

On the Complexity of Telephone Broadcasting from Cacti to Bounded Pathwidth Graphs

Aida Aminian ✉️🏠 

York University, Toronto, Canada

Shahin Kamali ✉️🏠 

York University, Toronto, Canada

Seyed-Mohammad Seyed-Javadi ✉️🏠 

York University, Toronto, Canada

Sumedha ✉️ 

York University, Toronto, Canada

Abstract

In TELEPHONE BROADCASTING, the goal is to disseminate a message from a given source vertex of an input graph to all other vertices in the minimum number of rounds, where at each round, an informed vertex can send the message to at most one of its uninformed neighbors. For general graphs of n vertices, the problem is NP-complete, and the best existing algorithm has an approximation factor of $\mathcal{O}(\log n / \log \log n)$. The existence of a constant factor approximation for the general graphs is still unknown.

In this paper, we study the problem in two simple families of sparse graphs, namely, cacti and graphs of bounded pathwidth. There have been several efforts to understand the complexity of the problem in cactus graphs, mostly establishing the presence of polynomial-time solutions for restricted families of cactus graphs (e.g., [5, 7, 15, 16, 19, 18]). Despite these efforts, the complexity of the problem in arbitrary cactus graphs remained open. We settle this question by establishing the NP-completeness of telephone broadcasting in cactus graphs. For that, we show the problem is NP-complete in a simple subfamily of cactus graphs, which we call snowflake graphs. These graphs are not only cacti but also have pathwidth 2. These results establish that, despite being polynomial-time solvable in trees, the problem becomes NP-complete in very simple extensions of trees.

On the positive side, we present constant-factor approximation algorithms for the studied families of graphs, namely, an algorithm with an approximation factor of 2 for cactus graphs and an approximation factor of $\mathcal{O}(1)$ for graphs of bounded pathwidth.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis; Theory of computation → Approximation algorithms analysis

Keywords and phrases Telephone Broadcasting, Approximation Algorithms, NP-Hardness, Graph Pathwidth, Cactus Graphs

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.10

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version:* [arXiv \[1\]](#)

Funding We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) [funding reference number DGECR-2018-00059].

1 Introduction

The TELEPHONE BROADCASTING problem involves disseminating a message from a single given source vertex to all other vertices in a network through a series of *telephone calls*. The network is often modeled as an undirected and unweighted graph of n vertices. Communication



© Aida Aminian, Shahin Kamali, Seyed-Mohammad Seyed-Javadi, and Sumedha; licensed under Creative Commons License CC-BY 4.0
52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025).
Editors: Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis
Article No. 10; pp. 10:1–10:17



Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



takes place in synchronous rounds. Initially, only the source is informed. During each round, any informed vertex can transmit the message to at most one of its uninformed neighbors via a “call”. The goal is to minimize the number of rounds required to inform the entire network. Hedetniemi *et al.* [20] identifies TELEPHONE BROADCASTING as a fundamental primitive in distributed computing and communication theory, forming the basis for many advanced tasks in these fields. Slater *et al.* [28] established the NP-completeness of TELEPHONE BROADCASTING. Nonetheless, efficient algorithms have been developed for specific classes of graphs. In particular, Fraigniaud and Mitjana [11] demonstrated that the problem is solvable in polynomial time for trees. There are polynomial algorithms for several other graph families; see, e.g., [6, 13, 22, 29, 10].

In this paper, we study cactus graphs, which are graphs in which any two cycles share at most one vertex. Cacti are a natural generalization of trees and ring graphs, providing a flexible model for applications such as wireless sensor networks, particularly when tree structures are too restrictive [3]. There have been several studies for broadcasting in specific families of cactus graphs [5, 7, 15, 16, 19, 18]. For instance, for unicyclic graphs, a simple subset of cactus graphs containing exactly one cycle, the problem is solvable in linear time [16]. Similarly, [18] proved that the chain of rings, which consists of cycles connected sequentially by a single vertex, has an optimal algorithm that runs in linear time. In k -restricted cactus graphs, where each vertex belongs to at most k cycles, for a fixed constant k , Čevnik and Žerovnik [5] proposed algorithms that compute the optimal broadcast scheme in $\mathcal{O}(n)$ time. Despite all these efforts, the complexity of TELEPHONE BROADCASTING for arbitrary cactus graphs has remained open, a question that we resolve in this paper.

