

Radio Aggregation Scheduling*

Rajiv Gandhi¹, Magnús M. Halldórsson^{2,3}, Christian Konrad², Guy Kortsarz¹,
and Hoon Oh¹

¹ Department of Computer Science, Rutgers University Camden, NJ.
{hoon,rajivg,guyk}@camden.rutgers.edu

² ICE-TCS, School of Computer Science, Reykjavik University, Iceland,
{mmh,christiank}@ru.is

³ RIMS, Kyoto University, Japan

Abstract. We consider the aggregation problem in radio networks: find a spanning tree in a given graph and a conflict-free schedule of the edges so as to minimize the latency of the computation. While a large body of literature exists on this and related problems, we give the first approximation results in graphs that are not induced by unit ranges in the plane. We give a polynomial-time $\tilde{O}(\sqrt{dn})$ -approximation algorithm, where \bar{d} is the average degree and n the number of vertices in the graph, and show that the problem is $\Omega(n^{1-\epsilon})$ -hard (and $\Omega((\bar{d}n)^{1/2-\epsilon})$ -hard) to approximate even on bipartite graphs, for any $\epsilon > 0$, rendering our algorithm essentially optimal. We target geometrically defined graph classes, and in particular obtain a $O(\log n)$ -approximation in interval graphs.

1 Introduction

Wireless sensor networks consist of autonomous sensors that typically monitor physical or environmental conditions. They use wireless communication to cooperatively aggregate the recorded data and forward it to a central location, the sink. The information desired is commonly in the form of a *compressible* function, such as “max” or “average”, in which in-network processing can be used to speed up the processing and greatly reduce transmission energy. At the same time, interference from simultaneous transmissions must be managed for successful reception.

In this paper, we consider the data aggregation problem in general graphs, or radio networks. The objective is to minimize the *latency*, or the longest time it takes for any message to reach the sink. The task is two-fold: a) to construct a directed spanning tree, an *in-arborescence*, and b) to form a *conflict-free schedule* of the transmissions (the edges) that obeys the ordering of the arborescence. A schedule is conflict free if whenever a node is to receive a message, none of its

* Gandhi and Oh are partly supported by NSF grants CCF 1218620 and CCF 1433220. Halldórsson and Konrad are supported by Icelandic Research Fund grant-of-excellence no. 120032011. Kortsarz is partly supported by NSF grant CCF 1218620.

other neighbors also transmit (causing interference), and a node can transmit to only one of its neighbors at a time.

This problem, which we dub *Radio Aggregation Scheduling* (RAS), has been widely studied under the name *Minimum Latency Aggregation Scheduling* in the wireless networking literature. Most of the existing works consider the setting where nodes are points in the plane with a fixed transmission radius, which corresponds to the case of *unit disc graphs* (UDG). It is, however, well known that wireless environments are *always* much more complicated [1, 21] — unless operating in vacuum in outer space. One popular approach in recent years has been to switch to the SINR model of interference, which is known to add more realism. However, its standard form also makes strong assumptions about the geometric nature of communicability and interference and thus ignores the unpredictability seen in practice. To go beyond these assumptions, we initiate here the study of aggregation in more pessimistic models, starting with general graphs. To emphasize the distinction of using graphs rather than planar positions, we refer to the problem as RAS.

By reversing the direction of the aggregation process, we can also view it as a *broadcasting* problem where:

1. [*one-on-one*] a node can only talk to one other node at a time, while
2. [*interference from neighbors*] a node can hear from its neighbor only if none of its other neighbors transmit.

We refer to this communication model as the *radio-unicast* model. It relates closely to two other classic broadcasting problems: *telephone broadcast*, where (1) holds but there are no conflicts from other neighbors (in essence, modeling aggregation in wired networks); and *radio broadcast*, where (2) holds, but a node can transmit to all its neighbors in the same time slot. As we shall see, however, RAS is significantly harder to solve in general than either of these problems.

In the telephone model, in each communication round, the successful transmissions form a (directed) matching. In the radio-unicast model, successful transmissions form what we call a *RAS-legal matching* (see Sec. 2 for precise definitions). For any two edges (s_1, r_1) and (s_2, r_2) in a RAS-legal matching connecting senders s_1, s_2 to receivers r_1, r_2 , it is required that neither (s_1, r_2) nor (s_2, r_1) are edges contained in the input graph, thus excluding all potential interference. This is closely related to the notion of an *induced matching*. A matching is induced if the edges of the subgraph induced by the matched vertices are precisely the edges of the matching. A RAS-legal matching hence lies somewhere between a matching and an induced matching, see Fig. 1.

Previous Work on RAS. All previous works on RAS consider the setting where nodes are points located in the plane with unit length transmission radii [7, 26, 25, 3, 15]⁴. This corresponds to the study of RAS in unit disc graphs, which has been shown to be NP-complete [7]. All algorithms known for unit disc graphs compute aggregation schedules of lengths $\Theta(\text{Diam} + \Delta)$, where *Diam* is the diameter of

⁴ In [15], unit interval graphs as well as grids and tori are considered, which are all subclasses of unit disc graphs.




Fig. 1. Left: A matching is a subset of vertex-disjoint edges. Center: The edges of the graph induced by the vertices of an induced matching are precisely the edges of the induced matching. Right: In a RAS-legal matching, every receiver is connected to precisely one sender.

the input graph and Δ the maximal degree. Since every aggregation schedule is of length at least $Diam$, these algorithms constitute $O(\Delta)$ -approximation algorithms which only give trivial approximation guarantees in graphs with large maximum degree (e.g. if $\Delta = \Theta(n)$). Despite the considerable effort put into the study of RAS on unit disc graphs, no better approximation ratios are known.

One difficulty in obtaining improved approximation ratios in unit disc graphs is to bound the length of an optimal aggregation schedule OPT in terms of properties of the input graph. For instance, in unit interval graphs, it is known that $OPT = \Omega(Diam + \omega(G))$, where $\omega(G)$ is the clique number (size of the largest clique) of the input graph [15]. It is also known how to compute an aggregation schedule of length $O(Diam + \omega(G))$, which hence constitutes an $O(1)$ -approximation algorithm (in [15], a 2-approximation is obtained). No interesting bounds on OPT are known for unit disc graphs or any other non-trivial graph class.

Our Contributions. We initiate a systematic study of RAS, starting with general graphs. We prove that it is NP-hard to approximate RAS within a factor of $n^{1-\epsilon}$ (**Theorem 4**) and $(\bar{d}n)^{1/2-\epsilon}$ (**Corollary 1**) even in bipartite graphs, for any $\epsilon > 0$, where n is the number of vertices of the input graph and \bar{d} is the average degree. On the positive side, we present a $\tilde{O}(\sqrt{\bar{d}n})$ -approximation⁵ algorithm for sparse general graphs (**Theorem 5**), almost matching our lower bound.

