all $u \in \tilde{C}'(v)$?
- 562 *Test 3:* Is it true that $N^-(u) \setminus \tilde{N}_{\text{ACY}}^-(u) \subseteq \tilde{C}(v)$ for all $u \in \tilde{C}'(v)$?

563 If at least one of these tests fails, v continues the execution of the SCC -subroutine on
 564 frequency j , proceeding to j -frame $r + 1$. If all tests pass, v aborts its SCC -subroutine
 565 altogether (thus aborting the simulations of $\text{SCCGOSSIP}_{j'}$ for all frequencies j') and switches
 566 to the ACY -subroutine, with its set of collected rumors being $\bigcup_{u \in \tilde{C}(v)} R(u)$. Let $\alpha_{\text{ACY}}(v) =$
 567 $\psi_{j,r+1}$ denote this time step. We call it the *ACY-activation time* of v .

568 *The ACY-subroutine.* The ACY -activation time $\alpha_{\text{ACY}}(v)$ of v , defined above, plays the
 569 role of v ’s activation time in protocol ACYGATHER . In this subroutine v will transmit
 570 at the ACY -frequencies and it simply executes ACYGATHER in its *ACY-activity* period

571 $[\alpha_{\text{ACY}}(v), \alpha_{\text{ACY}}(v) + \beta\theta)$. The activity stages and the transmissions of each node are defined
 572 in exactly the same way as in protocol ACYGATHER (except that we use $\alpha_{\text{ACY}}(v)$ instead of
 573 $\alpha(v)$). At each step any node is in one of three possible states: ACY-dormant, ACY-active,
 574 or ACY-expired. These concepts are natural adaptations of the corresponding concepts for
 575 protocol ACYGATHER. (The only difference is that now ACY-dormant nodes are not truly
 576 “dormant”, as they execute the SCC-subroutine.)

577 **Correctness.** To show correctness, we need to show that all rumors will eventually reach
 578 t , the target node of G . The basic structure of the proof is similar to the proof for Pro-
 579 tocol ACYGATHER in Section 3, namely we proceed by induction. The difference is that
 580 now in one step of the inductive argument we analyze progress in a whole sc-component,
 581 rather than a single vertex; that is we show that all nodes in any given sc-component A
 582 will collect all rumors from $\text{Anc}[A]$. This can be captured formally by considering an aux-
 583 iliary acyclic digraph $\text{SccDag}(G)$ that represents the structure of sc-components of G . The
 584 vertices of $\text{SccDag}(G)$ are the sc-components of G , and for any two sc-components A and
 585 A' , we include edge (A, A') in $\text{SccDag}(G)$ iff there are vertices $u \in A$ and $v \in A'$ with edge
 586 $(u, v) \in E$. The induction is then with respect to the values of $\delta(A)$, the maximum path
 587 length from a source sc-component to A in $\text{SccDag}(G)$. In the argument below we focus on
 588 analyzing the algorithm’s progress within one sc-component A (see Lemma 3 below), as the
 589 details of the argument showing that the nodes in A will receive the rumors from $\text{Anc}(A)$
 590 are essentially the same as in Section 3, that is they are based on the usage of ROUNDROBIN
 591 in Protocol ACYGATHER.

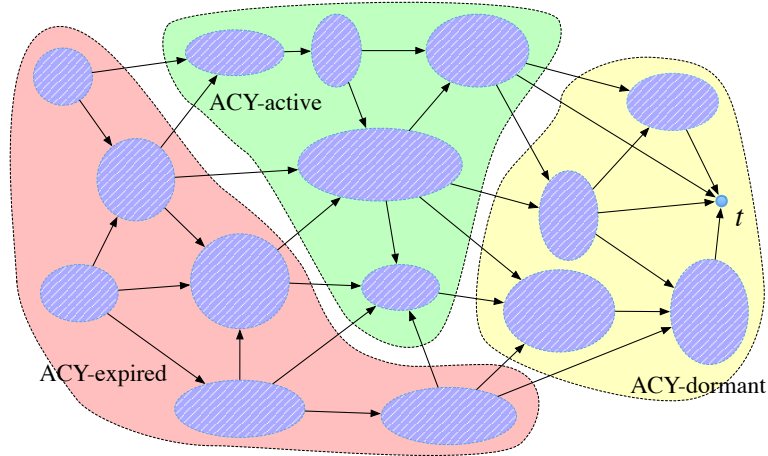
592 With each node v of G we will associate a time step called the *SCC-ready time*, denoted
 593 $\rho_{\text{SCC}}(v)$. Its definition is identical to the activation time in ACYGATHER: If $N_{\text{ACY}}^-(v) = \emptyset$
 594 then $\rho_{\text{SCC}}(v) = 0$. Otherwise, $\rho_{\text{SCC}}(v)$ is the last received value $\text{rws}_{u \rightarrow v}^1$ for $u \in N_{\text{ACY}}^-(v)$,
 595 where $\text{rws}_{u \rightarrow v}^1$ denotes the first rws_u value received by v from u . As explained earlier, these
 596 rws_u values will be received on ACY-frequencies. (We stress that the SCC-ready times are
 597 used only for the analysis. In fact, the value of $\rho_{\text{SCC}}(v)$ depends on $C(v)$, so it cannot even be
 598 computed before v determines $C(v)$.) We extend this definition naturally to sc-components;
 599 if A is an sc-component, we let $\rho_{\text{SCC}}(A) = \max_{u \in A} \rho_{\text{SCC}}(u)$.

