Distributed Deterministic Broadcasting in Wireless Networks of Weak Devices*

Tomasz Jurdzinski[1], Dariusz R. Kowalski[2], and Grzegorz Stachowiak[1]

[1] Institute of Computer Science, University of Wrocław, Poland
[2] Department of Computer Science, University of Liverpool, United Kingdom

Abstract. Many futuristic technologies, such as Internet of Things or nano-communication, assume that a large number of simple devices of very limited energy and computational power will be able to communicate efficiently via wireless medium. Motivated by this, we study broadcasting in the model of ad-hoc wireless networks of weak devices with uniform transmission powers. We compare two settings: with and without local knowledge about immediate neighborhood. In the latter setting, we prove $\Omega(n \log n)$-round lower bound and develop an algorithm matching this formula. This result could be made more accurate with respect to network density, or more precisely, the maximum node degree Δ in the communication graph. If Δ is known to the nodes, it is possible to broadcast in $O(D\Delta \log^2 n)$ rounds, which is almost optimal in the class of networks parametrized by D and Δ due to the lower bound $\Omega(D\Delta)$. In the setting with local knowledge, we design a scalable and almost optimal algorithm accomplishing broadcast in $O(D \log^2 n)$ communication rounds, where n is the number of nodes and D is the eccentricity of a network. This can be improved to $O(D \log g)$ if network granularity g is known to the nodes. Our results imply that the cost of "local communication" is a dominating component in the complexity of wireless broadcasting by weak devices, unlike in traditional models with non-weak devices in which well-scalable solutions can be obtained even without local knowledge.

1 Introduction

1.1 The Model

We consider a wireless network consisting of n *stations*, also called *nodes*, deployed into an Euclidean plane and communicating by a wireless medium. The *Euclidean metric* on the plane is denoted $\text{dist}(\cdot,\cdot)$. Each station v has its *transmission power* P_v, which is a positive real number. There are three fixed model parameters: path loss $\alpha > 2$, threshold $\beta \geq 1$, and ambient noise $\mathcal{N} > 0$.

* The full version of the paper is available at [13]. This work was supported by the Polish National Science Centre grant DEC-2012/06/M/ST6/00459.

F.V. Fomin et al. (Eds.): ICALP 2013, Part II, LNCS 7966, pp. 632–644, 2013.

The $SINR(v, u, \mathcal{T})$ ratio, for given stations u, v and a set of (transmitting) stations $\mathcal{T}$, is defined as follows:

$$SINR(v, u, \mathcal{T}) = \frac{P_v \mathrm{dist}(v, u)^{-\alpha}}{\mathcal{N} + \sum_{w \in \mathcal{T} \setminus \{v\}} P_w \mathrm{dist}(w, u)^{-\alpha}} \tag{1}$$

In the *weak devices model* considered in this work, a station u successfully receives a message from a station v in a round if $v \in \mathcal{T}$, $u \notin \mathcal{T}$, and:

a) $P_v \mathrm{dist}^{-\alpha}(v, u) \geq (1 + \varepsilon)\beta\mathcal{N}$, and
b) $SINR(v, u, \mathcal{T}) \geq \beta$,

where $\mathcal{T}$ is the set of stations transmitting at that time and $\varepsilon > 0$ is a fixed *signal sensitivity parameter* of the model.[1]

Ranges and Uniformity. The *communication range* r_v of a station v is the radius of the ball in which a message transmitted by the station is heard, provided no other station transmits at the same time. A network is *uniform*, when transmission powers P_v and thus ranges of all stations r_v are equal, or *nonuniform* otherwise. In this paper, only uniform networks are considered, i.e., $P_v = P$ and $r = r_v = (P/(\mathcal{N}\beta(1 + \varepsilon)))^{1/\alpha}$. The *range area* of a station v is defined to be the ball of radius r centered in v.

Communication Graph and Graph Notation. The *communication graph* $G(V, E)$, also called the *reachability graph*, of a given network consists of all network nodes and edges (v, u) such that u is in the range area of v. Note that the communication graph is symmetric for uniform networks. By a *neighborhood* of a node u we mean the set (and positions) of all neighbors of u in G, i.e., the set $\{w \mid (w, u) \in E(G)\}$. The *graph distance* from v to w is equal to the length of a shortest path from v to w in the communication graph, where the length of a path is equal to the number of its edges. The *eccentricity* of a node is the maximum graph distance from this node to any other node (note that the eccentricity is of order of the diameter). By Δ we denote the maximum degree of a node in the communication graph.