Next, we are interested in whether improved algorithms can be obtained for geometrically defined graph classes that contribute to metric-sensitive models of actual wireless environments. We focus here on interval graphs. They can be seen as one-dimensional projections of disc graphs that capture the aspect of different radii, and we present a highly non-trivial $O(\log n)$ -approximation algorithm (**Theorem 6**). The key part of our analysis is the identification of subgraphs that provide interesting lower bounds on the length of an optimal aggregation schedule.

Further Related Work. Aggregation problems have been extensively studied in the wireless literature; see the surveys [12, 17]. As previously mentioned, RAS

⁵ We use the notation $\tilde{O}(\cdot)$, which equals the usual $O(\cdot)$ notation where all poly-logarithmic factors are ignored.

has been considered in unit disc graphs [7, 26, 25, 3, 15] and $O(\Delta)$ -approximation algorithms are known. Furthermore, it has also been shown that, in unit disc graphs, if the interference radius is strictly larger than the transmission radius, then constant factor approximations can be obtained [25]. Optimal algorithms are known for grids and tori [14]. In trees, RAS is equivalent to the telephone broadcast problem, which has a textbook dynamic programming solution [19, Prob. 6.16]. This exhausts the list of previous work known on RAS.

A different setting for aggregation problems is where the nodes are located at points in the plane and can use power control to reach any other node. Kesselman and Kowalski [18] showed that aggregation can then be achieved in $O(\log n)$ slots. If interference and transmissions follow the geometric SINR model, Moscibroda and Wattenhofer [23] showed that poly-logarithmic slots suffice, which was improved to optimal $O(\log n)$ [16].

For broadcast in the radio model, Chlamtac and Weinstein [8] proved the first upper bound of $O(\text{Diam} \cdot \log^2 n)$, with Diam being the diameter of the graph, which was improved to $O(\text{Diam} \cdot \log n + \log^2 n)$ soon afterwards by Bar-Yehuda et al. [5]. The best bound known on the number of rounds, $O(\text{Diam} + \log^2 n)$, given by Kowalski and Pelc [22], is optimal in light of results of Alon et al. [2] and Elkin and Kortsarz [10].

The first approximation for telephone broadcast was an additive $O(\sqrt{n})$ approximation [20]. This was improved to a multiplicative $O(\log^2 n)$ -factor by [24], and then to $O(\log n)$ in [4]. The best approximation known for the problem is $O(\log n / \log OPT)$ [11], which is $O(\log n / \log \log n)$, since $OPT \geq \log_2 n$ always holds. The best lower bound known is a factor $3 - \epsilon$, given in [9].

Outline of the Paper. We give formal definitions of our problems in Sec. 2. Then, in Sec. 3, we present our hardness results for general graphs, and in Sec. 4, we present our algorithm for sparse general graphs. Finally, in Sec. 5, interval graphs are discussed.

Due to space restrictions, the proofs of lemmas, theorems, claims and observations marked by (*) can be found in the appendix.

2 Problem Definition and Notations

Radio Aggregation Scheduling. We are given as input a graph $G = (V, E)$ and a node $s \in V$ which is the *sink* node of the aggregation problem. We view G as a *bidirected graph*, i.e., all edges appear directed in both directions.

We seek a *schedule*, which is a sequence M_1, M_2, \dots, M_t of directed matchings in G . The union $\cup_i M_i$ of these matchings induces a directed spanning tree (*in-arborescence*) A directed toward s . Each matching M_i corresponds to a set of transmissions that can be successful simultaneously; namely, each matching must be *RAS-legal* in G : if $(u, v), (w, z) \in M_i$ then $(u, z), (w, v) \notin E(G)$. Finally, the edges of A occur in the matchings in order of precedence induced by the arborescence: if $(u, v) \in M_i$ and $(v, w) \in M_j$ then $i < j$. Namely, a node can only forward its message once it has heard from all of its children. Then an optimal

solution to the Radio Aggregation Scheduling problem (RAS) is a schedule of minimal length.

Broadcasting in the Radio-unicast Model. Since reversing the slots of an aggregation schedule gives a broadcast, and vice versa, both viewpoints can be used to tackle RAS. In the broadcast version of the problem, node $s \in V$ is the *source* node and holds a message that is to be sent to all other nodes $V \setminus \{s\}$ in the graph. In each round, we seek a RAS legal matching between the *informed nodes* (those that know the message) and the *uninformed nodes* (those that don't know the message yet). Initially, there is only a single informed node, the source node s . When an uninformed node receives the message, it joins the set of informed nodes and can serve as a sender in upcoming rounds. We denote this communication model where each round induces a RAS-legal matching as the *radio-unicast* model. An optimal solution to the broadcasting problem then is a broadcasting schedule that informs all nodes in the minimal number of rounds.

It turns out that the broadcasting perspective of RAS is more convenient when presenting our algorithms. All our algorithms solve the broadcasting problem in the radio-unicast model.

Notation. Let $G = (V, E)$ be the input graph. Unless stated differently, n denotes the number of vertices of G , \bar{d} the average degree, Δ the maximum degree, and $Diam$ the diameter. Those quantities may also appear as functions, e.g. $\Delta(H)$, $\bar{d}(H)$ and $Diam(H)$ denote the respective quantities of graph H .

We write $dist_G(u, v)$ for the number of hops between nodes u and v in graph G . Let $N_G(u)$ denote the set of neighbors of vertex u in G , and for a set S of vertices, let $N_G(S) = (\cup_{u \in S} N_G(u)) \setminus S$. We write $deg_G(u)$ the degree of u in G . Furthermore, for a graph G , we denote its vertex set by $V(G)$ and its edge set by $E(G)$. Given a subset of vertices $U \subseteq V$, we denote the subgraph of G induced by the vertices U by $G[U]$.

3 Approximation Hardness of RAS

In this section, we prove that RAS is hard to approximate within factors $n^{1-\epsilon}$ (Theorem 4) and $(\bar{d}n)^{1/2-\epsilon}$ (Corollary 1), for every $\epsilon > 0$. Before giving our lower bound construction, we introduce further required notations and definitions.

Further Definitions. We denote the chromatic number of a graph G with $\chi(G)$, and the independence number (size of a maximum independent set) with $\alpha(G)$. Our lower bound construction relies on *semi-induced matchings* and a specific *graph product* that we discuss first.

A matching is called an *induced matching* if there is no edge from one endpoint of an edge in the matching to an endpoint of another edge in the matching. The *semi-induced matching* has a general definition (see [6]) but we only give the definition for bipartite graphs that is simpler and all we need.

Definition 1 (Semi-induced Matching). Let $G = (U, V, E)$ be a bipartite graph with a total ordering u_1, \dots, u_n of U . A semi-induced matching is a matching so that if (u_i, a) and (u_j, b) are in the matching and $i < j$, then there is no edge between u_j and a .

Let $\text{IM}(G)$ be the size of the largest induced matching of G and $\text{SIM}(G)$ the size of the largest semi-induced matching. Observe that $\text{IM}(G) \leq \text{SIM}(G)$, for any graph G .