600 Any node v that ever becomes ACY-active is guaranteed to successfully transmit during
 601 the ACY-subroutine, because this subroutine involves a round of ROUNDROBIN. Thus the
 602 key difficulty in proving correctness is to show that, for any sc-component A , each node
 603 $v \in A$ will correctly complete the SCC-subroutine, meaning that it will collect all rumors
 604 from $\text{Anc}[v]$ before time $\alpha_{\text{ACY}}(v)$, when it switches to executing the ACY-subroutine. We
 605 make this more precise in the lemma below.

606 **► Lemma 3.** *Let A be an sc-component of G . Let n_A be its size and $j = \lceil \log n_A \rceil$. Then*
 607 *(i) $\rho_{\text{SCC}}(A)$ is well-defined (that is, finite).*
 608 *(ii) All nodes in A complete subroutine SCCGOSSIP at the same time. In other words, all*
 609 *values $\alpha_{\text{ACY}}(v)$, for $v \in A$, are equal. (Below we use notation $\alpha_{\text{ACY}}(A)$ for this common*
 610 *value of $\alpha_{\text{ACY}}(v)$ for $v \in A$.)*
 611 *(iii) At time $\alpha_{\text{ACY}}(A)$ each node in A has all rumors from $\text{Anc}[A]$.*
 612 *(iv) $\alpha_{\text{ACY}}(A) \leq \rho_{\text{SCC}}(A) + 4 \cdot T_{\text{SCC}}(j)$.*

613 **Proof.** As indicated earlier, the proof of Lemma 3 is by induction with respect to $\delta(A)$.
 614 That is, assuming that all nodes in $\text{Anc}(A)$ satisfy the claims (i)-(iv) of the lemma, we argue
 615 that they also hold for A .

616 We start with part (i). By the inductive assumption, at some time step all nodes in
 617 $\text{Anc}(A)$ will complete their execution of the SCC-subroutine and start the ACY-subroutine.



■ **Figure 6** Bird's-eye view of the snapshot of a computation of Protocol ARBGATHER. The striped ovals represent sc -components of G or, equivalently, vertices of $SccDag(G)$. The arrows are the edges of $SccDag(G)$, with two sc -components connected by an edge if they contain vertices connected by an edge of G . Note that in each sc -component all nodes are of the same type. This follows from Lemma 3 and from all ACY-activity intervals having the same length.

618 Since the ACY-subroutine involves ROUNDROBIN, eventually all nodes in A will receive
 619 messages of Protocol ACYGATHER from their in-neighbors outside A . (For details, see the
 620 argument in Section 3). This implies that $\rho_{SCC}(v)$ is well-defined for each $v \in A$, implying
 621 part (i).

622 In the proof of the remaining parts (ii)-(iv), the main idea is that once the nodes in
 623 $Anc(A)$ stop executing their SCC-subroutine, and thus do not interfere with A , the compu-
 624 tation of the SCC-subroutine in frequency j is guaranteed to succeed in A in time $T_{SCC}(j)$.
 625 One minor but complicating caveat is that this may not actually be true as stated, because
 626 a node in A can conceivably get “lucky” by having passed Tests 1-3 on some other frequency
 627 j' , in which case this node will terminate its execution of the whole SCC-subroutine in the
 628 middle of the current j -frame. In this case we need to argue that then in fact *all* nodes in
 629 A will succeed on frequency j' at that time.

630 To formalize this idea, we first show that at least one $v \in A$ will terminate its SCC-
 631 subroutine and become ACY-active at time step $\alpha_{ACY}(v) \leq \rho_{SCC}(A) + 4 \cdot T_{SCC}(j)$. Let
 632 r be the index such that $\psi_{j,r-1} \leq \rho_{SCC}(A) < \psi_{j,r}$. Since $\psi_{j,r+1} = \psi_{j,r-1} + 4 \cdot T_{SCC}(j) \leq$
 633 $\rho_{SCC}(A) + 4 \cdot T_{SCC}(j)$, it is sufficient to show that at least one $v \in A$ will have $\alpha_{ACY}(v) \leq \psi_{j,r+1}$.
 634 If some node $v \in A$ passes Tests 1-3 on some frequency $j' \neq j$ before time $\psi_{j,r+1}$ then
 635 $\alpha_{ACY}(v) \leq \psi_{j,r+1}$, and we are done. So suppose that this is not the case, that is all nodes
 636 in A remain ACY-dormant and execute $SCCGOSSIP_j$ until time $\psi_{j,r+1}$. By the choice of r ,
 637 at times $\tau \geq \psi_{j,r}$ the nodes in $Anc(A)$ are either ACY-active or ACY-expired, so they do
 638 not transmit at SCC-frequencies anymore. Then the two runs of $SCCGOSSIP_j$ in j -frame r
 639 will not experience any interference from outside A . Therefore, by the paragraph before the
 640 description of the algorithm, both these runs will complete correctly, namely for each node
 641 $v \in A$ all Tests 1-3 will pass and v will become ACY-active at time $\alpha_{ACY}(v) = \psi_{j,r+1}$, as
 642 needed. (Note that in this case this holds in fact for *all* nodes in A .)