Synchronization. It is assumed that algorithms work synchronously in rounds, each station can either act as a sender or as a receiver during a round. We do not assume global clock ticking.

Carrier Sensing. We consider the model *without carrier sensing*, that is, a station u has no other feedback from the wireless channel than receiving or not receiving a message in a round t.

[1] This model is motivated by the fact that it is too costly for weak devices to have receivers doing signal acquisition continuously, c.f., [7]. Therefore, in many systems they rather wait for an energy spike, c.f., condition (a), and once they see it, they start sampling and correlating to synchronize and acquire a potential packet preamble [19]. Once synchronized, they can detect signals, c.f., condition (b).

Knowledge of Stations. Each station has its unique ID from the set $[N]$,[2] where N is polynomial in n. Stations also know their locations, and parameters n, N. Some subroutines use the granularity g, defined as r divided by the minimum distance between any two stations (c.f., [5]). We distinguish between networks *without local knowledge (ad hoc)*, where stations do not know anything about the topology of the network, and networks *with local knowledge*, in which each station knows locations and IDs of its neighbors in the communication graph.

Broadcasting Problem and Complexity Parameters. In the broadcast problem, there is one distinguished node, called the *source*, which initially holds a piece of information (also called a source message or a broadcast message). The goal is to disseminate this message to all other nodes. The complexity measure is the worst-case time to accomplish the broadcast task, taken over all connected networks with specified parameters. Time, also called the *round complexity*, denotes the number of communication rounds in the execution of a protocol: from the round when the source is activated with its source message till the broadcast task is accomplished. For the sake of complexity formulas, we consider the following parameters: n, N, D, and g.

Messages and Initialization of Stations Other than Source. We assume that a single message sent in the execution of any algorithm can carry the broadcast message and at most polynomial, in the size of the network, number of control bits. (For the purpose of our algorithms, it is sufficient that positions of stations on the plane are stored with accuracy requiring $O(\log n)$ bits; therefore, we assume that each message contains the position of its sender.) A station other than the source starts executing the broadcast protocol after the first successful receipt of the source message; it is often called a *non-spontaneous wake-up model.*

1.2 Our Results

In this paper we present the first study of deterministic distributed broadcasting in wireless networks of weak devices with uniform transmission powers, deployed in the two dimensional Euclidean space. We distinguish between the two settings: with and without local knowledge about the neighbors in the communication graph. In the latter model, we developed an algorithm accomplishing broadcast in $O(n \log n)$ rounds, which matches the lower bound (Sections 2.1 and 2.3, resp.). Then, an algorithm accomplishing broadcast in time $O(D\Delta \log^2 n)$ is presented, where D is the eccentricity of the source and Δ is the largest degree of a node in the communication graph (Section 2.2). This algorithm is close to the lower bound $\Omega(D\Delta)$, see Section 2.3. Our solution for networks with local knowledge works in $O(D \log^2 n)$ rounds (Section 3), which provides only a small $O(\log^2 n)$ overhead over the straightforward lower bound of $\Omega(D)$, and is faster, in the worst case, than any algorithm designed for networks without local knowledge of eccentricity $D = o(n/ \log n)$ or maximal degree $\Delta = \omega(1)$. It also implies that the cost of learning neighborhoods by stations in wireless network

[2] We denote $[i] = \{1, 2, \ldots, i\}$, $[i, j] = \{i, i + 1, \ldots, j\}$ for $i, j \in \mathbb{N}$.

is much higher, by factor around $\min\{n/D, \Delta\}$, than the cost of broadcast itself (i.e., broadcast performed when such neighborhoods would be provided). If the granularity g is known, a complexity $O(D \log g)$ can be achieved by a variation of the algorithm mentioned above.

Our results rely on novel techniques which simultaneously exploit specific properties of conflict resolution in the SINR model (see e.g., [1]) and several algorithmic techniques developed for a different radio network model. In particular, we show how to efficiently combine a novel SINR-based communication technique, ensuring several simultaneous point-to-point communications inside the range area of one station (which is unfeasible to achieve in the radio network model), with strongly selective families and methods based on geometric grids developed in the context of radio networks. As a result, we are able to transform algorithms relying on the knowledge of network's granularity into algorithms of asymptotically similar performance (up to a $\log n$ factor) that do not require such knowledge; this is in particular demonstrated in the leader election algorithms.

Details of some algorithms and technical proofs can be found in the full version of the paper [13].

1.3 Previous and Related Results

To the best of our knowledge, this is the first theoretical study of the problem of distributed deterministic broadcasting in ad hoc wireless networks of weak devices. In what follows, we list most relevant results in the SINR-based model and in the older, but still related, radio network model.