Next, we make use of the following graph product:

Definition 2 (Inclusive Graph Product). *The inclusive graph product of $G = (V, E)$ and $H = (V', E')$, denoted by $G \vee H$, has vertices $\{(x_G, x_H) \mid x_G \in V, x_H \in V'\}$. A pair of vertices $(x_G, x_H) \in V(G \vee H)$ and $(y_G, y_H) \in V(G \vee H)$ is connected iff $(x_G, y_G) \in E$ or $(x_H, y_H) \in E'$.*

We denote $G^k = G \vee G \vee \dots \vee G$ when there are k copies of G . This graph has n^k vertices.

The following equalities are folklore for the specific product we chose:

$$\chi(G^k) = \chi(G)^k, \tag{1}$$

$$\alpha(G^k) = \alpha(G)^k. \tag{2}$$

Intermediate Problem: Induced Matching Cover. We shall consider a problem on bipartite graphs that is closely related to RAS. Given a bipartite graph $B = (U, V, E)$, let $\text{IMCOV}(B)$ denote the minimum number of induced matchings that together contain (or cover) all the vertices of V . Suppose that nodes U are informed and nodes V are uninformed. Then, it takes precisely $\text{IMCOV}(B)$ rounds in order to inform V . This is summarized in Observation 1.

Observation 1 (*) *Let $B = (U, V, E)$ be a bipartite graph. Suppose all the vertices in U know the message. Then, the minimum number of rounds it takes to inform V in the radio-unicast model equals $\text{IMCOV}(B)$.*

Lower Bound Construction. In order to prove our hardness result, we will use the construction of Feige and Kilian [13] which shows that it is hard to determine whether a graph G on n vertices has small chromatic number $\chi(G) \leq n^\epsilon$ (“yes instance”) or has a small independence number $\alpha(G) \leq n^\epsilon$ (“no instance”), for any $\epsilon > 0$.

Let G be a graph on n vertices as used in the construction of Feige and Kilian. From G , using a construction similar to the one in [6], we will construct a bipartite graph $B_e(G^k)$ on $\Theta(n^k)$ vertices so that:

$$\text{IMCOV}(B_e(G^k)) \leq \chi(G), \text{ and} \tag{3}$$

$$\text{IM}(B_e(G^k)) \leq k \cdot n + \alpha(G)^k. \tag{4}$$

Suppose now that one bipartition of $B_e(H^k)$ is informed and the other one is uninformed. Then, if G is a “yes instance” (i.e. it has small chromatic number), the whole graph can be informed quickly using Inequality 3 and Observation 1.

Suppose now that G is a “no instance” (i.e. it has small independence number). Then, by Inequality 4, $\text{IM}(B_e(G^k))$ is small, too. Using the obvious relationship $\text{IMCOV}(B_e(G^k)) \geq |V(B_e(G^k))|/\text{IM}(B_e(G^k))$, we conclude that $\text{IMCOV}(B_e(G^k))$ is large which implies that informing the whole graph takes many rounds.

The previous gap-reduction argument is made rigorous in the following. To this end, for a graph G , we first define the graph $B_e(G)$.

Definition 3. Given a graph $G = (V, E)$, the graph $B(G) = (V, \bar{V}, E_B)$ is a bipartite graph with a copy of V on each side. There is an edge $(v, \bar{u}) \in E_B$ if $(v, u) \in E$. The graph $B_e(G)$ results from $B(G)$ by adding the edges $\{(v, \bar{v}) : v \in V\}$.

Next, we prove Inequalities 3 and 4 in Claim 2 and Claim 3, respectively.

Claim 2 (*) Let $G = (V, E)$ be a graph. Then, $B_e(G) = (V, \bar{V}, E')$ can be decomposed into $\chi(G)$ induced matchings that are pairwise disjoint and together contain all of \bar{V} , i.e., $\text{IMCOV}(B_e(G)) \leq \chi(G)$.

Claim 3 (*) Let G be a graph, k an integer. Then, $\text{IM}(B_e(G^k)) \leq k \cdot n + \alpha(G)^k$.

Finally, we prove our hardness results in Theorem 4 and Corollary 1.

Theorem 4. The RAS problem is hard to approximate on bipartite graphs within a factor of $N^{1-\delta}$, for any $\delta > 0$, where N is the number of vertices.

Proof. We use the gap reduction of Feige and Kilian [13]: for any $\epsilon > 0$, it is hard to distinguish between the case (“yes” instance) when a graph G is n^ϵ -colorable, i.e., when $\chi(G) \leq n^\epsilon$, and the case (“no” instance) when there is no independent set of size at least n^ϵ , i.e., $\alpha(G) < n^\epsilon$.

Let ϵ be such that $1/\epsilon = 2\lceil 1/\delta \rceil$, and let $k = 1/\epsilon$. Consider $B_e(G^k) = (V_k, \bar{V}_k, E_k)$ and let H_k be the graph obtained by adding to $B_e(G^k)$ a complete binary tree of depth $O(\log |V_k|)$ whose set of leaves contains V_k . H_k is clearly bipartite, too. We show that it is hard to approximate the number of rounds in a RAS schedule of H_k .

Suppose that the root of the binary tree is the source node of the broadcast problem. Let OPT denote the length of a shortest broadcast schedule. Observe that informing the nodes of the complete binary tree, and thus also the nodes in V_k , requires only $O(\log n)$ slots. Informing \bar{V}_k after V_k has been informed takes $\text{IMCOV}(B_e(G^k))$ rounds, by Obs. 1. Thus, $OPT = \text{IMCOV}(B_e(G^k)) + O(\log n)$.

If G is a yes-instance, $\chi(G) \leq n^\epsilon$, so by Claim 2 and Inequality 1,

$$\text{IMCOV}(B_e(G^k)) \leq \chi(G^k) = \chi(G)^k \leq n^{k\epsilon} = n,$$

and hence

$$OPT = \text{IMCOV}(B_e(G^k)) + O(\log n) = O(n).$$

If G is a no-instance, $\alpha(G)^k \leq n^{k\epsilon} = n$, so by Claim 3, $\text{IM}(B_e(G^k)) = O(n)$, and

$$OPT \geq \text{IMCOV}(B_e(G^k)) \geq \frac{|V_k|}{\text{IM}(B_e(G^k))} = \Omega(n^{k-1}).$$

The ratio between the bounds for the two cases is $\Omega(n^{k-2})$. Recalling that the size of H_k is given by $N = |H_k| = \Theta(n^k)$, we get that the approximation hardness is $\Omega(n^{k-2}) = \Omega(N/n^2) = \Omega(N^{1-\frac{2}{k}}) = \Omega(N^{1-\delta})$. \square

Corollary 1 (*). *The RAS problem is hard to approximate on bipartite graphs within a factor of $(\bar{d}N)^{\frac{1}{2}-\delta}$, for any $\delta > 0$, where N is the number of vertices.*

Corollary 1 renders our $\tilde{O}(\sqrt{\bar{d}n})$ -approximation algorithm that we present in the next section essentially best possible.

