Tight Bounds for Deterministic h-Shot Broadcast in Ad-Hoc Directed Radio Networks

Aris Pagourtzis[1]
School of ECE, National Technical University of Athens, Athens, Greece
pagour@cs.ntua.gr
https://orcid.org/0000-0002-6220-3722

Tomasz Radzik[2]
Department of Informatics, King's College London, London, the United Kingdom
tomasz.radzik@kcl.ac.uk
https://orcid.org/0000-0002-7776-5461

Abstract

We consider the classical broadcast problem in ad-hoc (that is, unknown topology) directed radio networks with no collision detection, under the additional assumption that at most h transmissions (shots) are available per node. We focus on adaptive deterministic protocols for small values of h. We provide asymptotically matching lower and upper bounds for the cases $h = 2$ and $h = 3$. While for $h = 2$ our bound is quadratic, similar to the bound obtained for oblivious protocols, for $h = 3$ we prove a sub-quadratic bound of $\Theta(n^2 \log\log n/\log n)$, where n is the number of nodes in the network. The latter is the first result showing an adaptive algorithm which is asymptotically faster than oblivious h-shot broadcast protocols, for which a tight quadratic bound is known for every constant h. Our upper bound for $h = 3$ is constructive, making use of constructions of graphs with large girth. We also show an improved upper bound of $O(n^{1+\alpha/\sqrt{h}})$ for $h \geq 4$, where α is an absolute constant independent of h. Our upper bound for $h \geq 4$ is non-constructive.

2012 ACM Subject Classification Theory of computation → Distributed algorithms, Mathematics of computing → Extremal graph theory

Keywords and phrases Ad-hoc radio networks, wireless networks, deterministic broadcast, adaptive protocols, limited transmissions

Digital Object Identifier 10.4230/LIPIcs.MFCS.2018.80

1 Introduction

1.1 Model of broadcast with limited transmissions per node

In this paper a *transmission network* is a directed graph $G = (V, E)$ with the set of nodes $V = \{0, 1, \ldots, n-1\}$, where node 0 is the *source node*, denoted also by s, and all other nodes are reachable from this node. Initially each node knows only its identifier and the size n of the network. The source node knows also the *message*, which is to be broadcast to all other nodes. Let $\mathcal{G} \equiv \mathcal{G}^{(n)}$ denote the family of all transmission networks of size n.

We consider the following model of h-*shot broadcast*. Nodes of the network transmit in globally synchronized steps (counted from 1), with each node transmitting in at most h steps. If a node v transmits in a given step, then each node w such that $(v, w) \in E$ receives the

[1] Research mainly conducted while this author was visiting the Department of Informatics, King's College London, on sabbatical leave from the National Technical University of Athens.
[2] Research supported by EPSRC grant EP/M005038/1.

43rd International Symposium on Mathematical Foundations of Computer Science (MFCS 2018).
Editors: Igor Potapov, Paul Spirakis, and James Worrell; Article No. 80; pp. 80:1–80:13
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

transmitted message, unless a *collision* occurs at node w, that is, unless there is another edge (v', w), $v' \neq v$, with node v' transmitting in the same step. We assume that there is no collision detection: a node w cannot distinguish between no transmission from any of the neighbouring nodes and simultaneous transmissions by two or more neighbouring nodes. The only node transmitting in step 1 is the source node 0 and a node can transmit in the current step $t \geq 2$ only if it has already received the message in previous steps.

Most of the research on communication protocols for various models of radio networks has been concerned with minimizing the number of steps, without putting constraints on the number of transmissions by individual nodes or the total number of transmissions by all nodes. Limiting the maximum number of transmissions per node has received somewhat less attention, especially in the context of ad-hoc (that is, unknown) networks. This objective, however, may be important in practice, since it may mean limiting the maximum energy usage per node to keep all nodes alive for as long as possible.

An h-shot broadcast protocol can be viewed as a function $\Pi \equiv \Pi_n$ which for any node v, a time step $t \geq 1$, and the knowledge κ gathered by node v in steps $1, 2, \ldots, t-1$, tells node v whether it transmits in step t. The protocol has to ensure that, within all constraints of the model, for each transmission network $G \in \mathcal{G}$, all nodes eventually receive the message, that is, broadcast is always eventually completed. The design objective is to keep the worst-case completion time as small as possible.

An *oblivious h-shot protocol* is defined by a sequence of *transmission sets* $S_1 = \{0\}, S_2, S_3, \ldots$, which are subsets of the node set V. Once a node v receives the message in step t, it wakes up and transmits in the first h steps $\tau_i \geq t+1, 1 \leq i \leq h$, such that $v \in S_{\tau_i}$. The source node 0 is considered awake at time 0 and transmitting at step 1. (We remark that slightly different definitions of obliviousness may be used in other variants of radio network models.)

In a general *(adaptive) h-shot protocol*, nodes can take into account information which they have received in earlier steps when they decide whether to transmit in the current step. We do not put any limits on how much information can be transmitted in one step or stored in one node. In fact, for our lower bounds we assume that during a successful transmission from a node v to a node w, all knowledge accumulated so far by node v is transmitted to node w and is added to w's knowledge. We remark though that the (adaptive) protocols for our upper bounds include in the transmissions only the source message and the current count of step. They achieve a speed-up over oblivious protocols by using the current count of steps in a more subtle way.

1.2 Our results

We study adaptive deterministic protocols for h-shot broadcast (note that the term 'k-shot broadcasting' has been used in some literature for the same notion). We focus on small values of h and provide asymptotically matching lower and upper bounds on the (worst-case) number of steps for the cases $h = 2$ and $h = 3$, as well as improved upper bounds for larger values of h.

In particular, for $h = 2$ we provide a quadratic lower bound of $n^2/8 - O(n)$, showing that adaptive 2-shot broadcast protocols are not (asymptotically) faster than oblivious 2-shot protocols. On the other hand, for $h = 3$ we prove a sub-quadratic bound of $\Theta(n^2 \log\log n / \log n)$. To the best of our knowledge this is the first result showing an adaptive h-shot protocol which is asymptotically faster than oblivious h-shot protocols. For oblivious protocols a tight quadratic bound has been shown in [14] for every constant h. Our proof of existence of a $O(n^2 \log\log n / \log n)$-step 3-shot broadcast protocol is constructive, making use of constructions of graphs with large girth. The *girth of a graph* is the length of a shortest cycle.