SINR Models. In the model of (uniform) weak devices, distributed algorithms for building a backbone structure in $O(\Delta \text{ polylog } n)$ rounds were constructed in [11]. Unlike in our broadcast problem, in [11] it was assumed that all nodes simultaneously start building the backbone. That result combined with the results of this work implicates that there is an extra cost payed for the lack of initial synchronization. If devices are not weak (i.e., not restricted by the fact that the signal must be sufficiently strong in order to be noticed), broadcasting can be done in $O(D \log^2 n)$, as proved in [14]. Combined with results in this paper, it proves a complexity gap between the two models: weak and non-weak devices.

Under the SINR-based models in ad hoc setting, a few other problems were also studied, such as deterministic data aggregation [10] and *local* broadcasting [20], in which nodes have to inform only their neighbors in the corresponding reachability graph. The considered setting allowed power control by algorithms, in which, in order to avoid collisions, stations could transmit with any power smaller than the maximal one. Randomized solutions for contention resolution [15] and local broadcasting [8] were also obtained.

There is a vast amount of work on centralized algorithms under the SINR model. The most studied problems include connectivity, capacity maximization, and link scheduling types of problems; for recent results and references we refer the reader to the survey [9]. Multiple Access Channel properties were also recently studied under the SINR model, c.f., [18].

Radio Network Model. In this model, a transmitted message is successfully heard if there are no other simultaneous transmissions from the neighbors of the receiver in the communication graph. This model does not take into account the real strength of the received signals, and also the signals from outside of the close proximity. In the geometric ad hoc setting, Dessmark and Pelc [4] were the first who studied this problem. They analyzed the impact of local knowledge, defined as the range within which stations can discover the nearby stations. Emek et al. [5] designed a broadcast algorithm working in time $O(Dg)$ in Unit Disc Graphs (UDG) radio networks with eccentricity D and granularity g. Later, Emek et al. [6] developed a matching lower bound $\Omega(Dg)$. In the *graph-based* model of radio networks, in which stations are not explicitly deployed in a metric space, the fastest $O(n\log(n/D))$-round deterministic algorithm was developed by Kowalski [16], and almost matching lower bound was given by Kowalski and Pelc [17], who also studied fast randomized solutions (in parallel with [3]). The above results hold without assuming local knowledge. With local knowledge, Jurdzinski and Kowalski [12] showed a lower bound $\Omega(\sqrt{Dn\log n})$ on the number of rounds and an algorithm of complexity $O(D\sqrt{n}\log^6 n)$.

1.4 Technical Preliminaries

In the broadcast problem, a round counter could be easily maintained by already informed nodes by passing it along the network with the source message, so in all algorithms we in fact assume having a global clock. For simplicity of analysis, we assume that every message sent during the execution of our broadcast protocols contains the broadcast message; in practice, further optimization of a message content could be done in order to reduce the total number of transmitted bits in real executions. In a given round t we say that a station v transmits *c-successfully* in round t if v transmits a message in round t and this message is heard by each station u in the Euclidean distance at most c from v. We say that a station v transmits *successfully* in round t if it transmits r-successfully, i.e., each of its neighbors in the communication graph can hear its message. Finally, v transmits *successfully* to u in round t if v transmits a message and u receives this message in round t. We say that a station that received the broadcast message is *informed*.

Grids. Given a parameter $c > 0$, we define a partition of the 2-dimensional space into square boxes of size $c \times c$ by the grid G_c, in such a way that: all boxes are aligned with the coordinate axes, point $(0,0)$ is a grid point, each box includes its left side without the top endpoint and its bottom side without the right endpoint and does not include its right and top sides. We say that (i,j) are the coordinates of the box with its bottom left corner located at $(c\cdot i, c\cdot j)$, for $i,j\in\mathbb{Z}$. A box with coordinates $(i,j)\in\mathbb{Z}^2$ is denoted $C(i,j)$. As observed in [4,5], the *grid* $G_{r/\sqrt{2}}$ is very useful in the design of the algorithms for UDG (unit disk graph) radio networks, where r is equal to the range of each station. This follows from the fact that $r/\sqrt{2}$ is the largest parameter of a grid such that each station in a box is in the range of every other station in that box. We fix $\gamma = r/\sqrt{2}$ and call G_γ the *pivotal grid*. If not stated otherwise, our considerations will refer to (boxes of) G_γ. The boxes C, C' of the pivotal grid are *neighbors* in a network if there are stations $v \in C$ and $v' \in C'$ such that the edge (v, v') belongs to the communication graph. We define the set DIR $\subset [-2, 2]^2$ such that $(d_1, d_2) \in$ DIR iff it is possible that boxes $C(i, j)$ and $C(i + d_1, j + d_2)$ are neighbors.