The graphs used in the proofs of Theorem 4 and Corollary 1 have a diameter of $O(\log n)$. By adding additional edges, their diameters can be reduced to 2. This shows that unlike in the radio model, broadcasting in the radio-unicast model is no easier in graphs of low diameter.

4 $\tilde{O}(\sqrt{\bar{d}n})$ -approximation Algorithm

We now present a $\tilde{O}(\sqrt{\bar{d}n})$ -approximation algorithm for RAS in general graphs $G = (V, E)$ with average degree \bar{d} . We consider the broadcasting perspective in the radio-unicast model. Before presenting our algorithm, we discuss simulation results that allow us to reuse existing algorithms designed for the telephone and the radio models.

Simulation Between Models. We derive now (rather straightforward) bounds on RAS schedules, utilizing its relationship to better studied broadcast problems.

Recall that in the *telephone model*, there are no conflicts if two neighbors of a node both transmit. However, a node can only transmit to one of its neighbors in a given round. In the *radio model*, when a node transmits, its message goes to all of its neighbors. However, an uninformed neighbor receives the message only if exactly one of its neighbors is transmitting in that round.

Our problem shares the unicast transmission rule with the telephone model and the reception conflicts with the radio model. Algorithms for these models can be simulated in our models.

Lemma 1. *A round in the radio (telephone) model can be simulated in Δ ($2\Delta - 1$) rounds in the unicast-radio model, respectively.*

Proof. Suppose a set S of nodes transmits in a given round in the radio model. Assume without loss of generality that the neighbors of each node are ordered in an arbitrary order. We can then simulate it with Δ rounds, where in round i , each node in S forwards the message to its i -th neighbor.

Suppose a directed matching M that corresponds to the transmissions in of a round in the telephone model. For each edge $e \in M$, there are at most $2(\Delta - 1)$ edges in M within distance 2 in G . We can color the edges in M “first-fit” using $2\Delta - 1$ colors so that each color class induces a RAS-legal matching. \square

Simulating the algorithm of Kowalski and Pelc for radio broadcast [22], and using Lemma 1, we obtain the following corollary.

Corollary 2. *There is a polynomial-time algorithm for RAS that computes an aggregation schedule of length $O(\Delta(\text{Diam} + \log^2 n))$ and thus constitutes a $O(\Delta + \Delta \log^2(n)/\text{Diam})$ -approximation algorithm.*

In the previous corollary, we used the fact that $Diam$ is a trivial lower bound on the length of an optimal schedule. In light of the hardness results in Sec. 3, this bound is close to best possible.

Complete binary trees with degrees at least $\log^2 n$ provide examples showing that the $O(\Delta(Diam + \log^2 n))$ bound of Corollary 2 generally cannot be improved.

Center Selection. Our algorithm uses as a subroutine solutions to a classic facility location problem. In CENTER SELECTION, we are given a graph $G = (V, E)$, a set $X \subseteq V$ of possible sites for centers, a set $C \subseteq V$ of clients, and a parameter k . We wish to find a set $S \subseteq X$ of k centers, such that the maximum distance from a client to the nearest center is minimized. For a set of centers $S \subseteq X$, let $\rho(G, S, C) := \max_{v \in C} dist_G(v, S)$ be the *covering radius* of S in G . The objective of CENTER SELECTION is to find an $S \subseteq X$ of cardinality k which minimizes $\rho(G, S, C)$.

A greedy algorithm, which we denote by GREEDY-CS(G, X, C, k), gives a 3-approximation to this problem.⁶

Lemma 2 (*). GREEDY-CS is a 3-*appr.* algorithm for CENTER SELECTION.

RAS scheme. In Algorithm 1, we present an algorithm for the broadcast problem in the radio-unicast model. We assume that the optimal value OPT (length of a shortest broadcast scheme) is known by the algorithm. This can be ensured e.g. by running the algorithm multiple times trying the different values $\{\log n, \dots, n\}$ for OPT and returning the best solution ($\log n$ is an obvious lower bound).

Let $s \in V$ be the source node. To keep the presentation simple, we assume that $\deg_G(s) \geq \sqrt{dn}$. If this is not the case, then we first inform an arbitrary node s' of degree at least \sqrt{dn} in at most OPT rounds which then takes the role of s . Clearly, the length of a minimum length schedule of the modified instance with source s' is at most by OPT longer than the length of a minimum length schedule with source node s . Hence, by solving the instance with source node s' , we may lose an additive $2 \cdot OPT$ term. However, since our obtained approximation factor is polynomial, this factor is negligible. Last, if no node of degree at least \sqrt{dn} exists, then we simply apply the simulation result of Corollary 2, and we immediately obtain an $\tilde{O}(\sqrt{dn})$ -approximation algorithm.

First, our algorithm informs the large-degree nodes, i.e., nodes L of degree at least $K = \sqrt{dn}$. The number of large degree nodes is bounded by $|L| \leq K$, as otherwise the degree sum of the graph would be greater than $K^2 = \bar{dn} = 2|E(G)|$. Thus, by transmitting serially on shortest paths (with no transmissions occurring simultaneously), the nodes in L can be informed in time $O(K \cdot OPT)$. In order to inform the small-degree nodes $V \setminus L$, we simulate the radio-broadcast algorithm of [22] on the subgraph $G[C]$, where $C = V \setminus L$. To make this work in

⁶ While the result is surely well known, we were not aware of a reference for this particular version, and thus include the algorithm and a proof in the appendix for completeness.

Algorithm 1 Broadcast in the radio-unicast model for sparse general graphs

Require: $G = (V, E)$ input graph, let $K = \sqrt{dn}$; s source node of degree at least K

- 1: Let $L \leftarrow \{v : \deg_G(v) \geq K\}$, $C = V \setminus L$, and $X = N(L) \cap C$
 - 2: Inform the nodes in L sequentially along shortest paths from s
 - 3: Let $S \leftarrow \text{GREEDY-CS}(G[C], X, C, K \cdot \text{OPT})$
 - 4: Inform all nodes in S using single hops from L
 - 5: Simulate the radio broadcast algorithm of [22] on $G[C]$ until all nodes are informed
-

the desired number of rounds, we have to ensure that for each node in C , there is an informed node within distance $O(\text{OPT})$ in $G[C]$. In the following lemma, we show that the set S found by the greedy center selection algorithm guarantees this property.