Our improved upper bounds for $h \geq 4$ include a bound of $O(n^{1+\alpha/\sqrt{h}})$, where h is constant or grows (slowly) with n and α is an absolute constant independent of h. Our upper bounds for $h \geq 4$ are non-constructive and are based on hyper-graphs without small 2-covers. We give the precise definition of 2-covers in hyper-graphs in Section 4, noting here only that this notion can be viewed as a generalization of the notion of cycles in graphs.

1.3 Related previous work

Radio broadcasting with unlimited number of shots was first introduced by Chlamtac and Kutten [5] and has been extensively studied ever since. The first protocol, given by Bar-Yehuda, Goldreich and Itai [1], was randomized and worked in $O(D \log n + \log^2 n)$ expected time, where D is the diameter of the graph and n the number of nodes. Improved randomized protocols were later proposed in [10, 15] yielding a tight upper bound of $O(D \log(n/D)+\log^2 n)$ steps.

Deterministic radio broadcasting attracted much attention in the last two decades. Brusci and Del Pinto [4] proved a lower bound of $\Omega(D \log n)$ for undirected networks, which was subsequently improved for directed networks to $\Omega(n \log D)$ by Clementi et al. [9] and for undirected networks to $\Omega((n \log n)/\log(n/D))$ by Kowalski and Pelc [15]. The *round-robin* protocol, in which node i is the only node transmitting in steps $i + 1 + qn$, for each $q \geq 1$, gives a trivial $O(n^2)$ upper bound on deterministic broadcast. Chlebus et al. [6] presented the first sub-quadratic protocol of $O(n^{11/6})$ time complexity. The upper bound was then improved to $O(n^{5/3} \log^3 n)$ by De Marco and Pelc [17] and further by Chlebus et al. [7], who showed an $O(n^{3/2})$-time algorithm. Chrobak, Gąsieniec and Rytter [8] gave an $O(n \log^2 n)$ non-constructive protocol and De Marco [11] proved the best currently known upper bound of $O(n \log n \log \log n)$, again in a non-constructive manner.

Better upper bounds are known for undirected networks. Chlebus et al. [6] proposed a deterministic $O(n)$-time broadcasting algorithm, assuming *spontaneous wake-up* (that is, allowing the nodes to transmit before receiving the source message, learning that way the topology of the network). An optimal $O(n \log n)$-time broadcasting algorithm for undirected networks with non-spontaneous wake-up was given by Kowalski and Pelc [15].

Broadcasting with a limited number of shots ("h-shot broadcasting") in known-topology undirected networks was first studied by Gąsieniec et al. [12], who showed a lower bound of $D + \Omega((n - D)^{1/(2h)})$ and a randomized protocol which works in $D + O(hn^{1/(h-2)} \log^2 n)$ steps and has high probability of completing the broadcast. These lower and upper bounds were improved for the same setting (undirected known networks) by Kantor and Peleg [13] to $D + \Omega(h \cdot (n - D)^{1/2h})$ and $D + O(hn^{1/2h}) \log^{2+1/h} n)$, respectively. They also presented the first randomized h-shot broadcasting protocols for *unknown* undirected networks, which work in $O((D + \min\{Dh, \log n\})n^{1/(h-1)} \log n)$ steps for $h \geq 2$ and in $O(Dn^2 \log n)$ steps for $h = 1$. Still in the same setting, Berenbrink et al. [2] proposed, among other results, a randomized algorithm with optimal broadcasting time $O(D \log(n/D) + \log^2 n)$ that uses an expected number of $O(\log^2 n/ \log(n/D)$ transmissions per node.

The first work to address deterministic h-shot broadcasting in directed *ad hoc* radio networks is due to Karmakar et al. [14], who proved a lower bound of $\Omega(n^2/h)$ for oblivious protocols and a matching upper bound of $O(n^2/h)$ for each $h \leq \sqrt{n}$, as well as an upper bound of $O(n^{3/2})$ for $h > \sqrt{n}$. They also presented a lower bound of $\Omega\left(n^{1+1/h}\right)$ for adaptive broadcasting protocols, leaving open the question whether there are upper bounds for adaptive h-shot broadcast which are better than the $O(n^2/h)$ bound achieved by oblivious protocols.

2 Lower bounds

2.1 Layered networks

We show lower bounds using the following *layered networks*. We assume $n \geq 4$, and in addition to the source node $s = 0$, we also distinguish the node $d = n - 1$ as the "target" of the broadcast. Node d will be the last node of a layered network to receive the message. We derive lower bounds on the number of steps needed by a broadcast protocol to deliver the message from node s to node d in the worst case.

Consider a partition $L_0, L_1, \ldots, L_k$ of the set of nodes V into $k \geq 2$ sets called *layers*, such that $L_0 = \{0\}$, $L_k = \{n-1\}$ and $L_i \neq \emptyset$ for $0 \leq i \leq k$. These layers define the following acyclic broadcast network $G \equiv G(L_0, L_1, \ldots L_k)$. For each $0 \leq i \leq k-1$, the consecutive layers L_i and L_{i+1} are fully connected, that is, there is a directed edge from each node of L_i to each node of L_{i+1}, and there are no any other edges.

For any pairwise disjoint non-empty subsets $L_0, L_1, \ldots, L_j$ of $V \setminus \{n-1\}$, where $j \geq 1$ and $L_0 = \{0\}$, we denote by $\mathcal{G}_j \equiv \mathcal{G}_j(L_0, L_1, \ldots L_j)$ the family of all layered networks $G(L_0, L_1, \ldots L_j, L_{j+1} \ldots L_k)$, where $k > j$, $L_k = \{n-1\}$ and $L_{j+1} \ldots L_{k-1}$ are non-empty sets partitioning $V \setminus \left(\{n-1\} \cup \bigcup_{i=0}^{j} L_i\right)$. In other words, $\mathcal{G}_j$ is the family of all layered networks which have the same fixed initial layers $L_0, L_1, \ldots, L_j$. In particular, $\mathcal{G}_0$ is the family of all layered networks.

We use layered networks in order to show that for any given protocol, there is an assignment of nodes to layers which makes the progress of broadcast slow because of relatively long delays at each layer.

2.2 Conditional transmission sets

Let Π be any h-shot broadcast protocol for n-node networks and let $T_{\max}$ denote the maximum broadcast time of Π over all n-node networks. We define below *conditional transmission sets* for families of layered graphs described above.