Schedules. A (general) *broadcast schedule* $\mathcal{S}$ of length T wrt $N \in \mathbb{N}$ is a mapping from $[N]$ to binary sequences of length T. A station with identifier $v \in [N]$ *follows* the schedule $\mathcal{S}$ of length T in a fixed period of time consisting of T rounds, when v transmits a message in round t of that period iff the position $t \mod T$ of $\mathcal{S}(v)$ is equal to 1. For the tuples (i_1, j_1), (i_2, j_2) the relation $(i_1, j_1) \equiv (i_2, j_2) \mod d$ for $d \in \mathbb{N}$ denotes that $(|i_1 - i_2| \mod d) = 0$ and $(|j_1 - j_2| \mod d) = 0$. A set of stations A on the plane is δ*-diluted* wrt G_c, for $\delta \in \mathbb{N} \setminus \{0\}$, if for any two stations $v_1, v_2 \in A$ with grid coordinates (i_1, j_1) and (i_2, j_2), respectively, the relationship $(i_1, j_1) \equiv (i_2, j_2) \mod d$ holds. We say that δ*-dilution* is applied to a schedule S if each round of an execution of S is replaced with δ^2 rounds parameterized by $(i, j) \in [0, \delta - 1]^2$ such that a station $v \in C(a, b)$ can transmit a message only in the rounds (i, j) such that $(i, j) \equiv (a, b) \mod \delta$.

Proposition 1. *For each $\alpha > 2$ and $\varepsilon > 0$, there exists a constant d_0 such that the following properties hold. Assume that a set of n stations A is d-diluted wrt the grid G_x, where $x = \gamma/c$, $c \in \mathbb{N}$, $c > 1$ and $d \geq d_0$. Moreover, at most one station from A is located in each box of G_x. Then, if all stations from A transmit simultaneously, each of them transmits $\frac{2r}{c}$-successfully.*

Proposition 2. *For each $\alpha > 2$ and $\varepsilon > 0$, there exists a constant d satisfying the following property. Let A be a set of stations such that $\min_{u,v \in A}\{dist(u, v)\} = x \cdot \sqrt{2}$, where $x \leq \gamma$. If a station $u \in C(i, j)$ for a box $C(i, j)$ of G_x is transmitting in a round t and no other station in any box $C(i', j')$ of G_x such that $\max\{|i - i'|, |j - j'|\} \leq d$ is transmitting at that round, then v can hear the message from u at round t.*

Selective families. A family $S = (S_0, \ldots, S_{s-1})$ of subsets of $[N]$ is a (N, k)*-ssf (strongly-selective family)* of length s if, for every non empty subset Z of $[N]$ such that $|Z| \leq k$ and for every element $z \in Z$, there is a set S_i in S such that $S_i \cap Z = \{z\}$. It is known that there exists (N, k)-ssf of size $O(k^2 \log N)$ for every $k \leq N$, c.f., [2]. We identify a family of sets $S = (S_0, \ldots, S_{s-1})$ with the broadcast schedule S' such that the ith bit of $S'(v)$ is equal to 1 iff $v \in S_i$.

2 Algorithms without Local Knowledge

2.1 Size Dependent Algorithm

In this section we consider networks in which a station knows only n, N, its own ID and its coordinates in the Euclidean space. We develop an algorithm SizeUBr, which executes repeatedly two interleaved threads.

The **first thread** keeps combining stations into groups such that eventually, for any box C of the pivotal grid, all stations located in C form one group. Moreover, each group has the leader, and eventually each station should be aware of (i) which group it belongs to, (ii) which station is the leader of that group, and (iii) which stations belong to that group. Upon waking up, each station forms a group with a single element (itself), and the groups increase gradually by merging. The merging process builds upon the following observation. Let σ be the smallest distance between two stations and let u, v be the closest stations. Thus, there is at most one station in each box of the grid $G_{\sigma/\sqrt{2}}$. Then, if u transmits a message and no other station in distance $d \cdot \sigma$, for some constant d, transmits at the same time, then v can hear that message (see Prop. 2). Using a $(N, (2d+1)^2)$-strongly-selective family as a broadcast schedule S on the set of leaders of groups, c.f., [16], one can assure that such a situation occurs in each $O(\log N)$ rounds. If u can hear v and v can hear u during such a schedule, the groups of u and v can be merged. In order to coordinate the merging process, we implicitly build a matching among pairs (u, v) such that u can hear v and v can hear u during execution of S.