Lemma 3. *Each node in C is within distance at most $3 \cdot \text{OPT}$ from a node in S in the induced subgraph $G[C]$, i.e., $\rho(G[C], S, C) \leq 3 \cdot \text{OPT}$.*

Proof. Let Q be the set of nodes in C that are informed (directly) by nodes in L in the optimal broadcasting scheme. At most $|L|$ of them can be informed in a single round, so $|Q| \leq |L| \cdot \text{OPT} \leq K \cdot \text{OPT}$. The nodes $v \in C \setminus Q$ must then all satisfy $\text{dist}_{G[C]}(v, Q) \leq \text{OPT}$ and thus $\rho(G[C], Q, C) \leq \text{OPT}$. The center selection algorithm GREEDY-SC positions $K \cdot \text{OPT} \geq |Q|$ nodes, that by Lemma 3 yields a 3-approximation of the covering radius, giving $\rho(G[C], S, C) \leq 3 \cdot \rho(G[C], Q, C) \leq 3 \cdot \text{OPT}$. \square

The previous lemma is the main ingredient of the analysis of our main result:

Theorem 5 (*). *There is a polynomial time randomized approximation algorithm for RAS with approximation factor $\tilde{O}(\sqrt{dn})$.*

5 Interval Graphs

Let $G = (V, E)$ be an interval graph. For an interval $v \in V$, denote by $l(v)$ and $r(v)$ its left and right boundaries. For $x, y \in \mathbb{R}$, let $G[x, y]$ denote the subgraph of G induced by the intervals that are entirely contained in $[x, y]$, that is, $V(G[x, y]) = \{v \in V : l(v) \geq x \text{ and } r(v) \leq y\}$. Furthermore, denote by $\text{len}(v)$ the length of interval v . We write l_{\max} for the length of a longest interval in G . W.l.o.g., we assume that all interval boundaries are integers in $\{1, \dots, 2n\}$, and all interval boundaries are distinct (every interval graph has such a representation).

Before presenting our algorithm, we show that the clique number of an interval graph G (the size of a largest clique in G) provides a lower bound for the length of an optimal schedule. This lemma is similar to Lemmas 2 and 3 of [14].

Lemma 4 (*). $\text{OPT} \geq \omega(G)/2$.

Next, our algorithm relies on the subroutine $\text{DIAM-PATH}(G)$ that, given a connected interval graph G , returns a shortest-distance path that dominates all vertices of G .

DIAM-PATH(G). Let $u_1 \in V(G)$ be the interval with smallest left boundary, and let $u_2 \in V(G)$ be the interval with largest right boundary. Let $V_p \subseteq V(G)$ be the subset of *proper intervals*, that is, the set of intervals $v \in V(G)$ that are not contained in another interval. In other words, $v \in V_p$ if, and only if, there is no $v' \in V(G)$ with $l(v') < l(v) < r(v) < r(v')$. Since all interval boundaries are distinct, both u_1 and u_2 are proper intervals and hence in V_p . $\text{DIAM-PATH}(G)$ returns a shortest path from u_1 to u_2 in the graph $G[V_p]$. This “diameter path” has length at most $\text{Diam}(G)$.

Algorithm. Similar to our algorithm for sparse general graphs, we assume that the value of OPT is known. Furthermore, assume that the input graph G is connected. We will decompose G hierarchically as follows. Let $G_1 = G$ and let $P_1 = \text{DIAM-PATH}(G_1)$. Furthermore, for integers $i \geq 1$, let $U_i \subseteq V$ be the subset of intervals whose lengths are contained in $(\frac{1}{2})^i l_{\max}, (\frac{1}{2})^{i-1} l_{\max}$. Then, we define the subgraph $H_1 = G[V(P_1) \cup U_1]$ consisting of intervals of the largest length class plus a diameter path, where $V(P_1)$ denotes the intervals contained in path P_1 . As P_1 is a diameter path, $V(P_1)$ can be informed in $\text{Diam}(G)$ time. In Lemma 5, we will argue that the subgraph H_1 is 4-claw-free⁷, and, using this property, we will show in Lemma 6 that U_1 can be informed in $O(\text{OPT})$ rounds. Thus, overall in $O(\text{OPT})$ rounds, the nodes $V(H_1)$ are informed.

Next, given the subgraph G_i , we define inductively $G_{i+1} \subseteq G_i$ to be the subgraph induced by the set of yet uninformed intervals, that is, $G_{i+1} = G[V(G_i) \setminus V(H_i)]$. Let P_{i+1} be a collection of diameter paths of the connected components of G_{i+1} as computed by DIAM-PATH , and let $H_{i+1} = G_{i+1}[V(P_{i+1}) \cup (U_{i+1} \cap V(G_{i+1}))]$ consisting of yet uninformed intervals of length class $i+1$ and a collection of diameter paths, where $V(P_{i+1})$ denotes the intervals contained in the diameter paths P_{i+1} . Similar as before, once $V(P_{i+1})$ has been informed, by Lemma 6, we can inform $V(H_{i+1})$ in $O(\text{OPT})$ time. The key part of our argument is that $V(P_{i+1})$ can be informed by $V(P_i)$ in $O(\text{OPT})$ time, which is proved in Lemma 7. Our argument shows that given an interval $v \in V(P_i)$, there are at most $O(\text{OPT}^2)$ intervals in $V(P_{i+1})$ that intersect with v , and we prove that they can be informed in $O(\text{OPT})$ time. Thus, for every i , the nodes $V(H_i)$ can be informed in $O(\text{OPT})$ rounds.

As $l_{\max} \leq 2n$ and every interval is of length at least 1, there are $O(\log n)$ length classes. Hence, in $O(\log(n) \cdot \text{OPT})$ rounds, all nodes $V(G)$ can be informed.

Analysis. We are going to prove the following theorem:

Theorem 6. *There is a polynomial-time algorithm for RAS in interval graphs with approximation factor $O(\log n)$.*

The theorem follows from the previous description of the algorithm together with the main Lemmas, Lemma 6 and Lemma 7. In Lemma 6, we show that

⁷ A graph is 4-claw-free, if it doesn't contain the complete bipartite graph $K_{1,4}$ as an induced subgraph.

nodes $V(H_i)$ can be informed in $O(OPT)$ rounds if nodes $V(P_i)$ are informed, and in Lemma 7, we show that nodes $V(P_i)$ can be informed in $O(OPT)$ rounds if $V(P_{i-1})$ are informed.

We first state simple observations about the employed quantities in our algorithm.

Observation 7 *All intervals in subgraph G_i are of length at most $(\frac{1}{2})^{i-1}l_{max}$.*

Observation 8 *No interval in $V(H_i) \setminus V(P_i)$ contains an interval of P_i , that is, for every $v \in V(H_i) \setminus V(P_i)$ there is no $u \in V(P_i)$ such that $l(v) < l(u) < r(u) < r(v)$.*

Observation 8 follows by construction of P_i . The path P_i is constructed via algorithm DIAM-PATH which only chooses proper intervals.

Next, we show that the graphs H_i do not contain $K_{1,4}$ as an induced subgraph.

Lemma 5 (*). *For any i , the subgraph H_i is 4-claw-free.*

Last, we prove the main lemmas, Lemma 6 and Lemma 7, that show that the subtasks of our algorithm can all be performed in $O(OPT)$ rounds.