Let $i \geq 1$ and consider the family of networks $\mathcal{G}_{i-1}(L_0, L_1, \ldots, L_{i-1})$ for some arbitrary layer sets $L_0 = \{0\}, L_1, L_2, \ldots, L_{i-1}$, such that $|\bigcup_{j=0}^{i-1} L_j| \leq n-3$. This bound implies that the target node and at least two other nodes are still outside of the fixed layers. Protocol Π behaves in exactly the same way on any network $G \in \mathcal{G}_{i-1}(L_0, L_1, \ldots, L_{i-1})$ until (and including) the step T_{i-1} when the source message leaves layer $i-1$ for the first time. That is, T_{i-1} is the first step when a unique node in L_{i-1} transmits, sending the message simultaneously to all nodes in the next layer. Note that $T_0 = 1$ and step T_{i-1} is uniquely determined by the sets $L_0, L_1, \ldots, L_{i-1}$. We select the next layer L_i from the set

$$U_i = V \setminus \left(\{n-1\} \cup \bigcup_{j=0}^{i-1} L_j\right),$$

trying to maximize the *weighted delay* $(T_i - T_{i-1})/|L_i|$ at this layer.

For $t \geq 1$, the *conditional transmission set* $S_t \subseteq U_i$ contains a node $v \in U_i$, if and only if, node v transmits at step $T_{i-1} + t$, if v is included in the layer L_i. Set S_t is well defined since for each network $G \in \mathcal{G}_{i-1}(L_0, L_1, \ldots, L_{i-1})$ with $v \in L_i$, node v transmits in exactly the same steps, irrespectively of how the other nodes in $U_i \setminus \{v\}$ are distributed among the layers L_j, $j \geq i$. This follows from the fact that a node in one layer gets information, directly or indirectly, only from nodes in previous layers.

Since we consider h-shot protocols, each node $v \in U_i$ belongs to at most h conditional transmission sets S_t. We assume w.l.o.g. that v transmits in exactly h steps $T_{i-1} + t$, so it belongs to exactly h conditional transmission sets. (If v belongs to $k < h$ conditional sets, then add v to $h - k$ sets S_τ for $\tau = T_{\max} - T_{i-1} + 1, \ldots, T_{\max} - T_{i-1} + h - k$. This may create new transmission collisions, but only after step $T_{\max}$, that is, after the completion of broadcast.) For convenience, if it is clear from the context that we are discussing the selection of nodes for the layer L_i, then we will refer to the (global) transmission step $T_{i-1} + t$ as simply the transmission step t (the t-th step after step T_{i-1}). Also, "conditional transmission sets" will be abbreviated to "transmission sets".

At least one of the transmission sets S_t must be a singleton, or otherwise the message would never reach the target node in the network $G(L_0, L_1, \ldots, L_{i-1}, U_i, \{n-1\})$, that is, when layer i contains all remaining nodes (other than the target node $n-1$). Let $\tau_1 \geq 1$ be the smallest index of a singleton transmission set. Let $S_{\tau_1} = \{v_1\}$ and we also use $t_0(v_1)$ and $S_0(v_1)$ to denote τ_1 and S_{τ_1}, respectively.

Applying the same argument to sets $S'_t = S_t \setminus \{v_1\}$, we observe that there must be a singleton also among these sets. Indeed, if each non-empty set S'_t, $t \geq 1$, had size at least 2, then the message would never reach the target node in the network $G(L_0, L_1, \ldots, L_{i-1}, U_i \setminus \{v_1\}, \{v_1\}, \{n-1\})$. Let $S'_{\tau_2} = \{v_2\}$ be the first singleton among sets S'_t, and let $t_0(v_2)$ and $S_0(v_2)$ denote τ_2 and S_{τ_2}, respectively. We note that $S_0(v_2)$ is equal to either $\{v_2\}$ or $\{v_1, v_2\}$ and step $t_0(v_2)$ can be before or after step $t_0(v_1)$.

Continuing this way, we put all nodes of U_i in a sequence $v_1, v_2, \ldots, v_u$, where $u = |U_i| \geq 2$, and associate with them distinct transmission steps $t_0(v_j)$ and transmission sets $S(v_j)$ such that:

$$S_0(v_1) = \{v_1\} \text{ and } S_0(v_j) \setminus \{v_1, v_2, \ldots, v_{j-1}\} = \{v_j\}, \text{ for } 2 \leq j \leq u.$$

Note that for each $1 \leq j \leq u$, we have $\bigcup_{l=1}^{j} S_0(v_l) = \{v_1, v_2, \ldots, v_j\}$.

By construction, for any two distinct nodes v and w in U_i, the steps $t_0(v)$ and $t_0(w)$ are also distinct. Thus for at least $\lceil u/2 \rceil$ nodes in U_i, we have $t_0(v) > \lfloor u/2 \rfloor$. We denote the set of these nodes by U'_i, that is,

$$U'_i = \{v \in U_i : t(v) > \lfloor u/2 \rfloor\},$$

and let $u' = |U'_i| \geq \lceil u/2 \rceil$.

For each node $v \in U'_i$, we have designated one of the v's (conditional) transmission steps as the step $t_0(v) \geq |U_i|/2$ and we have denoted the corresponding transmission set by $S_0(v)$. We now further denote by $t_1(v) < t_2(v) < \cdots < t_{h-1}(v)$ and by $S_1(v), S_2(v), \ldots, S_{h-1}(v)$ the other $h-1$ steps when node v transmits (if in L_i) and the transmission sets at those steps. While for any two distinct nodes v' and v'' in U'_i, $t_0(v') \neq t_0(v'')$, we may have $t_q(v') = t_r(v'')$ for some $1 \leq q, r \leq h-1$.

The general idea for forcing a large weighted delay at layer i is to try to select for this layer a relatively small number of nodes $x_1, x_2, \ldots, x_k$ from U'_i which have transmission conflicts at all transmission steps $t_l(x_j)$, for $1 \leq j \leq k$ and $1 \leq l \leq h-1$. That is, for each $1 \leq j \leq k$ and $1 \leq l \leq h-1$, there is $1 \leq a \leq k$, $a \neq j$ such that $\{x_j, x_a\} \subseteq S_l(x_j)$. If we manage to select such a layer, then the progress of broadcast from this layer will have to rely on one of the steps $t_0(x_1), t_0(x_2), \ldots, t_0(x_k)$, so the weighted delay will be at least $\min\{t_0(x_1), t_0(x_2), \ldots, t_0(x_k)\}/k \geq |U_i|/(2k)$. This will be the basic case in our lower-bound analysis.