The **second thread** is supposed to guarantee that the broadcast message is transmitted from boxes containing informed stations to their neighbors. Each station determines its temporary ID (TID) as the rank of its ID in the set of IDs in its group. Using these TIDs, the stations apply round-robin strategy. Thus, if each group corresponds to all stations in the appropriate box, transmissions are successful (see Prop. 1), and thus they guarantee that neighbors of a box containing informed stations will also contain informed stations.

The main problem with implementation of these ideas is that, as long as there are many groups inside a box, transmissions in the second thread may cause unwanted interferences. Another problem is that the set of stations attending the protocol changes gradually, when new stations become informed and can join the execution of the protocol. These issues are managed by measuring the progress of a protocol using amortized analysis. The details of the implementation and analysis can be found in the full version of the paper.

Theorem 1. *Algorithm* SizeUBr *performs broadcasting in each n-node network in $O(n \log n)$ rounds, in the setting without local knowledge.*

2.2 Degree Dependent Algorithm

In this section we present the algorithm GenBroadcast whose complexity is optimized with respect to maximal degree of the communication graph. The core of this algorithm is a leader election procedure which, given a set of stations V, chooses exactly one station (the leader) in each box C of the pivotal grid containing at least one element from V. This procedure works in $O(\log n \cdot \log N) = O(\log^2 n)$ rounds and it is executed repeatedly during GenBroadcast. The set of stations attending a particular leader election execution consists of all stations which received the broadcast message and have not been chosen leaders of their boxes in previous executions of the leader election procedure. Moreover, at

Algorithm 1. LeaderElection(V, n)

```
For each v ∈ V: cand(v) ← true;
for i = 1, ..., log n + 1 do                                        ▷ Elimination
    for j, k ∈ [0, 2] do
        Execute S twice on the set:
         {w ∈ V | cand(w) = true and w ∈ C(j', k')
           such that (j' mod 2, k' mod 2) = (j, k)};
        Each w ∈ V determines and stores X_w during the first execution of S, and
         X_v, for each v ∈ X_w, during the second execution of S;
        for each v ∈ V do
            u ← min(X_v);
            if X_v = ∅ or v > min(X_u ∪ {u}) then cand(v) ← false; ph(v) ← i;
For each v ∈ V: state(v) ← active;                                  ▷ Selection
for i = log n, (log n) − 1, ..., 2, 1 do
    V_i ← GranLeaderElection({v ∈ V |
        ph(v) = i, state(v) = active}, 1/n);                        ▷ V_i – leaders
    Each element v ∈ V_i sets state(v) ← leader and
      transmits successfully using constant dilution (see Prop. 1);
    Simultaneously, for each v ∈ V which can hear u ∈ box(v): state(v) ← passive.
```

the end of each execution of the leader election procedure, each leader chosen in that execution transmits a message successfully — this can be done in a constant number of rounds, by using d-dilution with appropriate constant d (c.f., Prop. 1). In this way, each station receives the source message after $O(D\Delta \log^2 n)$ rounds. (Note that there are at most Δ stations in a box of the pivotal grid.)

In the following, we describe the leader election algorithm — its pseudo-code is presented as Algorithm 1. We are given a set of stations V of size at most n. The set V is not known to stations, each station knows merely whether it belongs to V or it does not belong to V. In the algorithm, we use (N, e)-ssf S of size $s = O(\log N)$, where $e = (2d + 1)^2$ and d is the constant depending merely on the parameters of the model, the same as in Section 2.1 (see also Prop. 2). Let X_v, for a given execution of S be the set of stations which belong to box(v) and v can hear them during that execution.

The following proposition combines properties of ssf with Prop. 2.

Proposition 3. *For each $\alpha > 2$ and $\varepsilon > 0$, there exists a constant k satisfying the following property. Let W be a 3-diluted (wrt the pivotal grid) set of stations and let C be a box of the pivotal grid. If $\min_{u,v \in C \cap W} = x \le r/n$ and $dist(u, v) = x$ for some $u, v \in W$ such that $box(u) = box(v) = C$, then v can hear the message from u during an execution of a (N, k)-ssf on W.*