Lemma 6. *Suppose that the vertices of P_i have been informed. Then, $V(H_i)$ can be informed in $O(OPT)$ rounds.*

Proof. We color the vertices of P_i alternately with four colors, where each color is used on every fourth vertex. Since P_i is a diameter path, nodes with the same color must have disjoint neighborhoods. Processing the colors in sequence, the nodes of each color inform their neighbors in H_i in parallel. Since H_i is 4-claw-free, each node has at most $3\omega(H_i) \leq 3\omega(G) \leq 6 \cdot OPT$ neighbors, and using Lemma 4, the lemma follows. \square

Lemma 7. *Nodes D_{i+1} can be informed by nodes D_i in $O(OPT)$ rounds.*

Proof. Let $\phi_{i+1} : D_{i+1} \rightarrow D_i$ be a mapping so that $\phi_{i+1}(v) = u \Rightarrow u \in N(v)$. Next, produce a 4-coloring of D_i with color classes D_i^1, \dots, D_i^4 , as in the proof of Lemma 6. Iterate now through the color classes D_i^j . In each iteration, all nodes $u \in D_i^j$ inform the nodes $\phi_{i+1}^{-1}(u)$ simultaneously as follows: Let $C_1 \dots C_k$ denote the connected components of $G[\phi_{i+1}^{-1}(u)]$. Node u informs every OPT -th interval of every connected component C_j . If $|C_j| < OPT$ then an arbitrary interval of C_j is informed. Thus, u requires $O(k + |\phi_{i+1}^{-1}(u)|/OPT)$ rounds. In Claim 9, we will prove that $k = O(OPT)$ and $|\phi_{i+1}^{-1}(u)| = O(OPT^2)$.

Claim 9 (*) $|\phi_{i+1}^{-1}(u)| = O(OPT^2)$ and the number of components of $G[\phi_{i+1}^{-1}(u)]$ is $O(OPT)$.

Thus, the previous step requires $O(OPT)$ rounds. Next, the informed nodes of $\phi_{i+1}^{-1}(u)$ inform the uninformed nodes of $\phi_{i+1}^{-1}(u)$. Since $\phi_{i+1}^{-1}(u)$ is a collection of paths, and since for every uninformed node of $\phi_{i+1}^{-1}(u)$ there is an informed node within distance OPT , this step clearly can be done in $O(OPT)$ rounds. This completes the proof. \square

References

1. Aguayo, D., Bicket, J., Biswas, S., Judd, G., Morris, R.: Link-level measurements from an 802.11 b mesh network. *ACM SIGCOMM Computer Communication Review* 34(4), 121–132 (2004)
2. Alon, N., Bar-Noy, A., Linial, N., Peleg, D.: A lower bound for radio broadcast. *Journal of Computer and System Sciences* 43, 290–298 (1991)
3. An, M.K., Lam, N.X., Huynh, D.T., Nguyen, T.N.: Minimum data aggregation schedule in wireless sensor networks. *I. J. Comput. Appl.* 18(4), 254–262 (2011)
4. Bar-Noy, A., Guha, S., Naor, J., Schieber, B.: Message multicasting in heterogeneous networks. *SIAM J. Comput.* 30(2), 347–358 (2000)
5. Bar-Yehuda, R., Goldreich, O., Itai, A.: On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization. *J. Comput. Syst. Sci.* 45(1), 104–126 (1992)
6. Chalermsook, P., Laekhanukit, B., Nanongkai, D.: Graph products revisited: Tight approximation hardness of induced matching, poset dimension and more. In: *SODA*. pp. 1557–1576. *SIAM* (2013)
7. Chen, X., Hu, X., Zhu, J.: Minimum data aggregation time problem in wireless sensor networks. In: *Mobile Ad-hoc and Sensor Networks*, pp. 133–142. *Springer* (2005)
8. Chlamtac, I., Weinstein, O.: Distributed "wave" broadcasting in mobil multi-hop radio networks. In: *ICDCS*. pp. 82–89 (1987)
9. Elkin, M., Kortsarz, G.: A combinatorial logarithmic approximation algorithm for the directed telephone broadcast problem. *SIAM J. Comput.* 35(3), 672–689 (2005)
10. Elkin, M., Kortsarz, G.: Polylogarithmic additive inapproximability of the radio broadcast problem. *SIAM J. Discrete Math.* 19(4), 881–899 (2005)
11. Elkin, M., Kortsarz, G.: Sublogarithmic approximation for telephone multicast. *J. Comput. Syst. Sci.* 72(4), 648–659 (2006)
12. Fasolo, E., Rossi, M., Widmer, J., Zorzi, M.: In-network aggregation techniques for wireless sensor networks: A survey. *IEEE Wireless Communications* 14(2), 70–87 (2007)
13. Feige, U., Kilian, J.: Zero knowledge and the chromatic number. *J. Comput. Syst. Sci.* 57(2), 187–199 (1998)
14. Gagnon, J., Narayanan, L.: Minimum latency aggregation scheduling in wireless sensor networks. In: *Algorithms for Sensor Systems - 10th International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics, ALGOSENSORS 2014, Wroclaw, Poland, September 12, 2014, Revised Selected Papers*. pp. 152–168 (2014)
15. Guo, L., Li, Y., Cai, Z.: Minimum-latency aggregation scheduling in wireless sensor network. *Journal of Combinatorial Optimization* pp. 1–32 (2014)
16. Halldórsson, M.M., Mitra, P.: Wireless Connectivity and Capacity. In: *SODA* (2012)
17. Incel, O.D., Ghosh, A., Krishnamachari, B.: Scheduling algorithms for tree-based data collection in wireless sensor networks. In: *Theoretical aspects of distributed computing in sensor networks*, pp. 407–445. *Springer* (2011)
18. Kesselman, A., Kowalski, D.: Fast distributed algorithm for convergecast in ad hoc geometric radio networks. In: *WONS*. pp. 119–124. *IEEE* (2005)
19. Kleinberg, J., Tardos, É.: *Algorithm design*. *Pearson Education* (2006)
20. Kortsarz, G., Peleg, D.: Approximation algorithms for minimum-time broadcast. *SIAM J. Discrete Math.* 8(3), 401–427 (1995)

21. Kotz, D., Newport, C., Gray, R.S., Liu, J., Yuan, Y., Elliott, C.: Experimental evaluation of wireless simulation assumptions. In: MSWiM. pp. 78–82 (2004)
22. Kowalski, D.R., Pelc, A.: Optimal deterministic broadcasting in known topology radio networks. *Distributed Computing* 19(3), 185–195 (2007)
23. Moscibroda, T., Wattenhofer, R.: The Complexity of Connectivity in Wireless Networks. In: INFOCOM (2006)
24. Ravi, R.: Rapid rumor ramification: Approximating the minimum broadcast time (extended abstract). In: FOCS. pp. 202–213 (1994)
25. Wan, P.J., Huang, S.C.H., Wang, L., Wan, Z., Jia, X.: Minimum-latency aggregation scheduling in multihop wireless networks. In: MOBIHOC. pp. 185–194. ACM (2009)
26. Xu, X., Wang, S., Mao, X., Tang, S., Li, X.Y.: An improved approximation algorithm for data aggregation in multi-hop wireless sensor networks. In: FOWANC. pp. 47–56. ACM (2009)

A Missing Proofs

Observation 1 Let $B = (U, V, E)$ be a bipartite graph. Suppose all the vertices in U know the message. Then, the minimum number of rounds it takes to inform V in the radio-unicast model equals $\text{IMCOV}(B)$.