TELEPHONE BROADCASTING is NP-hard for general graphs [12], and even for restricted graphs [23]. In particular, Tale recently showed that the problem remains NP-hard for graphs of pathwidth of 3 [30], a result that naturally extends to graphs with higher pathwidths. However, the complexity of the problem for graphs of pathwidth 2 has remained an open question, which we address in this paper.

Elkin and Kortsarz [8] proved that approximating TELEPHONE BROADCASTING within a factor of $3 - \epsilon$ for any $\epsilon > 0$ is NP-hard. Kortsarz and Peleg [21] showed that TELEPHONE BROADCASTING in general graphs has an approximation ratio of $\mathcal{O}(\log n / \log \log n)$ (improving upon [25, 2]). It is possible that there is a constant factor approximation for general graphs. However, constant-factor approximation exists only for restricted graph classes such as unit disk graphs [27] and certain sub-families of cactus graphs such as k -cycle graphs [4].

1.1 Contribution

This paper investigates the complexity and approximation algorithms for the TELEPHONE BROADCASTING problem for cacti and graphs of bounded pathwidth. The key contributions are summarized as follows:

- We present a simple polynomial-time algorithm, named *Cactus Broadcaster*, and prove it has an approximation factor of 2 for broadcasting in cactus graphs (Theorem 6). Constant-factor approximations are known for certain subfamilies of cactus graphs (e.g., k -cycle graphs [4]), and our result extends this to arbitrary cactus graphs. *Cactus Broadcaster* is reminiscent of the tree algorithm of [11] and leverages the separability of cactus graphs. In our analysis of *Cactus Broadcaster*, we use the k -broadcasting model of broadcasting [14, 20] as a reference point, where each vertex can inform up to two neighbors in a single round.

- Our main contribution is to establish the NP-completeness of TELEPHONE BROADCASTING problem in “snowflake graphs”, which are a subclass of cactus graphs and also have pathwidth at most 2, therefore resolving the complexity of the problem in these graph families (Theorem 18). For a formal definition of snowflake graphs, refer to Definition 3. This hardness result is achieved through a series of reductions that start from 3,4-SAT, a variant of the satisfiability problem that is NP-complete [31].
- We show the existence of a constant-factor approximation for graphs of bounded pathwidth. For that, we show the algorithm of Elkin and Kortsarz [9] has an approximation factor of $\mathcal{O}(4^w)$ for any graph of constant pathwidth w (Theorem 20). Note that this result does not rely on having a path decomposition certifying a bounded pathwidth. Constant-factor approximation algorithms are known for certain families of graphs of bounded path-width such as k -path graphs, which admit a 2-approximation algorithm [17], and our result extends this to any graph of bounded pathwidth.

1.2 Paper Structure

In Section 2, we present the preliminaries. In Section 3, we propose *Cactus Broadcaster*, which gives a 2-approximation for cactus graphs. In Section 4, we establish the NP-completeness of the problem in snowflake graphs. In Section 5, we present a constant-factor approximation algorithm for graphs of bounded pathwidth. We conclude in Section 6. All missing proofs are provided in the full version of the paper [1].

2 Preliminaries

For a positive integer n , we use notation $[n]$ to denote $\{1, 2, \dots, n\}$. Also, we use $[i, j]$ to refer to denote $\{i, i+1, \dots, j\}$. For a graph G , we use $G \setminus \{v\}$ to refer to the subgraph of G induced by all vertices of G except v .


► **Definition 1** ([20]). *An instance (G, s) of the TELEPHONE BROADCASTING problem is defined by a connected, undirected, and unweighted graph $G = (V, E)$ and a vertex $s \in V$, where s is the only informed vertex. The broadcasting protocol is synchronous and occurs in discrete rounds. In each round, an informed vertex can inform at most one of its uninformed neighbors. The goal is to broadcast the message as quickly as possible so that all vertices in V get informed in the minimum number of rounds.*

A broadcast scheme describes the order in which each vertex informs its neighbors. One can describe a broadcast scheme S with a *broadcast tree*, which is a spanning tree of G rooted at source s ; if a vertex u is informed through vertex v , then u will be a child of v in the broadcast tree. Given that the optimal broadcast scheme of trees can be computed in linear time, a broadcast tree can fully describe the broadcast scheme. We use $\mathbf{br}^*(G, s)$ to refer to the number of rounds in the optimal broadcast scheme.