2.3 Lower bound for 2-shot broadcast

We consider now the case when each node has only two transmissions, with a node $v \in U'_i$ transmitting in steps $t_0(v)$ and $t_1(v)$. Recall that $u' = |U'_i| \geq \lceil u/2 \rceil$ and consider two cases: either there is a node $v \in U'_i$ with $t_1(v) \geq \lceil u/2 \rceil$ or, by the pigeonhole principle, there are two distinct nodes v and w in U'_i such that $t_1(v) = t_1(w) < \lceil u/2 \rceil$. We set $L_i = \{v\}$ in the former case, to get $T_i = T_{i-1} + \min\{t_0(v), t_1(v)\} \geq T_{i-1} + \lceil u/2 \rceil$, and $L_i = \{v, w\}$ in the latter case, to get $T_i = T_{i-1} + \min\{t_0(v), t_0(w)\} \geq T_{i-1} + \lceil u/2 \rceil$. Thus we can force the delay at layer i of at least $\lceil u/2 \rceil$ by putting one or two nodes into this layer, so we have the following lemma.

▶ **Lemma 1.** *For each 2-shot broadcast protocol Π for n-node networks and a family of networks $\mathcal{G}(L_0, L_1, \ldots, L_{i-1})$, where $|\bigcup_{j=0}^{i-1} L_j| \leq n-3$, there exists $L_i \subseteq U_i = V \setminus \left(\{n-1\} \cup \bigcup_{j=0}^{i-1} L_j\right)$ such that $1 \leq |L_i| \leq 2$ and for each network in $\mathcal{G}(L_0, L_1, \ldots, L_{i-1}, L_i)$, the message does not leave layer i (that is, is not delivered to the next layer $i+1$) before the step $T_{i-1} + \lceil |U_i|/2 \rceil$.*

Using this lemma iteratively, we prove the following lower bound on the worst-case time of 2-shot broadcast protocols.

▶ **Theorem 2.** *For each 2-shot broadcast protocol Π for n-node networks, there exists a network in $\mathcal{G}_0$ on which Π needs at least $n^2/8 - O(n)$ steps to complete broadcast.*

Proof. Starting from $L_0 = \{0\}$, we apply Lemma 1 iteratively for $i = 1, 2, \ldots$ to obtained a network $G(L_0 = \{0\}, L_1, L_2, \ldots, L_{k-1}, L_k = \{n-1\})$ such that $k \geq n/2$, $1 \leq |L_i| \leq 2$, for each $1 \leq i \leq k-1$, and $T_i \geq T_{i-1} + |U_i|/2$. We have $|U_1| = n-2$ and $|U_i| \geq |U_{i-1}| - 2 \geq n - 2i$, for $2 \leq i \leq k-1$, so the worst-case number of steps needed by protocol Π to complete the broadcast is at least

$$1 + \sum_{1 \leq i \leq k-1} |U_i|/2 \;\geq \sum_{1 \leq i \leq (n/2)-1} (n-2i)/2 \;=\; n^2/8 - O(n). \qquad ◀$$

2.4 Lower bound for 3-shot broadcast

We consider now a 3-shot broadcast protocol Π and, as before, the family of networks $\mathcal{G}_{i-1}(L_0, L_1, \ldots, L_{i-1})$ for some arbitrary layer sets $L_0 = \{0\}, L_1, L_2, \ldots, L_{i-1}$. The message leaves layer $i-1$ at time step T_{i-1} and we want to select nodes for the next layer i to force a relatively large weigthed delay $(T_i - T_{i-1})/|L_i|$. We refer to the notation of (conditional) transmission sets and the related terminology introduced in Sections 2.1 and 2.2. A node v in the set $U' = U'_i$ transmits in the step $t_0(v) \geq u/2$ and in steps $t_1(v) < t_2(v)$, where $u = n - 1 - \left|\bigcup_{j=0}^{i-1} L_j\right|$ and $|U'| \geq u/2$.

For an integer parameter $1 \leq p \leq u/2$, which will be set later, we put each node $v \in U'$ into one of the sets V_0, V_1 and V_2, depending on when the v's transmission steps $t_1(v)$ and $t_2(v)$ are in relation to step p:

$$\begin{aligned} V_0 &= \{v \in U' : \, p < t_1(v) < t_2(v)\}, \\ V_1 &= \{v \in U' : \, t_1(v) \leq p \text{ and } t_2(v) > p\}, \\ V_2 &= \{v \in U' : \, t_1(v) < t_2(v) \leq p\}. \end{aligned}$$

For the set V_2, we construct an undirected (multi-)graph H_2 with vertices $t_l(v)$, where $v \in V_2$ and $l = 1, 2$, and edges $\{t_1(v), t_2(v)\}$ for all $v \in V_2$. More precisely, the vertex set and the edge (multi-)set of graph H_2 are

$$V(H_2) = \{t : t = t_l(v) \text{ for some } v \in V_2 \text{ and } 1 \le l \le 2\},$$
$$E(H_2) = \{\{t', t''\} : t' = t_1(v) \text{ and } t'' = t_2(v) \text{ for some } v \in V_2\}.$$

Thus graph H_2 has at most p vertices and exactly $|V_2| \le u'$ edges. There may be parallel edges in H_2 because there may be two nodes v', v'' in V_2 with $\{t_1(v'), t_2(v')\} = \{t_1(v''), t_2(v'')\}$. To avoid confusion, *nodes* are in the transmission *network*, while *vertices* are in the auxiliary *graph* H_2 (and in other similar auxiliary graphs constructed later). The vertices of graph H_2 correspond to (some) steps of the protocol and the edges of H_2 correspond to (some) nodes in the transmission network.

The underlying idea in our lower-bound argument is that if the number of edges in graph H_2 is relatively large in relation to p, that is, if H_2 is sufficiently dense, then it must contain a short cycle $\Gamma = (\tau_0, \tau_1, \ldots, \tau_{k-1}, \tau_k = \tau_0)$. Let $v_i \in V_2$ be the node in the transmission network which corresponds to the edge $\{\tau_i, \tau_{i+1}\}$ in this cycle, for $i = 0, 1, \ldots, k-1$. (Two parallel edges in H_2 would form a cycle of length $k = 2$.) If we set the next layer $L_i = \{v_0, v_1, \ldots, v_{k-1}\}$, then these nodes transmit in steps $\tau_0, \tau_1, \ldots, \tau_{k-1}$, but in each of these steps exactly two nodes in L_i transmit, resulting in a collision. This means that the progress of broadcast has to rely on one of the steps $t_0(v_0), t_0(v_1), \ldots, t_0(v_{k-1})$, but each of these steps is at least $u/2$, so the weighted delay at layer i is at least $u/(2k)$. Therefore we have the following lemma.

▶ **Lemma 3.** *If graph H_2 has a cycle Γ of length k, then taking for the layer L_i the set of transmission nodes which correspond to the edges of Γ gives the weighted delay at layer i at least $u/(2k)$.*

To proceed with our analysis, we need an upper bound on the *girth of a graph*, that is, on the length of a shortest cycle. The asymptotic bounds given below in Lemma 4 and in its corollary are widely known and sufficient for us, but we note that more precise bounds are available in the literature, for example, in [3].