Proof. Consider a RAS-legal matching that contains the edges (x, a) and (y, b) , where $x, y \in U$ and $a, b \in V$. Note that it is required that $(x, b), (y, a) \notin E$ and hence the RAS-legal matching is an induced matching. Conversely, given an induced matching, all the vertices in V in the matching receive the message as there is no interference. \square

Claim 2 Let $G = (V, E)$ be a graph. Then, $B_e(G) = (V, \bar{V}, E')$ can be decomposed into $\chi(G)$ induced matchings that are pairwise disjoint and together contain all of \bar{V} , i.e., $\text{IMCOV}(B_e(G)) \leq \chi(G)$.

Proof. Let $C_1, C_2, \dots, C_{\chi(G)}$ be the color classes of G . For $i \geq 1$, define a matching M_i between C_i and their copies \bar{C}_i using the edges (v, \bar{v}) , for each $v \in C_i$. Note that since C_i is an independent set, M_i is an induced matching. By definition, the matchings $M_1, M_2, \dots, M_{\chi(G)}$ are vertex disjoint and cover all the vertices. \square

Claim 3 Let G be a graph, k an integer. Then, $\text{IM}(B_e(G^k)) \leq k \cdot n + \alpha(G)^k$.

Proof. We will use the following inequalities, which appear as Lemma 5.3 and Corollary 5.1 in [6], respectively.

$$\text{SIM}(B_e(G)) \leq \text{SIM}(B(G)) + \alpha(G), \quad (5)$$

$$\text{SIM}(B(G^k)) \leq k \cdot \text{SIM}(B(G)). \quad (6)$$

Applying (5) to $B_e(G^k)$, followed by Inequality 6 and Inequality 2 gives

$$\text{SIM}(B_e(G^k)) \leq \text{SIM}(B(G^k)) + \alpha(G^k) \leq k \cdot \text{SIM}(B(G)) + \alpha(G)^k \leq k \cdot n + \alpha(G)^k.$$

The claim then follows from the relationship $\text{IM}(G) \leq \text{SIM}(G)$ (that holds for any graph G). \square

Corollary 1 The RAS problem is hard to approximate on bipartite graphs within a factor of $(\bar{d}N)^{\frac{1}{2}-\delta}$, for any $\delta > 0$, where N is the number of vertices.

Proof. Consider the graph H_k from the proof of Theorem 4, and let $n = |V(H_k)|$. Let $m = |E(H_k)|$. Let \hat{H}_d be the graph obtained from H_k by adding a complete binary tree with $\Theta(m/d)$ vertices to the graph and connect the root of the binary tree to the source node.

Then, $N = |V(\hat{H}_d)| = n + \Theta(m/d) = \Theta(m/d)$, while the number of edges is $m + \Theta(m/d) = \Theta(m(1 + 1/d))$. The average degree of \hat{H}_d is hence $\Theta(d)$. Note that the introduced binary tree can be informed in $\Theta(\log(m/d)) = \Theta(\log(n/d))$ rounds. Since in any graph $\text{OPT} = \Omega(\log n)$, the introduced binary tree hence doesn't change the hardness of RAS and it is still hard to approximate it within a factor of $n^{1-\epsilon}$. Since $dN = \Theta(m) = O(n^2)$, the problem is also hard to approximate within $(d\hat{n})^{(1-\epsilon)/2}$. \square

Theorem 5 There is a polynomial time randomized approximation algorithm for RAS with approximation factor $\tilde{O}(\sqrt{dn})$.

Proof. Suppose that OPT is known to the algorithm. Recall from above that $|L| \leq K$. As any node can be informed in OPT time along a shortest path, the set L is informed in time $OPT \cdot K$ (Line 2). The center selection algorithm GREEDY-SC chooses $K \cdot OPT$ centers S that are adjacent to L in G . Informing those in Line 4 takes time at most $K \cdot OPT$, since each requires only a single transmission from a node in L .

Consider now the graph $G[C]$. By construction, the maximum degree in $G[C]$ is at most K . As shown in Lemma 3, the distance in $G[C]$ from an arbitrary node to an informed node (a node in S) is at most $3 \cdot OPT$. Suppose we form the graph H consisting of $G[C]$ along with a new node s' that is adjacent to all the nodes in S . By the above argument, the diameter of H is $O(OPT)$, so the radio broadcast algorithm of [22] uses $O(\text{Diam}(H) + \log^2 |V(H)|) = O(OPT + \log^2 n)$ rounds to broadcast information from s' . Running the algorithm on H when all the nodes in S have been informed will certainly not take more time. Thus, we can apply our radio broadcast simulation of Lemma 1 to obtain a RAS broadcast on $G[C]$ in time $O((OPT + \log^2 n)K) = \tilde{O}(OPT \cdot K)$. \square

Lemma 4 $OPT \geq \omega(G)/2$.

Proof. Let C be a largest clique of size $\omega(G)$ and let $x \in \mathbb{R}$ such that for every $u \in C : l(u) \leq x \leq r(u)$, that is, every interval of the clique intersects x . Suppose for the sake of a contradiction that three intervals of C are informed in the same round, that is, there are distinct informed intervals $u_1, u_2, u_3 \in V$ and distinct uninformed intervals $v_1, v_2, v_3 \in C$ such that, for $i \in \{1, 2, 3\}$, u_i informs v_i . This implies that for every i , v_i is adjacent to u_i but not to $\{u_1, u_2, u_3\} \setminus \{u_i\}$. As the intervals v_1, v_2, v_3 all intersect the point x , one of the three intervals v_i , $i \in \{1, 2, 3\}$ is so that $\min\{l(v_j) : j \neq i\} < l(v_i) < r(v_i) < \max\{l(v_j) : j \neq i\}$. Then, however, at least one interval v_j with $j \neq i$ is also adjacent to u_i , a contradiction. Thus, at most two intervals of C are informed in the same round which proves the lemma. \square

Lemma 5 For any i , the subgraph H_i is 4-claw-free.

Proof. $V(H_i)$ consists of the intervals of the diameter path P_i , and a subset of U_i . As the lengths of intervals in U_i differ at most by a factor of 2, the subgraph of H_i induced by the vertices $V(H_i) \cap U_i$ cannot induce a 4-claw. Next, by Observation 8, no interval of P_i is contained in any interval of U_i . Thus, a 4-claw in H_i could potentially only exist if an interval $v \in P_i$ had four independent neighbors in $V(H_i) \cap U_i$. This, however, implies that $\text{len}(v) \geq 2 \cdot \min\{\text{len}(u) : u \in V(H_i) \cap U_i\} + 2$, since two of the four intervals have to be fully contained in v and the other two have to overlap. The bound can be bounded from below by $2 \cdot (\frac{1}{2})^i l_{max} + 2 = (\frac{1}{2})^{i-1} l_{max} + 2$, a contradiction to Observation 7. Hence, H_i is 4-claw-free. \square

Claim 9 $|\phi_{i+1}^{-1}(u)| = O(OPT^2)$ and the number of components of $G[\phi^{-1}(u)]$ is $O(OPT)$.