► **Definition 2.** *Cactus graphs are connected graphs in which any two simple cycles have at most one vertex in common.*

► **Definition 3.** *A tree T is said to be a reduced caterpillar if there are three special nodes x, y , and z in T such that every node in T is either located on the path between x and y or is connected to z .*

A graph G is said to be a snowflake, if and only if it has a center vertex c such that $G \setminus \{c\}$ is a set of disjoint reduced caterpillars such that c is connected to exactly two vertices in any of these caterpillars, none being special vertices.



■ **Figure 1** An illustration of reduced caterpillar and snowflake graphs.

An example of snowflake graphs and one of its corresponding reduced caterpillar components is shown in Figure 1. Informally, a snowflake graph is formed by a set of cycles that have a common center c ; moreover, each cycle has a special vertex (z vertices). Any vertex in G is either i) a part of one of the cycles or ii) a part of a “dangling path” connected to a neighbor of c or iii) finds a special vertex as its sole neighbor.

► **Definition 4** ([26]). A path decomposition D of a given graph $G = (V, E)$ is a sequence $\langle B_1, B_2, \dots, B_k \rangle$, where each B_i is called a bag and contains a subset of V , such that every vertex $v \in V$ appears in at least one bag and, for every edge $(u, v) \in E$, there exists a bag B_i containing both u and v . Furthermore, if a vertex v appears in B_i and in B_j , it must appear in any B_k where $k \in [i, j]$. The width of the path decomposition D is the maximum cardinality of its bags minus 1. Now, G is said to have pathwidth w if it has a path decomposition of width at most w .

► **Observation 5.** Snowflake graphs have a pathwidth of at most 2.

Proof. Removing the center from a snowflake graph results in a collection of disjoint caterpillar graphs, each having a pathwidth of 1 (caterpillars are graphs of pathwidth 1 [24]). A valid path decomposition for a snowflake graph can be achieved by adding the center to all bags of path decompositions of these caterpillars. ◀

3 Cactus Broadcaster: A 2-Approximation for Cactus Graphs

In this section, we present our 2-approximation algorithm for cactus graphs. We use ideas from k -broadcasting model with parameter k [14], where a vertex can inform up to k of its neighbors in a single round via a *super call*. It is easy to see that if one can complete k -broadcasting in m rounds, then it is possible to complete broadcasting (in the classic setting) within km rounds. This can be achieved by “simulating” a super-call with up to k regular calls [1]. For our algorithm, we use k -broadcasting with $k = 2$. In particular, we design an algorithm for 2-broadcasting in a cactus graph G and show that if it completes within m rounds, then any broadcast scheme for classic broadcasting in G takes at least m rounds [1]. Therefore, if we simulate every super call with two regular calls (in arbitrary order), the broadcasting completes within $2m$ rounds, and thus, we achieve an approximation factor of 2. In what follows, we describe the algorithm in detail.

3.1 Cactus Broadcaster Algorithm

This section presents *Cactus Broadcaster*, a 2-broadcasting algorithm that yields our 2-approximation algorithm for the TELEPHONE BROADCASTING problem in cactus graphs. *Cactus Broadcaster* is a mutually recursive algorithm with two main methods, namely, **single-br**(G, s) and **double-br**(G, s_1, s_2). The **single-br**(G, s) method is used for 2-broadcasting a connected subgraph G starting from a source vertex s . The **double-br**(G, s_1, s_2) method is used for 2-broadcasting in a connected subgraph G using two sources s_1 and s_2 under an assumption that both s_1 and s_2 have the message at time 0 and there is a unique path between them in G .