▶ **Lemma 4.** *Every graph with p vertices and the minimum degree $d = d(p) \ge 3$ contains a cycle of length $O(\log p/\log d)$.*

Proof. Consider any graph H with p vertices and the minimum degree $d \ge 3$. Let v be any vertex in H, $k \ge 1$ and $H(v,k)$ the subgraph of G induced by the vertices within distance at most k from v. If H does not have a cycle of length $2k$ or less, then $H(v,k)$ is a tree and has more than $(d-1)^k$ vertices. This means that for $k = \lceil \log n/\log(d-1) \rceil$, $H(v,k)$ is not a tree and contains a cycle of length at most $2k = O(\log n/\log d)$. ◀

▶ **Corollary 5.** *Every graph of average degree d with p vertices contains a cycle of length $O(\log p/\log d)$.*

Proof. Any graph G of average degree d must contain a nonempty subgraph G' of minimum degree at least $d/2$. To see this, repeatedly remove from G all vertices of degree strictly less than $d/2$. Not all vertices can be removed in this process because otherwise G would contain fewer than $pd/2$ edges altogether, a contradiction. By applying Lemma 4 to G' the claim follows. ◀

▶ **Lemma 6.** *Let Π be any 3-shot broadcast protocol Π for n-node networks and consider any family of networks $\mathcal{G}(L_0, L_1, \ldots, L_{i-1})$, where $|\bigcup_{j=0}^{i-1} L_j| \le n/2$. There exists the next i-th layer $L_i \subseteq U_i = V \setminus \left(\bigcup_{j=0}^{i-1} L_j \cup \{n-1\}\right)$ such that for each network in $\mathcal{G}(L_0, L_1, \ldots, L_{i-1}, L_i)$, the weighted delay at layer i is $\Omega(n \log\log n/\log n)$.*

Proof. We set $p = n \log\log n/\log n$ and consider sets V_0, V_1 and V_2. If V_0 is not empty, then we take $L_i = \{v\}$ where v is an arbitrary node in V_0. All three steps $t_0(v), t_1(v)$ and $t_2(v)$ when node v transmits are at least p, so the weighted delay at layer i is at least p.

If $|V_1| > p$, then there must be two nodes v' and v'' in V_1 such that $t_1(v') = t_1(v'') \leq p$, but all other steps $t_0(v'), t_2(v'), t_0(v'')$ and $t_2(v'')$ when v' or v'' transmits are at least p. Taking $L_i = \{v', v''\}$ gives the weighted delay at layer i at least $p/2$.

If V_0 is empty and V_1 has fewer than p nodes, then V_2 has more than $u' - p$ nodes, so graph H_2 has $|V_2| > u' - p \geq |U_i|/2 - p \geq n/5$ edges but at most p vertices. This means that the average degree in H_2 is greater than $(2/5)n/p$, so, from Corollary 5, H_2 has a cycle Γ of length $O(\log p/\log(n/p)) = O(\log n/\log\log n)$. Thus Lemma 3 implies that taking for the layer L_i the set of transmission nodes which correspond to the edges of Γ gives the weighted delay at layer i at least $\Omega(n \log\log n/\log n)$. ◀

We are now ready to prove the lower bound for the 3-shot case.

▶ **Theorem 7.** *For each 3-shot broadcast protocol Π for n-node networks, there exists a network in $\mathcal{G}_0$ on which Π needs $\Omega(n^2 \log\log n/\log n)$ steps to complete broadcast.*

Proof. Starting from $L_0 = \{0\}$, we use Lemma 6 iteratively, obtaining layers $L_1, L_2, \ldots, L_m$ and stopping when $\bigcup_{0 \leq i \leq m} |L_i| > n/2$. From Lemma 6, there is a constant $c > 0$ such that for each layer $i = 1, 2, \ldots, m$, the weighted delay $(T_i - T_{i-1})/|L_i|$ is at least $cn \log\log n/\log n$. Therefore,

$$T_{\max} \geq T_m = 1 + \sum_{1 \leq i \leq m} (T_i - T_{i-1}) \geq \sum_{1 \leq i \leq m} (|L_i| cn \log\log n/\log n)$$
$$\geq (c/2) n^2 \log\log n/\log n.$$ ◀

3 Upper bounds for h-shot broadcast for $h \leq 3$

For the 2-shot case a trivial upper bound which matches asymptotically the $\Omega(n^2)$ lower bound of Section 2.3 is given by the oblivious Round Robin (which is actually a 1-shot broadcast protocol).

We provide in this section an upper bound of $O(n^2 \log\log n/\log n)$ for 3-shot broadcast, which matches our lower bound and shows that in contrast to the 2-shot case, the fastest adaptive 3-shot protocols are faster than the best oblivious protocols by a factor $\omega(1)$.[3] We base our approach on graph-theoretic results [16] showing that it is possible to construct relatively dense graphs of high girth. We use such graphs to specify appropriate transmission sets as detailed below.

To define the sequence of transmission sets in our protocol, we use a graph $H = H(n, p, g)$ with n edges, p vertices and girth g. Any graph $H(n, p, g)$ would do for the correctness of our protocol, but to achieve fast (worst case) broadcast, we need a graph with relatively small number of nodes p and high girth g. More precisely, to have asymptotically fastest broadcast, we need a graph $H(n, p, g)$ with $p = \Theta(n \log\log n/\log n)$ and $g = \Theta(\log n/\log\log n)$.

We identify the edge set $E(H)$ of graph H with the node set $V(G)$ of the transmission network G, and we number the vertices in H from 1 to p (in an arbitrary way). Let H_i denote the set of edges in H which are incident to vertex i. The sets $H_1, H_2, \ldots, H_p$ are (some of) the transmission sets of our protocol. Clearly, for any node v of G, v belongs to two

[3] Recall that the oblivious bound is $\Theta(n^2/k)$ for k-shot protocols and $k \leq \sqrt{n}$.

sets H_i and H_j, where $\{i, j\} \in E(H)$ is the edge identified with node v. Node v transmits in two steps with transmissions sets H_i and H_j, while the third transmission is within one Round-Robin sequence.

Formally, our protocol $\Pi(H)$ is defined by the repeated Round-Robin sequence $\langle R\rangle = (\{0\}, \{1\}, \ldots, \{n-1\})$ interleaved with the repeated sequence $\langle H\rangle = (H_1, H_2, \ldots, H_p)$. Let's say that we use the odd steps of the protocol for repeating the Round-Robin sequence and the even steps for repeating the sequence $\langle H\rangle$. If a node v receives the message in step t, then it transmits in its step of the first Round-Robin sequence which starts after step t, and in the steps H_i and H_j of the first copy of the sequence $\langle H\rangle$ which starts after step t, where $\{i, j\} \in E(H)$ is the edge identified with node v.