643 So now we know that at least one $v \in A$ became ACY-active at time $\psi_{j,r+1}$ or earlier.
 644 Let v be the *first* node in A that became ACY-active, and let j' be the frequency for which
 645 v passed Tests 1-3, say after executing two runs of $SCCGOSSIP_{j'}$ in some j' -frame s . Thus

646 $\alpha_{\text{ACY}}(v) = \psi_{j',s+1} \leq \psi_{j,r+1}$. We need to show that then all other nodes in A pass these tests
647 in the same j' -frame s .

648 To this end, we claim first that at time $\psi_{j',s+1}$ we have $\tilde{C}(v) = A$. Indeed, Tests 1-2
649 imply that each $u \in \tilde{C}(v)$ and v are reachable from each other, and therefore $\tilde{C}(v) \subseteq A$. And
650 if we had $A \setminus \tilde{C}(v) \neq \emptyset$ then, by strong connectivity of A , there would be a vertex $z \in A \setminus \tilde{C}(v)$
651 with an out-neighbor u in $\tilde{C}(v)$. By the choice of v , this z is still ACY-dormant, so it has
652 not yet done any transmissions on ACY-frequencies. This means that $z \in N^-(u) \setminus \tilde{N}_{\text{ACY}}^-(u)$
653 and $z \notin \tilde{C}(u)$, so Test 3 would fail for this u , and v could not terminate its execution of the
654 SCC-subroutine — contradiction. We thus conclude that $\tilde{C}(v) = A$, as claimed.

655 Next, we claim that the two runs of $\text{SCCGOSSIP}_{j'}$ in A in j' -frame s did not experience
656 any interference from nodes outside A . Consider any $u \in A$. By the choice of v , the nodes in
657 A are not ACY-active before j' -frame s , and therefore $\tilde{N}_{\text{ACY}}^-(u) \cap A = \emptyset$. Also, since u passed
658 v 's Test 3, we have $N^-(u) \setminus \tilde{N}_{\text{ACY}}^-(u) \subseteq A$. This implies that $\tilde{N}_{\text{ACY}}^-(u) = N_{\text{ACY}}^-(u)$. Therefore
659 all nodes $z \in N_{\text{ACY}}^-(u)$, since they transmitted on ACY-frequencies before j' -frame s , must
660 be either ACY-active or ACY-expired throughout j' -frame s , so they do not transmit at
661 SCC-frequencies. This shows that there is no interference from outside of A during j' -frame
662 s , as claimed.


663 With no outside interference, both the exploration and dissemination runs of $\text{SCCGOSSIP}_{j'}$
664 in A will complete successfully for *all* nodes $v \in A$. This implies (ii) and (iii). As for (iv),
665 it now follows from the bounds we established earlier: $\alpha_{\text{ACY}}(A) = \alpha_{\text{ACY}}(v) \leq \psi_{j,r+1} \leq$
666 $\rho_{\text{SCC}}(A) + 4 \cdot T_{\text{SCC}}(j)$. ◀

667 **Running time.** Next, we estimate the running time. The argument follows the reasoning
668 in Section 3, but now we need to account for the contribution of the SCC-subroutine. To this
669 end, we apply Lemma 3, which gives us that for an sc-component A of size n_A and for $j =$
670 $\lceil \log n_A \rceil$ we have $\alpha_{\text{ACY}}(A) - \rho_{\text{SCC}}(A) \leq 4 \cdot T_{\text{SCC}}(j) = \tilde{O}(n_A^{4/3})$. The quantity $\alpha_{\text{ACY}}(A) - \rho_{\text{SCC}}(A)$
671 represents the contribution of A to slowing down the ACY-subroutine. Accumulated over
672 all sc-components, this slowdown works out to be $\tilde{O}(n^{4/3})$. On the other hand, the ACY-
673 subroutine itself, which simply executes ACYGATHER , contributes the total of $O(n^{1.5} \log^3 n)$
674 to the running time. Putting it all together, we obtain the $O(n^{1.5} \log^3 n)$ upper bound on
675 the running time of Algorithm ARBGATHER .