Broadcasting from a single source. We explain how **single-br**(G, s) operates. Observe that removing the source vertex s partitions G into one or more disjoint connected components. By the definition of cactus graphs, each of these connected components contains at most two neighbors of s (see [1] for details). We refer to a connected component with one neighbor (respectively two neighbors) of s as a *single-neighbor* (respectively *double-neighbor*) component.


single-br(G, s) computes the broadcast time of each single-neighbor component C recursively by computing **single-br**(C, u), where u is the single neighbor of s in C . Similarly, it computes the broadcast time of each double-neighbor component C by computing **double-br**(C, u, v), where u, v are neighbors of s in C . After calculating all these broadcasting times, **single-br**(G, s) sorts these values in non-increasing order and informs the neighbors of s in this order. For that, it uses a regular call for single-neighbor components and a super-call for double-neighbor components. In other words, **single-br**(G, s), informs the component with the largest broadcasting time in the first round, the component with the second largest broadcast time next, and so on. This is reminiscent of broadcasting in trees [11], except that super-calls are used for double components. Let b_i denote the broadcast time of the i 'th component C_i in this order.

The total broadcasting time for **single-br**(G, s) is then computed as $\max_i(i + b_i)$, which will be the output of **single-br**(G, s).

Broadcasting from two sources. Next, we describe **double-br**(G, s_1, s_2). Let P be the unique path between s_1 and s_2 . In any broadcast tree T of **double-br**, some vertices on P are informed through s_1 and some through s_2 . Thus, exactly one edge $e \in P$ is excluded from the tree. After removing such e , the graph G will be partitioned into two disjoint connected components G_1^e and G_2^e , which are respectively informed through s_1 and s_2 (See Figure 2). The method **double-br** works by removing the edge e^* with a minimum value of $\max\{\mathbf{single-br}(G_1^e, s_1), \mathbf{single-br}(G_2^e, s_2)\}$ over all $e \in P$.

Next, we explain how **double-br** finds e^* . An exhaustive approach that tries all edges $e \in P$ and calls **single-br** twice per edge may take exponential time. To fix this, **double-br** works in two phases:

- In Phase 1, the method pre-computes the broadcast time for the connected subgraphs of G after removing vertices of P from G . Every resulting connected component after removing P has at most two neighbors in P [1]. For any single-neighbor component C with vertex s connected to P , **single-br**(C, s) is computed recursively. Similarly, for any double-neighbor component with vertices s_1, s_2 connected to P , **double-br**(C, s_1, s_2) is computed recursively.
- Next, we describe Phase 2. At each iteration of Phase 2, we fix an edge $e = (u, v) \in P$ and aim to compute **single-br**(G_1^e, s_1) and **single-br**(G_2^e, s_2), where $u \in G_1^e$ and $v \in G_2^e$. We explain how to efficiently compute **single-br**(G_1^e, s_1). Finding **single-br**(G_2^e, s_2)



■ **Figure 2** An illustration of Phase 2 of the **double-br** method. Different components in iteration $i = 3$ of Phase 2 (for u_3) are highlighted in different colors.

is done similarly. Let $\langle u(= u_1), u_2, \dots, u_k(= s_1) \rangle$ be the sequence of vertices in the unique path from u to s_1 . Let H_i^e be the connected component of G that contains u_i after removing the edge (u_i, u_{i+1}) from G_1^e . Note that $H_k^e = G_1^e$. We compute $b_i = \text{single-br}(H_i^e, u_i)$ using the values computed in Phase 1 as follows in an iterative manner from $i = 1$ to $i = k$.

Consider all connected components of H_i^e after removing u_i . Note that, for $i > 1$, H_{i-1}^e is one of these components. Consider the set B of broadcast times for all these components, with the neighbor(s) of u_i as the source(s) of broadcast. All these values are computed in Phase 1 except for b_{i-1} , which is computed in the previous iteration. **double-br** sorts B in non-increasing order of broadcast times and informs neighbors of u_i accordingly (similar to **single-br**). As a result, in the i 'th iteration, the broadcast time b_i is computed as $\max_{j \in [B]} (j + B[j])$. At iteration $i = k$, we compute b_k , which is indeed $\text{single-br}(G_1^e, s_1)$ [1].

Figure 2 provides an illustration.

► **Theorem 6.** *Cactus Broadcaster runs a polynomial-time for TELEPHONE BROADCASTING and achieves an approximation factor of 2 on cactus graphs.*

Sketch of the proof. First, we show that *Cactus Broadcaster* runs in $\mathcal{O}(n^3 \log n)$. This can be established using induction on the size of the input graph, based on the recursive nature of the algorithm. Details can be found in [1].