Proof. Let u_1, \dots, u_q denote the intervals of $\phi_{i+1}^{-1}(u)$ ordered from left to right. Since u_3 does not intersect with u_1 and u_1 intersects with u , u_3 is entirely contained in u . By a similar argument, u_{q-2} is entirely contained in u . Hence, all intervals u_3, \dots, u_{q-2} are entirely contained in u .

Let x be the left boundary of u_3 , and let y be the right boundary of u_{q-2} . Then, $y - x \leq \text{len}(u) \leq (\frac{1}{2})^{i-1} l_{max}$, where the second inequality is due to Observation 7.

Consider now the graph $G[x, y]$. Note that as $y - x \leq \text{len}(u) \leq (\frac{1}{2})^{i-1} l_{max}$, none of the nodes of $\bigcup_{j \leq i-1} U_j$ are contained in $V(G[x, y])$. Furthermore, as for every j , P_j consists of proper intervals in G_j , none of the intervals $\bigcup_{j \leq i} P_j$ are included in u and hence in $V(G[x, y])$. Thus, the only nodes outside $V(G_{i+1})$ that could potentially be contained in $G[x, y]$ are nodes of U_i . Let $V' = G[x, y] \cap U_i$.

Let C_1, \dots, C_k be the components of $G[x, y] - V'$. Those components have to be informed by the nodes $N(V(G[x, y] - V'))$. Note that for every $w \in N(V(G[x, y] - V'))$, either $w \in V'$, w intersects x , w intersects y , or w intersects both x and y . The key of our argument is that at most four⁸ intervals of $N(V(G[x, y] - V'))$ can inform intervals of $V(G[x, y] - V')$ simultaneously in one round. To see this, observe first that it is impossible that two intervals that both intersect x (or y) simultaneously inform two intervals of $V(G[x, y] - V')$ (see left side of Figure 2). Then, since every interval of V' is of length at least $(\frac{1}{2})^i l_{max}$ and hence at least of length $\frac{1}{2}(y - x)$, at most 2 intervals of V' may inform intervals of $V(G[x, y] - V')$ simultaneously (see right side of Figure 2).




Fig. 2. Left: Illustration of the fact that two intervals a, b intersecting x cannot inform two intervals u, v of $G[x, y]$ simultaneously since either u or v is adjacent to both a and b (in the illustration, u is adjacent to a, b). Right: No three intervals a, b, c of $G[x, y]$ of sizes at least $\frac{1}{2}(y - x)$ can inform three intervals u, v, w of $G[x, y]$ simultaneously (in the illustration, v is adjacent to both a and b).

Thus, in OPT rounds, at most $4 \cdot OPT$ intervals of $V(G[x, y] - V')$ can be informed. This immediately proves the second part of the lemma, that is, the number of components of $G[\phi_{i+1}^{-1}(u)]$ is $O(OPT)$.

To prove the first part, for the sake of a contradiction, suppose that $|\phi_{i+1}^{-1}(u)| > C \cdot OPT^2$ for a large enough C . Since $G[\phi_{i+1}^{-1}(u)]$ is a collection of paths and the

⁸ By a more precise argument, three can also be argued. Any constant is enough for our purposes.

fact that at most $4 \cdot OPT$ intervals of $\phi_{i+1}^{-1}(u)$ have been informed by nodes outside $V(G[x, y] - V')$, there exists a node $v \in \phi_{i+1}^{-1}(u)$ that has not been informed by $V \setminus \phi_{i+1}^{-1}(u)$ and is at distance at least $C \cdot OPT^2 / (4 \cdot OPT) = C \cdot OPT/4$ from an informed node. As C is chosen large enough, this implies that more than OPT rounds are required to inform v , a contradiction. Hence, we have $\phi_{i+1}^{-1}(u) = O(OPT^2)$. \square

B Center Selection Algorithm

We believe that the following algorithm together with Lemma 2 are well-known. However, since we couldn't find a reference for this particular version of the center selection problem, we include the algorithm and its analysis for completeness.

Algorithm 2 Center Selection algorithm GREEDY-CS(G, X, C, k)

Require: Graph $G = (V, E)$, potential sites for centers $X \subseteq V$, clients $C \subseteq V$, number of centers to be placed k

- 1: $S \leftarrow \{\text{an arbitrary node in } C\}$
 - 2: **for** $i = 1 \dots k - 1$ **do**
 - 3: $c \leftarrow \arg \max_{c' \in C} \text{dist}(c', S)$
 - 4: $x \leftarrow \arg \min_{x' \in X} \text{dist}(x', c)$
 - 5: $S \leftarrow S \cup \{x\}, X \leftarrow X \setminus \{x\}$
 - 6: **end for**
 - 7: **return** S
-

Lemma 2 GREEDY-CS is a 3-approximation algorithm for CENTER SELECTION.

Proof. Let $r = \rho(G, S, C)$ be the covering radius of the set S as computed by GREEDY-CS. Let S^* denote an optimal solution and let $r^* = \rho(G, S^*, C)$ be its covering radius.

First, suppose that there are two centers $x_1, x_2 \in S$ with $\text{dist}(x_1, x_2) \leq \frac{2}{3}r$. W.l.o.g. suppose that x_1 was inserted into S before x_2 . Consider the iteration when x_2 was inserted and denote by c the client that was chosen in this iteration in Line 3. Since c was chosen, we have $\text{dist}(c, x_1) \geq r$. Using this fact and the assumption $\text{dist}(x_1, x_2) \leq \frac{2}{3}r$, by the triangle inequality, we obtain $\text{dist}(x_2, c) \geq \frac{1}{3}r$. Note that x_2 is the node that minimizes the distance to c , and thus we have $r^* \geq \text{dist}(x_2, c)$ which implies $r^* \geq \frac{1}{3}r$ and proves the lemma for this case.

Assume now that for every two centers $x_1, x_2 \in S$, we have $\text{dist}(x_1, x_2) \geq \frac{2}{3}r$. Let $x \in S$ be any node and denote by c the selected client when x was inserted into S . Then, $\text{dist}(x, c) \leq r^*$. As c is covered in S^* within distance r^* , there exists an $x' \in S^*$ s.t. $\text{dist}(x, x') \leq 2r^*$. Suppose that $r > 3r^*$. Under this assumption and using the fact that two centers $x_1, x_2 \in S$ are at least a distance $\frac{2}{3}r$ apart, there exists an injective mapping $\phi : S \rightarrow S^*$ so that $\text{dist}(x, \phi(x)) \leq 2r^*$. As $|S| = |S^*|$, this mapping is a bijection. This, however, implies that $r \leq 3r^*$, a

contradiction. Hence, the assumption that $r > 3r^*$ was wrong and we deduce that $r \leq 3r^*$ which proves the lemma. \square