The leader election algorithm consists of two stages. The first stage gradually eliminates elements from the set of candidates for the leaders of boxes in consecutive executions of the ssf S in the first for loop. Therefore, we call this stage *Elimination*. Let *phase* l of Elimination stage denote the executions of S for $i = l$. Each station v "eliminated" in phase l has assigned the value $ph(v) = l$. Let $V(l) = \{v \mid ph(v) > l\}$ and $V_C(l) = \{v \mid ph(v) > l \text{ and box}(v) = C\}$ for $l \in \mathbb{N}$

and C being a box of the pivotal grid. That is, $V_C(l)$ is the set of stations from C which are *not* eliminated until phase l. The key property of the sets $V_C(l)$ is that $|V_C(l+1)| \leq |V_C(l)|/2$ and the granularity of $V_C(l_C^\star)$ is *smaller* than n for each box C and $l \in \mathbb{N}$, where $l_C^\star$ is the largest $l \in \mathbb{N}$ such that $V_C(l)$ is not empty. Therefore, we can choose the leader of each box C by applying (simultaneously in each box) the granularity dependent leader election algorithm GranLeader-Election, described later in Section 3.2 on the set $V_C(l_C^\star)$ and with upper bound n on granularity. Note that we can elect the leaders in $O(\log N) = O(\log n)$ rounds in this way. However, the stations in C do not necessary know the value of $l_C^\star$. Therefore, the second stage (called *Selection*) applies the granularity dependent leader election on $V(\log n)$, $V(\log n - 1)$, $V(\log n - 2)$ and so on. When the leader of a box C is chosen, all stations in C become silent (state *passive* in line 18), i.e., they do not attend the following executions of GranLeaderElection. It is important that a station becomes silent after the leader of its box is chosen, since granularity of $V_C(i)$ might be larger than n for $i < l_C^\star$. Activity of stations from such a box C for $i < l_C^\star$ during the Selection stage could cause large interferences preventing other boxes from electing leaders.

Recall that each leader broadcasts successfully at the end of the execution of LeaderElection in which it is elected. If each station attends consecutive leader election executions until it becomes a leader in its box, the broadcasting message is transmitted from a box C to all its neighbors in $O(\Delta \log^2 n)$ rounds, since there are at most Δ station in each box of the pivotal grid. Therefore, we obtain a GeneralBroadcast algorithm providing the following result.

Theorem 2. *Algorithm GeneralBroadcast completes broadcast in $O(D\Delta \log^2 n)$ rounds in any network without local knowledge.*

2.3 Lower Bounds

Theorem 3. *There exist: (i) an infinite family of n-node networks requiring $\Omega(n \log n)$ rounds to accomplish deterministic broadcast, and (ii) for every $D\Delta = O(n)$, an infinite family of n-node networks of diameter D and maximum degree Δ requiring $\Omega(D\Delta)$ rounds to accomplish deterministic broadcast.*

Proof (Sketch). We describe a family of networks $\mathcal{F}$ such that broadcasting requires time $\Omega(D \log N)$. By L_i we denote the set of stations in distance i from the source in the communication graph. Each element of $\mathcal{F}$ is formed as a sequential composition of D networks $V_1, \ldots, V_D$ of eccentricity 3 each, such that:

- the source s is connected with two nodes v_1, v_2 in L_1 with arbitrary IDs and fixed positions;
- v_1, v_2 are connected with w, the only element of L_2, and satisfy the condition:

$$P \cdot \mathrm{dist}(v_1, w)^{-\alpha} = P \cdot \mathrm{dist}(v_2, w)^{-\alpha} - \mathcal{N}/2 \ . \qquad (2)$$

Sequential composition of networks $V_1, \ldots, V_D$ stands for identifying the element w of network component V_i with the source s of network component V_{i+1}.

Note that if v_1 and v_2 transmit simultaneously in a network component V_i, the message is *not* received by w. Using simple counting argument, one can force such choice of IDs of v_1 and v_2 that $\Omega(\log N)$ rounds are necessary until a round in which exactly one of v_1, v_2 transmits successfully a message to w under the SINR model of weak devices. Since $D = \Theta(n)$ in the above construction and $\log n = \Theta(\log N)$, the bound $\Omega(n \log N)$ holds.

The above proof can be extended to obtain lower bounds $\Omega(D\Delta)$, by considering the following class of network components V_i: the source s, located in the origin point $(0,0)$, is the only element of L_0; L_1 consists of Δ nodes $v_0, \ldots, v_{\Delta-1}$, where the position of v_i is $(\gamma \cdot \frac{i}{\Delta}, \gamma)$ for $0 \leq i \leq \Delta - 1$; and L_2 contains only one node w_j with coordinates $(\gamma \cdot \frac{j}{\Delta}, \gamma + r)$, i.e., w_j can receive a message only from v_j. □

This result can also be transformed to the case of randomized algorithms. We sketch an idea of these transformations by considering networks from the family $\mathcal{F}$ described in the proof of Theorem 3. Recall that each element of the layer L_1 should transmit as the only element of L_1 in order to guarantee that the only element of L_2 is informed, regardless of its location. However, by simple counting arguments, the expectation of the number of steps after which some of elements of L_1 transmit as the only one is $\Omega(\log n)$ or $\Omega(\Delta)$, respectively.