We now proceed with the analysis of protocol $\Pi(H)$. Consider any n-node transmission network G with source s and an arbitrary node $v \neq s$. Let $k \geq 1$ denote the distance from s to v. In order to upper-bound the time needed for the message to go from source node s to node v, we consider the partitioning $L_v(G)$ of the nodes within distance k to v into layers. These are breadth-first-search layers constructed from node v following the edges of G in reverse direction. For $0 \leq i \leq k$, the layer L_i is the set of all nodes in G with distance $k-i$ to v. Thus $L_k = \{v\}$, L_{k-1} is the set of all nodes with edges to v, and so on. The source node s belongs to layer L_0.

Note that for each $1 \leq i \leq k$, each node $u \in L_i$ and each edge (x, u), $x \in L_j$ for some $j \geq i-1$. Thus the message reaches layer L_i (any node in layer L_i) for the first time during a transmission by a node from layer L_{i-1}. We use T_i to denote the time step at which the message first reaches layer L_i. We have $T_1 = 1$ (layer L_1 must have at least one out-neighbour of the source) and the following lemma gives an upper bound on the delays at layers of relatively small cardinality.

▶ **Lemma 8.** *Consider an n-node transmission network G with source s, an arbitrary node v, the layers $L_0, L_1, \ldots, L_k$ corresponding to this node and the protocol $\Pi(H)$ defined by a graph $H = H(n, p, g)$. During the execution of this protocol, if $|L_i| < g$, then the time needed to transmit the message from L_i to L_{i+1}, that is, $T_{i+1} - T_i$, is at most $4p$.*

Proof. Let $L'_i \subseteq L_i$ be the set of nodes in L_i that have received the message by time $T_i + t$, where t is the smallest integer such that $(T_i + t) \bmod (2p) = 0$. Only nodes in L'_i will be transmitting at even steps between $T_i + t + 1$ and $T_i + t + 2p$.

Since $|L'_i| \leq |L_i| < g$, the edges corresponding to nodes of L'_i form an acyclic subgraph T of H, so for each vertex w_j in H with degree 1 in T (there must be at least two such vertices) the transmission set H_j contains exactly one node from L'_i. During each such step, the message is transmitted from layer L_i to layer L_{i+1}. Hence $T_{i+1} \leq T_i + t + 2j \leq T_i + 4p$. ◀

▶ **Theorem 9.** *Protocol $\Pi(H)$ defined by a graph $H = H(n, p, g)$ completes broadcast in an arbitrary n-node transmission network G within $O(n^2/g + np)$ steps.*

Proof. We take an arbitrary node $v \neq s$ and consider its layers $L_0, L_1, \ldots, L_k$. There can be at most n/g layers of size at least g. For each such layer L_i, when a message arrives at this layer, then it will reach the next layer L_{i+1} by the time the next full Round Robin is completed. That is, in this case $T_{i+1} \leq T_i + 4n$. Combining this with Lemma 8 gives the claimed bound on the number of steps, since the number of layers of size smaller than g is at most $n-1$. ◀

To minimize the upper bound $O(n^2/g+np) = O(n^2/\min\{g, d_{ave}\})$, where d_{ave} is the average degree in graph $H(n, p, g)$, we have to find a graph with n edges and $\min\{g, d_{ave}\}$ as large as possible. Corollary 5 implies that for all graphs, $\min\{g, d_{ave}\} = O(\min\{\log n/\log d_{ave},$

$d_{ave}\}) = O(\log n/\log\log n)$. It turns out that there are explicitly constructed graphs with n edges for which $\min\{g, d_{ave}\} = \Theta(\log n/\log\log n)$. We use the construction given in [16].

▶ **Theorem 10** ([16]). *For each positive odd integer $k \geq 3$ and a power of a prime q, there is an explicit construction of a q-regular bipartite graph $H(q,k)$ with $2q^k$ vertices and girth at least $k+5$.*

▶ **Corollary 11.** *There exists an explicit construction of a graph H with n edges, $p = \Theta(n\log\log n/\log n)$ vertices and girth $g = \Theta(\log n/\log\log n)$.*

Proof. For a given sufficiently large n, let $q \geq 4$ be the largest power of 2 not greater than $\log n/\log\log n$ and let $k = q-1 \geq 3$. Let $H(q,k)$ be the graph from Theorem 10. This graph has $2q^{q-1}$ vertices, $q^q \leq n$ edges and girth at least $q+4$.

Let H be a graph with exactly n edges obtained by taking copies of graph $H(q,k)$ as connected components. We remove (arbitrarily) some edges from the last copy of $H(q,k)$ so that the total number of edges is exactly n. We need $\lceil n/q^q \rceil$ copies of $H(q,k)$, so the number of vertices in graph H is at most $2q^{q-1}(n/q^q+1) \leq 4n/q \leq 8n\log\log n/\log n$. Graph H has the same girth as $H(q,k)$, so at least $q+4 \geq (1/2)\log n/\log\log n$. ◀

Using in protocol $\Pi(H)$ the graph H from Corollary 11, Theorem 9 gives us the following result.

▶ **Corollary 12.** *There exists a constructive 3-shot broadcast protocol which completes broadcast on any graph G with n nodes in time $O(n^2\log\log n/\log n)$.*

4 Upper bounds for h-shot broadcast for $h \geq 4$

It was shown in [14] that for any $h \geq 1$, an h-shot broadcast protocol requires $\Omega\left(n^{1+1/h}\right)$ steps. In previous sections, we improved this lower bound and provided matching upper bounds for the cases when h is equal to 2 and 3. In this section, we show upper bounds for $h \geq 4$. In particular, if h is a sufficiently large constant or is slowly growing with n, then we prove that there exist h-shot broadcast protocols with $O(n^{1+\alpha/\sqrt{h}})$ steps, where α is an absolute constant independent of h.

The general idea for h-shot broadcast protocols for $h \geq 4$ is similar to the idea of using a large girth graph to construct a 3-shot protocol. Now, however, we need to define $r = h-1 \geq 3$ transmission slots for each node (in addition to the transmissions defined by Round-Robin), so we use r-uniform hyper-graphs instead of graphs $H(n,p,g)$. Let $H_r = H_r(n,p,k)$ be an r-uniform hyper-graph (each edge is a set of r vertices) with n (hyper-)edges, p vertices, and no 2-cover of size k or smaller. A 2-*cover* of a hyper-graph is a non-empty subset $\mathcal{A}$ of edges such that each node which belongs to an edge in $\mathcal{A}$ belongs to at least two edges in $\mathcal{A}$. The notion of 2-covers in hyper-graphs generalizes the notion of cycles in graphs: minimal 2-covers in graphs are (simple) cycles.