676 To make this argument more precise, we extend the definition of a critical path that was
677 introduced in Section 3. In this section, the *critical path* (see Figure 7) will be defined as a
678 sequence of nodes $P = (v_0, w_0, v_1, w_1, \dots, v_p, w_p = t)$ determined as follows:

- 679 ■ For any $a = p, p-1, \dots, 0$, suppose that w_a has already been defined, and let $C_a = C(w_a)$.
680 If $\text{Anc}(C_a) \neq \emptyset$, then let $v_a \in C_a$ be the node for which $\rho_{\text{SCC}}(v_a) = \rho_{\text{SCC}}(C_a)$. In other
681 words, v_a is the node in C_a for which $\rho_{\text{SCC}}(v_a)$ is maximum. (It could happen that
682 $v_a = w_a$.) On the other hand, if $\text{Anc}(C_a) = \emptyset$ (that is, C_a is a source sc-component),
683 then $a = 0$ and $v_0 \in C_0$ is arbitrary; for example we can take $v_0 = w_0$.
- 684 ■ For any $a = p-1, p-2, \dots, 0$, suppose that v_{a+1} has already been defined. Then w_a is
685 the node in $N_{\text{ACY}}^-(v_{a+1})$ for which $\rho_{\text{SCC}}(v_{a+1}) = \text{rus}_{w_a \rightarrow v_{a+1}}^1$. In other words, w_a is the
686 in-neighbor of v_{a+1} outside of C_{a+1} whose message was received last by v_{a+1} .

687 Denote by T the running time of protocol ARBGATHER on G . For $a = 0, 1, \dots, p$, let also
688 $n_a = |C_a|$ and $j_a = \lceil \log n_a \rceil$. By Lemma 3(iii), at time $\alpha_{\text{ACY}}(C_p)$ all rumors from G will
689 be collected by t ; thus $T \leq \alpha_{\text{ACY}}(C_p)$. By definition, $\rho_{\text{SCC}}(C_0) = 0$. So we can bound T as



■ **Figure 7** Illustration of the time analysis for arbitrary digraphs.

690 follows

$$\begin{aligned}
 691 \quad T &\leq \alpha_{\text{ACY}}(C_p) - \rho_{\text{SCC}}(C_0) \\
 692 \quad &= \sum_{a=0}^p [\alpha_{\text{ACY}}(C_a) - \rho_{\text{SCC}}(C_a)] + \sum_{a=0}^{p-1} [\rho_{\text{SCC}}(C_{a+1}) - \alpha_{\text{ACY}}(C_a)]. \\
 693
 \end{aligned}$$

694 We estimate the two sums separately. From Lemma 3(iv) we have $\alpha_{\text{ACY}}(C_a) \leq \rho_{\text{SCC}}(C_a) +$
 695 $4 \cdot T_{\text{SCC}}(j_a)$, so the first sum is

$$\begin{aligned}
 696 \quad \sum_{a=0}^p [\alpha_{\text{ACY}}(C_a) - \rho_{\text{SCC}}(C_a)] &\leq 4 \cdot \sum_{a=0}^p T_{\text{SCC}}(j_a) \\
 697 \quad &= 4 \cdot \sum_{a=0}^p \tilde{O}(n_a^{4/3}) = \tilde{O}(n^{4/3}), \\
 698
 \end{aligned}$$

699 because $\sum_{a=0}^p n_a \leq n$. To estimate the second sum, by the definition of v_{a+1} and w_a and
 700 by Lemma 3(ii) we have

$$701 \quad \rho_{\text{SCC}}(C_{a+1}) - \alpha_{\text{ACY}}(C_a) = \rho_{\text{SCC}}(v_{a+1}) - \alpha_{\text{ACY}}(w_a) = rws_{w_a \rightarrow v_{a+1}}^1 - \alpha_{\text{ACY}}(w_a).$$

702 We can now apply the charging argument identical to that in Section 3, namely we charge the
 703 delay $rws_{w_a \rightarrow v_{a+1}}^1 - \alpha_{\text{ACY}}(w_a)$ to stage index increments. (Specifically, these will be charged
 704 to the stage-index increments in the set U of in-neighbors of v_{a+1} , defined analogously as in
 705 the proof of Lemma 1.) This will give us a bound of $O(n^{1.5} \log^3 n)$ on the second sum. We
 706 thus obtain the main result of this paper:

707 ► **Theorem 4.** *Let G be an arbitrary digraph with n vertices and a designated target node*
 708 *reachable from all other nodes. Algorithm ARBGATHER completes information gathering in*
 709 *G in time $O(n^{1.5} \log^3 n)$.*

710 5 $\tilde{O}(n)$ -Time Protocol With Acknowledgements for Acyclic Graphs

711 We now consider the problem of gathering in acyclic graphs with a weak form of acknow-
 712 ledgment of transmission success. In this model, following each transmission from a node
 713 v , v receives a single bit of information indicating whether at least one node successfully
 714 received that transmission. Thus v does not learn which specific node, or how many nodes in
 715 total, received its transmission. Our main goal in this section is to show that this single bit

716 is enough to allow for gathering to be performed in time $O(n \log^2 n)$ on acyclic graphs with
 717 n vertices. The key idea here will be that nodes which have successfully transmitted can at
 718 least temporarily stop transmitting, making it easier for other nodes to succeed. In order
 719 for this to work, though, we need to guarantee that successful transmissions are occurring
 720 at a reasonable rate. The following combinatorial object will be our main tool for this.