Second, we show that *Cactus Broadcaster* has an approximation factor of 2. Intuitively, **single-br** with super-calls runs no longer than an optimal broadcast scheme without super-calls; this is because every super call only expedites informing one of the two sources of broadcasting in double-neighbor components. As mentioned earlier, a broadcast scheme with super calls can be simulated with a scheme (without super-calls) that completes no later than twice the number of rounds. On the other hand, the selection of edge e^* by **double-br** ensures that $\max(\text{single-br}(G_1^{e^*}, s_1), \text{single-br}(G_2^{e^*}, s_2)) \leq \max(\text{single-br}(G_1^{e^+}, s_1), \text{single-br}(G_2^{e^+}, s_2))$, where e^+ is the edge that is absent in the optimal broadcast tree of G . A formal proof follows from an inductive argument on the input size [1]. ◀

The approximation factor given by Theorem 6 is tight. Consider a graph G formed by t triangles that all share an endpoint s . The broadcasting time of *Cactus Broadcaster* on (G, s) is $2t$ as its broadcast tree will be a star with $2t$ leaves, while the optimal broadcast tree completes broadcasting in $t + 1$ rounds (s will have degree t). Thus, the approximation factor tends to 2 for large t .

4 Hardness Proof of Snowflake Graphs

The NP-hardness proof of TELEPHONE BROADCASTING in snowflake graphs proceeds through successive reductions. First, we reduce 3,4-SAT, which is known to be NP-hard [31] to a new problem that we call TWIN INTERVAL SELECTION (Lemma 9), which itself reduces to another new problem DOME SELECTION WITH PREFIX RESTRICTIONS (Lemma 15). Finally, we present a reduction from DOSEPR to TELEPHONE BROADGUESS, which establishes our main result (Theorem 18).

4.1 Hardness of Twin Interval Selection

The first step in reducing 3,4-SAT to TELEPHONE BROADGUESS is to establish a reduction from 3,4-SAT to TWIN INTERVAL SELECTION (TWIS). For that, we first explain how to construct a TWIS instance from a given 3,4-SAT instance and demonstrating the equivalence of solutions between the two problems.

► **Definition 7** ([31]). *The 3,4-SAT problem is a special type of the classic 3-satisfiability (3-SAT) problem. The input is a boolean CNF formula $\phi = (X, C)$, where X is the set of variables and C is the set of clauses such that every clause $C_i \in C$ contains exactly three literals $\ell_k^i \in C_i$ for $k \in [3]$, and each variable $x \in X$ appears in at most four clauses. The decision problem asks whether there is a satisfying assignment of ϕ .*


The 3,4-SAT problem is known to be NP-complete [31]. In what follows, we call a pair of intervals $(I_i, \bar{I}_i)$ for $i \in [n]$ a *twin interval*, assuming that I_i and $\bar{I}_i$ are non-crossing and have equal lengths with endpoints in $[m]$ for some positive m . The endpoints of I_i are less than the endpoints of $\bar{I}_i$. We refer to I_i and $\bar{I}_i$ the *left* and the *right* interval of the twin, respectively.

► **Definition 8.** *An instance of the TWIN INTERVAL SELECTION (TWIS) is formed by a tuple (I, r, m) , where I is a set of twin interval pairs and r is a restriction function with domain $[m]$, where m is referred to the horizon of the instance. We have $I = \{(I_1, \bar{I}_1), \dots, (I_n, \bar{I}_n)\}$, where for each $i \in [n]$, $(I_i, \bar{I}_i)$ are non-crossing intervals with endpoints in $[m]$.*

The objective is to select exactly one of I_i and $\bar{I}_i$, while respecting the restriction imposed by the restriction function as follows. The restriction function $r : [m] \rightarrow [n]$ requires that for any $t \in [m]$, the number of selected intervals that contain t be at most $r(t)$. The decision problem asks whether there is a valid selection that satisfies these restrictions.

For a solution S and $t \in [m]$, let $\Gamma_S(t)$ be the number of selected intervals in S containing t .