3 Algorithms for Networks with Local Knowledge

In this section we assume that each station knows n, N as well as IDs and locations of all stations in its range area. We start with presenting a generic algorithmic scheme. Next, we describe an algorithm for networks with additionally known granularity bound g. Finally we provide a solution for the general setting when granularity g is not known in advance.

3.1 Generic Algorithmic Scheme

In the first round the source sends the broadcast message. Then, we repeat the generic procedure Inter-Box-Bdcst, whose ith repetition is aimed at transmitting the broadcast message from boxes of the pivotal grid containing at least one station that has received the broadcast message in the previous execution of Inter-Box-Bdcst (or from the source) to boxes which are their neighbors. The specific implementation of procedure Inter-Box-Bdcst depends on the considered setting.

Each station v is in state $s(v)$, which may be equal to one of the following three values: asleep, active, or idle. At the beginning, the source sends the source message and all stations of its box in the pivotal grid set their states to active, while all the remaining stations are in the asleep state. The states of stations change only at the end of Inter-Box-Bdcst, according to the following rules:

Rule 1: All stations in state active change their state to idle.
Rule 2: A station u changes its state from asleep to active if it has received the broadcast message from a station v in the current execution of Inter-Box-Bdcst

such that either v was in state active or v belongs to the same box of the pivotal grid as u. That is, let C be a box of the pivotal grid, let $u \in C$ be in state asleep at the beginning of Inter-Box-Bdcst. The only possibility that u receives a message and it does not change its state from asleep to active at the end of Inter-Box-Bdcst is that each message received by u is sent by a station v which is in state asleep when it sends the message and $v \notin C$.
The intended properties of an execution of Inter-Box-Bdcst are:

(I) For each box C of the pivotal grid, states of all stations in C are equal.
(P) The broadcast message is (successfully) sent from each box C containing stations in state *active* to all stations located in boxes which are neighbors of C.

The following proposition easily follows from the above stated properties.

Proposition 4. *If (I) and (P) are satisfied, the source message is transmitted to the whole network in $O(D \cdot T)$ rounds, where T is the number of rounds in a single execution of Inter-Box-Bdcst.*

3.2 A Granularity-Dependent Algorithm

First, we develop a broadcasting algorithm with known granularity g of a network. The main ingredient of this protocol is a leader election algorithm, called GranLeaderElection(A, g), which, given a set of stations A chooses the leader in each box of the pivotal grid containing stations from A (at the beginning, each station knows only whether it belongs to A or not). The idea of the leader election procedure is as follows. Granularity g implies that each station is the leader of a box of G_x, where $x = \gamma/h$ for $h = \min(2^i \mid 2^i \geq g)$. Then, leaders of boxes of $G_{2^i x}$ are chosen among leaders of boxes of $G_{2^{i-1}x}$ for $i = 1, 2, \ldots, \lceil \log h \rceil$ in constant number of rounds with help of Prop. 1. Thus, leaders in boxes of the pivotal grid can be chosen in $O(\log g)$ rounds.

Given the above (local) leader election procedure, the procedure Inter-Box-Bdcst is implemented as follows. For each direction $(d_1, d_2) \in \text{DIR}$, the leaders are elected in all boxes among station in state active which have neighbors in the direction (d_1, d_2). Then, these leaders send messages successfully using dilution (see Prop. 1). Moreover, since each station knows all stations in its box, the station with smallest ID among newly informed in each box sends the broadcast message which is delivered to all stations from its box. In this way, the procedure Inter-Box-Bdcst satisfying the invariants (I) and (P) working in time $O(\log g)$ is obtained which gives the broadcasting algorithm working in time $O(D \log g)$.

Theorem 4. *Algorithm* GranUBr *accomplishes broadcast in any n-node network of diameter D and granularity g in $O(D \log g)$, in the setting with local knowledge.*

3.3 General Algorithm

In this section we develop Algorithm DiamUBr, which also builds on the generic scheme from Section 3.1. Procedure Inter-Box-Bdcst required by the generic algorithm is implemented as in Section 3.2, the only difference is that GranLeaderElection with complexity $O(\log g)$ is replaced with the procedure LeaderElection from Section 2.2 (Alg. 1). By Prop. 4, and by the round complexity $O(\log^2 n)$ of algorithm LeaderElection, we obtain the following result.