Similarly as in the previous subsection, we identify the edge set $E(H_r)$ of the hyper-graph H_r with the node set $V(G)$ of the transmission network G. We number the vertices in H_r from 1 to p in an arbitrary order and denote by $H_r^{(i)}$ the set of edges in H_r which are incident to vertex i. If we use the sequence $\langle H_r\rangle = \langle H_r^{(1)}, H_r^{(2)}, \ldots, H_r^{(p)}\rangle$ as a sequence of transmission sets, then for each nonempty subset W of at most k nodes in G, one of these transmission sets has exactly one node from W – otherwise the set of edges in H_r corresponding to the nodes in W would form a 2-cover in H_r of size at most k.

The following simple counting argument shows how large k can be in an $H_r(n,p,k)$ hyper-graph.

▶ **Lemma 13.** *There is a constant C such that for each $n \geq 1$ and for each $p \geq r \geq 3$, there exists a hyper-graph $H_r = H_r(n,p,k)$ with $k = \lfloor p/(Crn^{2/r}) \rfloor$.*

Proof. We consider a random r-uniform hyper-graph $\mathcal{H}$ with p vertices and n edges (independently and uniformly selected from the family of sets of r vertices) and show that for k defined in the lemma (where constant C will come out from the calculations) and for a fixed $2 \leq q \leq k$, the probability that $\mathcal{H}$ has a 2-cover of size q is at most $1/2^q$. By summing up over all $2 \leq q \leq k$, we get the conclusion that there must exist a hyper-graph $H_r = H_r(n,p,k)$.

A 2-cover $\mathcal{A}$ of size q covers at most $qr/2$ vertices, or otherwise there would be a vertex belonging to exactly one edge in $\mathcal{A}$. Thus the probability that $\mathcal{H}$ has a 2-cover of size q is at most the probability that there exists in $\mathcal{H}$ a set $\mathcal{A}$ of q edges and a set X of $qr/2$ vertices such that each edge in $\mathcal{A}$ is a subset of X. Using the union bound over all possible $\mathcal{A}$ and X, the probability of the latter event is at most

$$\binom{n}{q}\binom{p}{qr/2}\frac{\binom{qr/2}{r}^q}{\binom{p}{r}^q} \leq \left(\frac{en}{q}\right)^q \left(\frac{2ep}{qr}\right)^{qr/2} \frac{(eqr/2)^{qr}/r^{qr}}{p^{qr}/r^{qr}} \leq \frac{1}{q^q}\left(\frac{Cqrn^{2/r}}{p}\right)^{qr/2} \leq \frac{1}{2^q},$$

where the second inequality holds for $C = (2e)^2$ and the last one holds for any $2 \leq q \leq p/(Crn^{2/r})$. For the first inequality, we use $\left(\frac{a}{b}\right)^b \leq \binom{a}{b} \leq \left(\frac{ea}{b}\right)^b$. ◀

For a hyper-graph $H_r = H_r(n,p,k)$, the protocol $\Pi(H_r)$ which interleaves repeated copies of $\langle H_r \rangle$ with copies of a Round-Robin sequence $\langle R \rangle$ is an h-shot broadcast protocol with $O(n^2/k + np)$ steps. This can be shown in an analogous way as in the proof of Theorem 9, by considering separately the layers with sizes at most k and the layers with sizes greater than k. If we consider hyper-graphs $H_r = H_r(n,p,k)$ with $k = \lfloor p/(Crn^{2/r}) \rfloor$, whose existence is guaranteed by Lemma 13, and take $p = r^{1/2}n^{1/2+1/r}$ to minimize $O(n^2/k + np)$, then we obtain an h-shot broadcast protocol with $O(h^{1/2}n^{3/2+1/(h-1)})$ steps. This gives, for example, upper bounds $O(n^{11/6})$ and $O(n^{7/4})$ for 4-shot and 5-shot broadcast, respectively, but no better bound than $O(n^{3/2})$ even if h grows to infinity.

To obtain upper bounds with the exponent at n decreasing to 1 for increasing values of h, we combine hyper-graphs $H_r(n,p,k)$ for a number of different values of k. More specifically, for $h = \rho^2/2+1$, where ρ is an even integer at least 4, let $H_{\rho,j} = H_\rho(n, C\rho n^{2j/\rho}, n^{2(j-1)/\rho})$, for $j = 1, 2, \ldots, J = \rho/2$, where C is the constant from Lemma 13. Our h-shot broadcast protocol Π_h is defined by the sequence of transmission sets obtained by interleaving $\rho + 1$ sequences $(\langle H_{\rho,1}\rangle, \langle H_{\rho,1}\rangle, \ldots)$, $(\langle H_{\rho,2}\rangle, \langle H_{\rho,2}\rangle, \ldots), \ldots, (\langle H_{\rho,J}\rangle, \langle H_{\rho,J}\rangle, \ldots)$, and $(\langle R\rangle, \langle R\rangle, \ldots)$, and by the following transmission schedule. For a node v in the transmission network G, if v receives the message for the first time in step t, then let $\overline{\langle H_{\rho,j}\rangle}$, for $j = 1, 2, \ldots, J$, and $\overline{\langle R\rangle}$ be, respectively, the first copies of $\langle H_{\rho,1}\rangle, \langle H_{\rho,2}\rangle, \ldots, \langle H_{\rho,J}\rangle$ and $\langle \mathcal{R}\rangle$ which start after step t. Node v transmits in the steps corresponding to the transmission sets in $\overline{\langle H_{\rho,1}\rangle}, \overline{\langle H_{\rho,2}\rangle}, \ldots, \overline{\langle H_{\rho,J}\rangle}$ and $\overline{\langle R\rangle}$ which include v. Thus v transmits in $\rho \cdot (\rho/2) + 1 = h$ steps.

▶ **Theorem 14.** *For $h = \rho^2/2+1$, where ρ is an even integer at least 4, the (non-constructive) protocol Π_h is an h-shot broadcast protocol with $O(hn^{1+\sqrt{8/(h-1)}})$ steps.*

Proof. By the definition of protocol Π, no node transmits more then h times. We show now the claimed bound on the number of steps.