Construction. Suppose we are given an instance $\phi = (X, C)$ of 3,4-SAT. We construct an instance $\mathcal{I}_\phi$ of TWIS as follows. For each variable $x_i \in X$, form a pair of $(X_i, \bar{X}_i)$ of twin intervals, each of length 7, where $X_i = [16i - 15, 16i - 8]$ and $\bar{X}_i = [16i - 7, 16i]$. We refer to X_i (respectively $\bar{X}_i$) as the *principal* interval of literal x_i (respectively $\neg x_i$). The length 7 of these intervals allows for placing up to 4 non-crossing unit intervals, each of length 1, within the principal interval; these intervals will be used for clause gadgets, as we will explain. For each clause C_j , we define three interval twin pairs $(I_k^j, \bar{I}_k^j)$, each associated with one of the literals of $\ell_k^j \in C_j$ for each $k \in [3]$. Suppose ℓ_k^j is the p 'th literal where its corresponding variable $x_i \in X$ appears ($p \in [4]$). If ℓ_k^j is a positive (respectively negative) literal of variable x_i , then define I_k^j as a unit interval starting at $16i - 7 + 2(p - 1)$ (respectively, starting at $16i - 15 + 2(p - 1)$). Intuitively, I_k^j is defined as one of the four-unit intervals within the



■ **Figure 3** Construction of the instance $\mathcal{I}_\phi$ of the TWIS problem from the 3,4-SAT instance $\phi = (x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2 \vee \neg x_3) \wedge (\neg x_1 \vee \neg x_2 \vee x_3)$. The restriction function $r(t)$ is defined as $r(t) = 1$ for $t \leq 48$ and $r(t) = 2$ otherwise. A satisfying assignment for ϕ is $\sigma(x_1) = 1$ and $\sigma(x_2) = \sigma(x_3) = 0$; the selected intervals in the corresponding solution for $\mathcal{I}$ are shown in solid color.

principal interval of $\neg x$. Meanwhile, $\overline{I_k^j}$ is $[16n + 2j, 16n + 2j + 1]$ for all $k \in [3]$. Finally, to define the restriction function, we let $r(t) = 1$ for all $t \leq 16n$ and $r(t) = 2$, otherwise. Figure 3 illustrates an example of the above reduction.

Correctness. We provide an intuitive explanation of why this reduction works. The formal proof can be found in [1]. We argue that a satisfying assignment to instance ϕ of the 3,4-SAT problem bijects to a satisfying interval selection for the instance $\mathcal{I}_\phi$ of TWIS. The bijection requires that for any literal ℓ that is true in a satisfying assignment of ϕ , the principal interval of ℓ is selected in the solution for $\mathcal{I}_\phi$.

The restrictions imposed by r imply that whenever a principal interval of a literal ℓ is selected, the unit intervals associated with (up to 4) occurrences of $\neg \ell$ cannot be selected (because $r(t) = 1$ in their intersecting points). Thus, whenever the principal interval of a literal ℓ is selected (i.e., its twin, which is the principal interval of $\neg \ell$ is not selected), one can select the unit intervals associated with any occurrence of ℓ in clauses where it appears; this is because they intersect with the principal interval of $\neg \ell$, which is not selected. Selecting these unit intervals means their twin unit intervals. This yields an equivalent between a satisfying assignment in the SAT formula and the TWIS. We further note that the number of twin intervals in $\mathcal{I}_\phi$ is polynomial in $|X|$ and $|C|$. We can conclude the following lemma, which establishes the hardness of TWIS in the strong sense based on the above reduction.

► **Lemma 9.** *Answering instances (I, r, m) of the TWIS problem is NP-hard even if the m is polynomial in $|I|$.*

4.2 Hardness of Dome Selection with Prefix Restrictions

We present a polynomial-time reduction from TWIS to a new problem that we refer to as DOME SELECTION WITH PREFIX RESTRICTIONS (DoSEPR). Before defining the DoSEPR problem, we first define the notion of a *dome* as follows.