721 We say that a collection $(S_0, S_1, \dots, S_{b-1})$ of label sets forms a (n, k) -half-selector if for
 722 every $X \subseteq [n]$ with $|X| \leq k$ there are at least $|X|/2$ choices of $x \in X$ for which there is
 723 an index i with $S_i \cap X = \{x\}$. (This is in contrast to strong selectors where we want this
 724 property to hold for *every* choice of x). It is a consequence of Lemma 1 in [6] that for every
 725 n and $k \leq n$ there exists an (n, k) -half-selector of size $O(k \log n)$.

726 For all $j = 0, 1, \dots, \log n$, by 2^j -HALFSELECT = $(S_0^j, S_1^j, \dots, S_{b_j-1}^j)$ we will denote an
 727 $(n, 2^j)$ -half-selector of size $b_j = O(2^j \log n)$. Without any loss of generality we can also
 728 assume that $b_{j+1} = 2b_j$ for all $j \leq \log n - 1$, implying that $b_j = \gamma 2^j \log n$ for some absolute
 729 constant γ .

730 As in the previous sections, our algorithm is presented in the relaxed radio-network
 731 model that satisfies Assumptions (MFC), (SRT) and (INN) from Section 2. The algorithm
 732 will use $\kappa = \log n + 2$ frequencies. The intuition here is that for $0 \leq j \leq \kappa - 2$ frequency j
 733 will be used to handle potential interferences involving at most 2^j vertices.

734 **Algorithm ACYGATHERACK.** At any given time step, a node can be either *dormant* or
 735 *active*. Initially the source nodes (with no in-neighbors) will be active and the remaining
 736 nodes will be dormant. Any active node transmits according to 2^j -HALFSELECT on each
 737 frequency $j = 0, 1, \dots, \kappa - 2$, and according to ROUNDROBIN on frequency $\kappa - 1$. An active
 738 node which receives an acknowledgement of a successful transmission moves to the dormant
 739 state, and a dormant node which receives a transmission becomes active.

740 We remind the reader that when we write that “an active node transmits according to
 741 2^j -HALFSELECT”, we mean that at a step τ this node, say v , will transmit if and only if
 742 $v \in S_{\tau \bmod b_j}^j$. (Thus, on any given frequency, at any given time step all nodes use the same
 743 transmission set to determine if they are supposed to transmit or not.) The meaning of
 744 transmitting according to ROUNDROBIN is analogous.

745 Observe that, unlike in the previous algorithms, it is now possible for a node to become
 746 active multiple times during the process as it continually receives new messages. Also, the
 747 target node, once it receives the first message, remains in the active state forever.

748 **Analysis.** We now prove correctness of Algorithm ACYGATHERACK and establish an upper
 749 bound of $O(n \log n)$ on its running time. To this end, we show that after at most $O(n \log n)$
 750 steps each rumor will reach the target node t .

751 Our proof uses again the layered structure of G introduced in Section 3. Recall that,
 752 for a node v , $\delta(v)$ denotes the length of the *longest* directed path from some source node to
 753 v , and that $\delta^* = \max_{v \in G} \delta(v) = \delta(t)$. For $i = 0, 1, \dots, \delta^*$, B_i denotes the set of nodes with
 754 $\delta(v) = i$. (See Figure 3 for illustration.)

755 Let $\tau_i = 4\gamma \sum_{p < i} |B_p| \log n$ for all $i = 0, 1, \dots, \delta^*$. (In particular, $\tau_0 = 0$.) Our argument
 756 is based on the lemma below:

757 **► Lemma 5.** *The following two properties hold for every $i = 0, 1, \dots, \delta^*$:*

- 758 (i) *All nodes in $\bigcup_{p < i} B_p$ remain dormant at all times after τ_i (inclusive).*
 759 (ii) *At time τ_i each rumor is in some active node in $\bigcup_{p \geq i} B_p$.*

760 **Proof.** We establish Lemma 5 inductively. Both parts (i) and (ii) of the claim hold vacuously
 761 for $i = 0$. Now we assume that parts (i) and (ii) hold for some i and we consider the

762 computation of the nodes in $\bigcup_{p \geq i} B_i$, beginning at time τ_i . The proof involves three claims
 763 below that capture fundamental properties of Algorithm ACYGATHERACK.