Similarly to the analysis of the 3-shot protocol in Section 3, we consider an arbitrary node v and its in-neighbourhood layers $L_0, L_1, \ldots, L_k$, where $s \in L_0$ and $v \in L_k$. The delay at layer L_i, that is, the number of steps between the time when the first node in L_i receives the message and the time when the first node in L_{i+1} receives the message (from one of the nodes in L_i), depends on the size of this layer. If $n^{2(j-2)/\rho} < |L_i| \le n^{2(j-1)/\rho}$, for some $1 \le j \le J$, then the message is delivered from (one of the nodes of) layer L_i to (one of the nodes of) the next layer by the next copy of $\langle H_{\rho,j} \rangle$, so within $C\rho^2 n^{2j/\rho}$ steps. (Recall that the transmission sets of each $\langle H_{\rho,i} \rangle$ are scheduled every $\rho/2+1$ steps, hence the additional factor of ρ).

For a layer L_i such that $n^{2(J-1)/\rho} < |L_i|$, the message is delivered to the next layer by the next copy of Round-Robin, so within ρn steps. Thus the delay at each layer L_i is at most $C\rho^2 n^{4/\rho}|L_i|$ steps, so node v receives the message within $O(\rho^2 n^{1+4/\rho}) = O(hn^{1+\sqrt{8/(h-1)}})$ steps. ◀

We defined protocols Π_h only for values $h = \rho^2/2+1$, where ρ is an even integer at least 4. Since the h-shot broadcast protocol Π_h is also an h'-shot broadcast protocol for any $h' \ge h$, then Theorem 14 implies the following corollary.

▶ **Corollary 15.** *There is a constant α such that for any $1 \le h = O(\log n)$, there exists an h-shot broadcast protocol with $O(\min\{n^2, n^{1+\alpha/\sqrt{h}}\})$ steps.*

Proof. It is enough to consider $h \ge 9$, since the case when $h < 9$ can be covered by taking sufficiently large α. For $h \ge 9$, take $\rho = \lfloor\sqrt{2(h-1)}\rfloor$, $\tilde{h} = \rho^2/2+1 \le h$ and the protocol $\Pi_{\tilde{h}}$, which is an h-shot broadcast protocol. Theorem 14 implies that protocol $\Pi_{\tilde{h}}$ works in $O(\rho^2 n^{1+4/\rho})$ steps, which is $O(n^{1+5/\sqrt{h}})$ for $9 \le h = O(\log n)$. ◀

For the cases $h = 2$ and $h = 3$, we have obtained asymptotically matching lower and upper bounds on the number of steps in h-shot broadcast protocols. For $h \ge 4$, however, we still have a gap between the lower bound of $\Omega\left(n^{1+1/h}\right)$ shown by Karmakar et al. [14] and our upper bounds.

References

1 Reuven Bar-Yehuda, Oded Goldreich, and Alon Itai. On the time-complexity of broadcast in radio networks: An exponential gap between determinism and randomization. In *PODC '87*, pages 98–108, 1987.

2 Petra Berenbrink, Colin Cooper, and Zengjian Hu. Energy efficient randomised communication in unknown adhoc networks. *Theoretical Computer Science*, 410(27-29):2549–2561, 2009. doi:10.1016/j.tcs.2009.02.002.

3 Norman Biggs. *Algebraic Graph Theory.* Cambridge University Press, Cambridge, 2nd edition, 1993.

4 Danilo Brusci and Massimiliano Del Pinto. Lower bounds for the broadcast problem in mobile radio networks. *Distributed Computing*, 10(3):129–135, 1997.

5 Imrich Chlamtac and Shay Kutten. On broadcasting in radio networks-problem analysis and protocol design. *IEEE Transactions on Communications 33*, pages 1240–1246, 1985.

6 Bogdan S. Chlebus, Leszek Gąsieniec, Alan Gibbons, Andrzej Pelc, and Wojciech Rytter. Deterministic broadcasting in unknown radio networks. In *Proc. 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'00)*, pages 861–870. ACM/SIAM, 2000.

7 Bogdan S. Chlebus, Leszek Gąsieniec, Anna Östlin, and John Michael Robson. Deterministic radio broadcasting. In *Proc. 27th International Colloquium on Automata, Languages and Programming (ICALP'00)*, volume 1853 of *Lecture Notes in Computer Science*, pages 717–728. Springer Verlag, 2000.

8 Marek Chrobak, Leszek Gąsieniec, and Wojciech Rytter. Fast broadcasting and gossiping in radio networks. *Journal of Algorithms*, 43(2):177–189, 2002.

9 Andrea E. F. Clementi, Angelo Monti, and Riccardo Silvestri. Distributed broadcast in radio networks of unknown topology. *Theoretical Computer Science*, 302(1-3):337–364, 2003. `doi:10.1016/S0304-3975(02)00851-4`.

10 Artur Czumaj and Wojciech Rytter. Broadcasting algorithms in radio networks with unknown topology. In *Proc. 44th Symposium on Foundations of Computer Science (FOCS'03)*, pages 492–501. IEEE Computer Society, 2003.

11 Gianluca De Marco. Distributed broadcast in unknown radio networks. *SIAM Journal on Computing*, 39(6):2162–2175, March 2010.

12 Leszek Gąsieniec, Erez Kantor, Dariusz R. Kowalski, David Peleg, and Chang Su. Time efficient k-shot broadcasting in known topology radio networks. *Distributed Computing*, 21(2):117–127, 2008. `doi:10.1007/s00446-008-0058-0`.

13 Erez Kantor and David Peleg. Efficient k-shot broadcasting in radio networks. *Discrete Applied Mathematics*, 202:79–94, 2016.

14 Sushanta Karmakar, Paraschos Koutris, Aris Pagourtzis, and Dimitris Sakavalas. Energy-efficient broadcasting in ad hoc wireless networks. *Journal of Discrete Algorithms*, 42:2–13, 2017. `doi:10.1016/j.jda.2016.11.004`.

15 Dariusz R. Kowalski and Andrzej Pelc. Broadcasting in undirected ad hoc radio networks. *Distributed Computing*, 18(1):43–57, 2005. `doi:10.1007/s00446-005-0126-7`.

16 Felix Lazebnik and Vasiliy A. Ustimenko. Explicit construction of graphs with an arbitrary large girth and of large size. *Discrete Applied Mathematics*, 60(1-3):275–284, 1995. `doi:10.1016/0166-218X(94)00058-L`.

17 Gianluca De Marco and Andrzej Pelc. Faster broadcasting in unknown radio networks. *Inf. Process. Lett.*, 79(2):53–56, 2001. `doi:10.1016/S0020-0190(00)00178-2`.