► **Definition 10.** *A dome is formed by four positive integers (a, b, c, d) such that $a \leq b < c \leq d$ and $b - a = d - c$. We refer to (a, d) (respectively (b, c)) as an arc with endpoints a and d (respectively b and c). We refer to (a, d) and (b, c) as the outer arc and inner arc of the dome, respectively. When $a = b$ and $c = d$, we call the dome a singleton dome and otherwise call it a regular dome.*

In our figures, a singleton dome (a, b) is shown as $\textcircled{a} \textcircled{b}$ and regular dome (a, b, c, d) is shown as $\textcircled{a} \textcircled{b} \textcircled{c} \textcircled{d}$.

► **Definition 11.** An instance (D, m) of the DOME SELECTION WITH PREFIX RESTRICTIONS (DoSEPR) is defined by a multiset $D = \{D_1, \dots, D_n\}$ of domes and a positive integer m , where $D_i = (a_i, b_i, c_i, d_i)$ s.t. $d_i \leq m$ for some integer m that we refer to as horizon. The decision problem asks whether there is a multiset S of arcs with exactly one arc from each dome D_i , such that, for any $t \in [m]$, it holds that $\mathcal{N}_S(t) \leq t$, where $\mathcal{N}_S(t)$ denotes the number of arc endpoints in S with value at most t , counting each endpoint as many times as it appears across different arcs in S .

► **Example 12.** Consider an instance of the DoSEPR with domes $D = \{(2, 3, 4, 5), (3, 4), (4, 5, 9, 10), (7, 8)\}$ and $m = 10$. A valid solution for this instance is $S = \{(2, 5), (3, 4), (5, 9), (7, 8)\}$. For example, when $t = 5$, we have $\mathcal{N}_S(t) = 5$, which is no more than t . In particular, selected endpoints $\leq t$ are $\{2, 5, 3, 4, 5\}$. Note that there are two endpoints with the value of 5, which belong to two different arcs.

Next, we explain how the TWIS problem reduces to the DoSEPR problem.

Construction. Given an arbitrary instance $\mathcal{I} = (I, r, m')$ of TWIS, where $|I| = n'$ for some positive n' , we construct the corresponding instance $\mathcal{D} = (D, m)$ of DoSEPR as follows. For each twin pair of intervals $(I_i, \bar{I}_i) \in I$, where $I_i = (a', b')$ and $\bar{I}_i = (c', d')$, we construct a regular dome $D_i = (a_i, b_i, c_i, d_i)$ where $a_i = 6n'a'$, $b_i = 6n'b' + 3n'$, $c_i = 6n'c'$, and $d_i = 6n'd' + 3n'$. It is easy to verify that D_i is indeed a dome, that is $b_i - a_i = d_i - c_i$ [1].

We define m to be a large enough horizon. In particular, we let $m = 164(n'm')^2$. For any $t \in [m]$ that is a multiple of $6n'$ and $t \leq 3n'(2m' + 1)$, we further add $d(t)$ some singleton domes that all start at t and end at m . In other words, we add $d(t)$ identical singleton domes (t, m) . Next, we explain how $d(t)$ is defined. For convenience, we let $d(t) = 0$ if t is not a multiple of $6n'$ or $t > 3n'(2m' + 1)$. Suppose t is indeed a multiple of $6n'$. The idea is to define $d(t)$ in a way to project the requirements imposed by the restriction function r in $\mathcal{I}$ to the requirement $\mathcal{N}_S(t) \leq t$ in DoSEPR.

Let $\mathfrak{D}_i^t$ be the set of regular domes with exactly i endpoints before or at point t , and define $c(t) = 2|\mathfrak{D}_4^t| + |\mathfrak{D}_3^t| + |\mathfrak{D}_2^t|$.

► **Example 13.** For $t = 18$, dome $(1, 4, 5, 8) \in \mathfrak{D}_4^t$, dome $(11, 14, 16, 19) \in \mathfrak{D}_3^t$, dome $(9, 10, 26, 27) \in \mathfrak{D}_2^t$, dome $(17, 20, 21, 24) \in \mathfrak{D}_1^t$, and dome $(25, 26, 27, 28) \in \mathfrak{D}_0^t$.