Theorem 5. *Algorithm* DiamUBr *completes broadcast in any n-node network of diameter D in $O(D \log^2 n)$ rounds, in the setting with local knowledge.*

References

1. Avin, C., Emek, Y., Kantor, E., Lotker, Z., Peleg, D., Roditty, L.: Sinr diagrams: towards algorithmically usable sinr models of wireless networks. In: PODC, pp. 200–209 (2009)
2. Clementi, A.E.F., Monti, A., Silvestri, R.: Selective families, superimposed codes, and broadcasting on unknown radio networks. In: SODA, pp. 709–718 (2001)
3. Czumaj, A., Rytter, W.: Broadcasting algorithms in radio networks with unknown topology. In: FOCS, pp. 492–501 (2003)
4. Dessmark, A., Pelc, A.: Broadcasting in geometric radio networks. J. Discrete Algorithms 5(1), 187–201 (2007)
5. Emek, Y., Gasieniec, L., Kantor, E., Pelc, A., Peleg, D., Su, C.: Broadcasting in udg radio networks with unknown topology. Distributed Computing 21(5), 331–351 (2009)
6. Emek, Y., Kantor, E., Peleg, D.: On the effect of the deployment setting on broadcasting in euclidean radio networks. In: PODC, pp. 223–232 (2008)
7. Goldsmith, A.J., Wicker, S.B.: Design challenges for energy-constrained ad hoc wireless networks. IEEE Wireless Communications 9(4), 8–27 (2002)
8. Goussevskaia, O., Moscibroda, T., Wattenhofer, R.: Local broadcasting in the physical interference model. In: DIALM-POMC, pp. 35–44 (2008)
9. Goussevskaia, O., Pignolet, Y.A., Wattenhofer, R.: Efficiency of wireless networks: Approximation algorithms for the physical interference model. Foundations and Trends in Networking 4(3), 313–420 (2010)
10. Hobbs, N., Wang, Y., Hua, Q.-S., Yu, D., Lau, F.C.M.: Deterministic distributed data aggregation under the SINR model. In: Agrawal, M., Cooper, S.B., Li, A. (eds.) TAMC 2012. LNCS, vol. 7287, pp. 385–399. Springer, Heidelberg (2012)
11. Jurdzinski, T., Kowalski, D.R.: Distributed backbone structure for algorithms in the SINR model of wireless networks. In: Aguilera, M.K. (ed.) DISC 2012. LNCS, vol. 7611, pp. 106–120. Springer, Heidelberg (2012)
12. Jurdzinski, T., Kowalski, D.R.: On the complexity of distributed broadcasting and MDS construction in radio networks. In: Baldoni, R., Flocchini, P., Binoy, R. (eds.) OPODIS 2012. LNCS, vol. 7702, pp. 209–223. Springer, Heidelberg (2012)
13. Jurdzinski, T., Kowalski, D.R., Stachowiak, G.: Distributed deterministic broadcasting in wireless networks of weak devices under the sinr model. CoRR, abs/1210.1804 (2012)
14. Jurdzinski, T., Kowalski, D.R., Stachowiak, G.: Distributed deterministic broadcasting in uniform-power ad hoc wireless networks. CoRR, abs/1302.4059 (2013)

15. Kesselheim, T., Vöcking, B.: Distributed contention resolution in wireless networks. In: Lynch, N.A., Shvartsman, A.A. (eds.) DISC 2010. LNCS, vol. 6343, pp. 163–178. Springer, Heidelberg (2010)
16. Kowalski, D.R.: On selection problem in radio networks. In: PODC, pp. 158–166 (2005)
17. Kowalski, D.R., Pelc, A.: Broadcasting in undirected ad hoc radio networks. Distributed Computing 18(1), 43–57 (2005)
18. Richa, A., Scheideler, C., Schmid, S., Zhang, J.: Towards jamming-resistant and competitive medium access in the sinr model. In: Proc. 3rd ACM Workshop on Wireless of the Students, by the Students, for the Students, S3 2011, pp. 33–36 (2011)
19. Schmid, S., Wattenhofer, R.: Algorithmic models for sensor networks. In: IPDPS. IEEE (2006)
20. Yu, D., Wang, Y., Hua, Q.-S., Lau, F.C.M.: Distributed local broadcasting algorithms in the physical interference model. In: DCOSS, pp. 1–8 (2011)