764 ► **Claim 4.** If some rumor is in an active node in $\bigcup_{p \geq i} B_p$ at some time step τ , then it will
 765 be in some active node in $\bigcup_{p \geq i} B_p$ at any time step $\tau' \geq \tau$.

766 To justify Claim 4 note that if some rumor is in an active node $w \in \bigcup_{p \geq i} B_p$ then w
 767 remains active until it successfully transmits its message, which includes this rumor, to at
 768 least one of its out-neighbors. And once an out-neighbor of w , say u , receives this message
 769 from w , it gets immediately activated, if it's not already active. Since $u \in \bigcup_{p \geq i} B_p$ as well,
 770 Claim 4 follows.

771 ► **Claim 5.** Let A be the set of nodes in B_i that are active at time τ_i . Then (i) Any node
 772 in $B_i \setminus A$ remains dormant forever, and (ii) any node in A becomes permanently dormant
 773 right after its first successful transmission.

774 The proof of Claim 5 is quite simple: Each node in A that successfully transmits is
 775 immediately made dormant by the algorithm. The nodes in B_i will not receive any rumors
 776 after time τ_i since, by the inductive hypothesis (i) and Observation 1, none of their in-
 777 neighbors will be active after time τ_i . So any node in B_i that is dormant at some time
 778 $\tau \geq \tau_i$ will remain dormant forever.

779 ► **Claim 6.** Each node in A will succeed in transmitting its message and become dormant
 780 before time τ_{i+1} .

781 To prove Claim 6, choose j such that $2^{j-1} < |B_i| \leq 2^j$. Trivially, $|A| \leq |B_i| \leq 2^j$.
 782 Since the algorithm runs 2^j -HALFSELECT on frequency j , at least $|A|/2$ nodes in A will
 783 have a time step in the interval $[\tau_i, \tau_i + b_j)$ when they will successfully transmit, at which
 784 point they become permanently dormant, by Claim 5. Thus, if A' is the set of nodes in B_i
 785 that are active at time $\tau_i + b_j$, then $|A'| \leq |A|/2 \leq 2^{j-1}$. Next, we look at time interval
 786 $[\tau_i + b_j, \tau_i + b_j + b_{j-1})$. Since the algorithm runs 2^{j-1} -HALFSELECT on frequency $j-1$,
 787 using the same argument, if A'' is the set of nodes in B_i active at time $\tau_i + b_j + b_{j-1}$ then
 788 $|A''| \leq |A'|/2 \leq 2^{j-2}$. Continuing inductively, all the nodes in A will succeed (and become
 789 dormant) no later than at time

$$\begin{aligned}
 790 \quad \tau_i + \sum_{q=0}^j b_q &= \tau_i + \gamma(\sum_{q=0}^j 2^q) \log n \\
 791 \quad &< \tau_i + \gamma 2^{j+1} \log n \\
 792 \quad &< \tau_i + 4\gamma |B_i| \log n = \tau_{i+1}, \\
 793
 \end{aligned}$$

794 completing the proof of Claim 6.

795 By Claims 5 and 6, all nodes in $\bigcup_{p < i+1} B_p$ will be permanently dormant starting no later
 796 than at time τ_{i+1} . Each successful transmission from B_i arrives at a node in $\bigcup_{p \geq i+1} B_p$, so
 797 at time τ_{i+1} all rumors are in some active nodes in $\bigcup_{p \geq i+1} B_p$, by Claims 6 and 4. This
 798 completes the inductive step and the proof of Lemma 5. ◀

799 Taking $i = \delta^*$ in Lemma 5, we obtain that the algorithm delivers all rumors to t in at most
 800 $\tau_{\delta^*} = 4\gamma \sum_{p=0}^{\delta^*-1} |B_p| \log n \leq 4\gamma n \log n$ time steps. This proves correctness and establishes an
 801 $O(n \log n)$ bound on the running time in our relaxed communication model with κ frequencies
 802 and assumptions (MFC), (SRT), and (INN). Since $\kappa = O(\log n)$, as explained in Section 2,
 803 this implies an $O(n \log^2 n)$ -time bound on the running time of Algorithm ACYGATHERACK
 804 in the standard single-frequency model. This is summarized in the theorem below.

805 ► **Theorem 6.** *Let G be an acyclic directed graph with n vertices and a designated target*
806 *node reachable from all other nodes. Using acknowledgements of successful transmissions,*
807 *Algorithm ACYGATHERACK completes information gathering in G in time $O(n \log^2 n)$.*

808 6 Final Comments

809 In this paper we provided an $\tilde{O}(n^{1.5})$ -time protocol for information gathering in ad-hoc radio
810 networks, improving the trivial upper bound of $O(n^2)$. For the model with transmissions
811 acknowledgments we gave a $\tilde{O}(n)$ -time protocol for acyclic digraphs.

812 We hope that some ideas behind our algorithms will lead to further improvements, and
813 perhaps find applications to other communication dissemination problems in ad-hoc radio
814 networks. One idea that is particularly promising is the amortization technique in Section 3,
815 where a failure of a node in transmitting its message is charged to stage-index increments
816 of the interfering nodes. Another idea is the technique for integrating a gossiping protocol
817 (applicable only to strongly connected digraphs) with an information gathering protocol for
818 acyclic digraphs, to obtain an information gathering protocol for arbitrary digraphs. Using
819 this technique, improving the upper bound to below $\tilde{O}(n^{1.5})$ should be possible by designing
820 an appropriate protocol that beats the $\tilde{O}(n^{1.5})$ bound for acyclic graphs.

821 Several open problems remain. The two most intriguing problems are about the time
822 complexity of gossiping and information gathering, as for both problems the best known
823 lower bounds are only $\Omega(n \log n)$, the same as for broadcasting.

824 There are a number of other natural questions about information gathering protocols for
825 radio networks that deserve study. For example, how does the complexity of information
826 gathering depend on the graph diameter D and maximum degree Δ ? Some initial work
827 in this direction was done in [28], where refined bounds for the case of trees were given.
828 Another natural direction of research would be to analyze the complexity of information
829 gathering when the graph topology is known. This has been well studied for broadcasting
830 and gossiping – see, for example, the survey in [15].

831 **Acknowledgements.** We are very grateful to the anonymous reviewers for pointing out
832 several deficiencies in the original submission and for insightful comments that helped us
833 improve the quality of presentation.

834 — References —

- 835 1 Noga Alon, Amotz Bar-Noy, Nathan Linial, and David Peleg. A lower bound for radio
836 broadcast. *J. Comput. Syst. Sci.*, 43(2):290–298, 1991.
- 837 2 Danilo Bruschi and Massimiliano Del Pinto. Lower bounds for the broadcast problem in
838 mobile radio networks. *Distributed Computing*, 10(3):129–135, 1997.
- 839 3 Bogdan S. Chlebus, Leszek Gąsieniec, Alan Gibbons, Andrzej Pelc, and Wojciech Rytter.
840 Deterministic broadcasting in ad hoc radio networks. *Distributed Computing*, 15(1):27–38,
841 2002.
- 842 4 Marek Chrobak and Kevin P. Costello. Faster information gathering in ad-hoc radio tree
843 networks. *Algorithmica*, 80(3):1013–1040, 2018.
- 844 5 Marek Chrobak, Kevin P. Costello, Leszek Gąsieniec, and Dariusz R. Kowalski. Information
845 gathering in ad-hoc radio networks with tree topology. *Information and Computation*,
846 258:1–27, 2018.
- 847 6 Marek Chrobak, Leszek Gąsieniec, and Wojciech Rytter. Fast broadcasting and gossiping
848 in radio networks. *Journal of Algorithms*, 43(2):177–189, 2002.

- 849 **7** Marek Chrobak, Leszek Gasieniec, and Wojciech Rytter. A randomized algorithm for
850 gossiping in radio networks. *Networks*, 43(2):119–124, 2004.
- 851 **8** Andrea E. F. Clementi, Angelo Monti, and Riccardo Silvestri. Selective families, superim-
852 posed codes, and broadcasting on unknown radio networks. In *Proc. 12th Annual Sym-
853 posium on Discrete Algorithms (SODA'01)*, pages 709–718, 2001.
- 854 **9** Andrea E. F. Clementi, Angelo Monti, and Riccardo Silvestri. Distributed broadcast in
855 radio networks of unknown topology. *Theor. Comput. Sci.*, 302(1-3):337–364, 2003.
- 856 **10** Artur Czumaj and Peter Davies. Faster deterministic communication in radio networks.
857 In *Proc. 43rd International Colloquium on Automata, Languages, and Programming (IC-
858 ALP'16)*, pages 139:1–139:14, 2016.
- 859 **11** Artur Czumaj and Wojciech Rytter. Broadcasting algorithms in radio networks with un-
860 known topology. *Journal of Algorithms*, 60(2):115 – 143, 2006.
- 861 **12** Gianluca De Marco. Distributed broadcast in unknown radio networks. In *Proc. 19th
862 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'08)*, pages 208–217, 2008.
- 863 **13** Gianluca De Marco. Distributed broadcast in unknown radio networks. *SIAM J. Comput.*,
864 39(6):2162–2175, 2010.
- 865 **14** Antonio Fernández Anta, Miguel A. Mosteiro, and Jorge Ramón Muñoz. Unbounded
866 contention resolution in multiple-access channels. *Algorithmica*, 67(3):295–314, 2013.
- 867 **15** Leszek Gasieniec. On efficient gossiping in radio networks. In *Proc. 16th Int. Colloquium on
868 Structural Information and Communication Complexity (SIROCCO'09)*, pages 2–14, 2009.
- 869 **16** Leszek Gasieniec. Deterministic broadcasting in radio networks. In *Encyclopedia of Al-
870 gorithms*, pages 529–530. Springer US, 2016.
- 871 **17** Leszek Gasieniec and Igor Potapov. Gossiping with unit messages in known radio networks.
872 In *Proc. 2nd IFIP International Conference on Theoretical Computer Science (TCS'02)*,
873 pages 193–205, 2002.
- 874 **18** Leszek Gasieniec, Tomasz Radzik, and Qin Xin. Faster deterministic gossiping in directed
875 ad hoc radio networks. In *Proc. Scandinavian Workshop on Algorithm Theory (SWAT'04)*,
876 pages 397–407, 2004.
- 877 **19** Alon Itai. Randomized broadcasting in radio networks. In *Encyclopedia of Algorithms*,
878 pages 1734–1738. Springer US, 2016.
- 879 **20** Tomasz Jurdzinski and Dariusz R. Kowalski. Wake-up problem in multi-hop radio networks.
880 In *Encyclopedia of Algorithms*, pages 2352–2354. Springer, 2016.
- 881 **21** Dariusz R. Kowalski and Andrzej Pelc. Faster deterministic broadcasting in ad hoc radio
882 networks. *SIAM J. Discrete Math.*, 18(2):332–346, 2004.
- 883 **22** Eyal Kushilevitz and Yishay Mansour. An $\Omega(D \log(N/D))$ lower bound for broadcast in
884 radio networks. *SIAM J. Comput.*, 27(3):702–712, 1998.
- 885 **23** Ding Liu and Manoj Prabhakaran. On randomized broadcasting and gossiping in radio
886 networks. In *Proc. 8th Annual Int. Conference on Computing and Combinatorics (CO-
887 COON'02)*, pages 340–349, 2002.
- 888 **24** Gianluca De Marco and Dariusz R. Kowalski. Contention resolution in a non-synchronized
889 multiple access channel. In *Proc. of the 27th Int. Symposium on Parallel Distributed Pro-
890 cessing (IPDPS)*, pages 525–533, 2013.
- 891 **25** Gianluca De Marco and Dariusz R. Kowalski. Fast nonadaptive deterministic algorithm
892 for conflict resolution in a dynamic multiple-access channel. *SIAM Journal on Computing*,
893 44(3):868–888, 2015.
- 894 **26** Gianluca De Marco and Dariusz R. Kowalski. Contention resolution in a non-synchronized
895 multiple access channel. *Theor. Comput. Sci.*, 689:1–13, 2017.
- 896 **27** David Peleg. Time-efficient broadcasting in radio networks: A review. In *Proc. Distributed
897 Computing and Internet Technology, 4th International Conference, ICDCIT'07*, pages 1–
898 18, 2007.

- 899 28 Parker Williams. *Information Gathering on Bounded Degree Trees and Properties of Random Matrices*. PhD thesis, University of California, Riverside, 2017.
- 900
- 901 29 Ying Xu. An $O(n^{1.5})$ deterministic gossiping algorithm for radio networks. *Algorithmica*, 36(1):93–96, 2003.
- 902

903 **A** About Protocol SccGossip

904 The analysis of Algorithm SCCGOSSIP in [18] (called GOSSIP2 in that paper) assumes that
905 N is bounded polynomially in \bar{n} . Since in this case $\log N = O(\log \bar{n})$, the bound on the
906 running time of their algorithm (see Theorem 2 of that paper) does not explicitly separate
907 the dependence on \bar{n} and N . We now explain how the bound given in our Property (scc)
908 follows from the analysis in [18]. In essence, the running time depends on N in three ways:
909 as the label range of selectors, as the range of binary search, or as the range of a doubling (or
910 halving) process. All these three contributions are only logarithmic in N . (In fact, in some
911 cases the range of binary search or doubling can be reduced to \bar{n} , but this is not relevant to
912 our application.) A more detailed explanation follows.

913 The algorithm in [18] uses a broadcasting algorithm from [6] in their procedure DISPERSE()
914 (Section 3). In [6] it is assumed that the label set is $[\bar{n}]$ and the running time is given as
915 $O(\bar{n} \log^2 \bar{n})$. However, all this algorithm does is to repeatedly run selectors, so for the label
916 range $[N]$ its running time can be expressed as $O(\bar{n} \log^2 N)$. Procedure DISPERSE() in [18]
917 also involves a binary search, introducing another factor of $O(\log N)$. Thus the running
918 time of DISPERSE() (Lemma 2 in [18]) can be restated as $O((\bar{n}/x) + r)\bar{n} \log^3 N$. Phase I of
919 Algorithm GOSSIP2 in Section 3.2 in [18] involves a halving process that executes a selector
920 in each iteration. Phase II involves $O(\log \bar{n})$ iterations, each executing a selector. For this
921 reason, only logarithmic factors in the running time will depend on N .