Intuitively, this definition of c implies that there are $c(t)$ arc endpoints with value at most t in any valid solution S (a set of arcs), regardless of the choices made to form S . This is because:


- (i) all arc endpoints of domes in $\mathfrak{D}_4^t$ are at most t , and two of them (associated with one arc) contribute to $c(t)$.
- (ii) three arc endpoints of domes in $\mathfrak{D}_3^t$ are at most t , and any such dome contributes at least 1 to $c(t)$.
- (iii) the left endpoints of both arcs of domes in $\mathfrak{D}_2^t$, and since one arc is in S , the dome contributes exactly 1 to $c(t)$.

Finally, we let

$$d(t) = t - c(t) - \sum_{j < t} d(j) - r(t/(6n')).$$

The scaling argument that is applied when forming the DoSEPR instance from TWIS instance ensures that $d(t)$ is indeed non-negative for any $t \in [m]$ (see [1] for details). This completes our construction of the DoSEPR instance. In a nutshell, we have added one regular dome per twin interval in the TWIS instance, and for any $t \in [m]$, we added extra identical singleton domes to capture the requirements imposed by the restriction function in the TWIS instance.

10:10 On the Complexity of Telephone Broadcasting



■ **Figure 4** An illustration of the construction of an instance $\mathcal{D}$ of DoSEPR from an instance $\mathcal{I}$ of TwIS, as explained in Example 14. Solid lines show the selected intervals in a valid solution of $\mathcal{I}$ and the selected arcs in the corresponding solution in $\mathcal{D}$.

► **Example 14.** Consider an instance $\mathcal{I} = (I, r, m')$ of the TwIS problem with $I = \{((1, 3)(4, 6)), ((2, 3), (5, 6)), ((2, 3), (4, 5))\}$ and $m' = 6$. Suppose $r(1) = r(2) = r(6) = 3$, $r(3) = r(4) = 1$, and $r(5) = 2$. The corresponding instance $\mathcal{D} = (D, m)$ of DoSEPR has regular domes $D = \{(18, 63, 72, 117), (36, 63, 90, 117), (36, 63, 72, 99)\}$ and $m = 164(n'm')^2 = 164 \cdot (3 \cdot 6)^2 = 53,136$ (see Figure 4). In addition, $3n'(2m' + 1) = 117$ for any $t \in [1, 117]$ that is a multiple of $6n' = 18$, $d(t)$ singleton domes (t, m) are added. Some examples of $d(t)$ are shown as follows.

$$\begin{array}{ll}
 d(18) = 18 - 0 - 0 - 3 = 15 & c(18) = 0 \\
 d(36) = 36 - 0 - 15 - 3 = 18 & c(36) = 0 \\
 d(54) = 54 - 0 - (18 + 15) - 1 = 20 & c(54) = 0 \\
 d(72) = 72 - 3 - (15 + 18 + 20) - 1 = 15 & c(72) = 3
 \end{array}$$

Correctness. We provide some intuitions about the correctness of the reduction. A formal proof can be found in [1]. We show that any valid solution S for instance $\mathcal{D}$ of DoSEPR bijects to a valid solution S' of $\mathcal{I}$. Note that S contains exactly one arc, the inner or the outer arc, of each regular dome D_i of $\mathcal{D}$ (in addition to singleton domes, which are contained in any valid solution for $\mathcal{D}$). Now, S' selects the left interval I_i when the outer arc of D_i is selected in S and the right interval $\bar{I}_i$ when the inner arc of D_i is selected in S (see Figure 4). We let $\Gamma_{S'}(t')$ denote the number of intervals selected by S' that include $t' \in [m']$. We show that S is a valid solution for $\mathcal{D}$ if and only if S' is a valid solution to $\mathcal{I}$.

Suppose S is a valid solution for DoSEPR. Then we must have $\mathcal{N}_S(t) \leq t$ for any $t \in [m]$, where $\mathcal{N}_S(t)$ is the number of arc endpoints in S with value at most t . As mentioned earlier, $c(t)$ endpoints will be present in any valid solution regardless of the choices to form S . Moreover, S must contain all endpoints of singleton domes in $\mathcal{D}$ for $j \leq t$. That is, any valid solution for $\mathcal{D}$, including S , must contain $c(t) + \sum_{j \leq t} d(j)$ arc endpoints. In addition to these fixed points, there will be more arcs in S , which are contributed by the domes in $\mathfrak{D}_1^t$ and $\mathfrak{D}_3^t$ as follows. Let D_i be a regular dome in $\mathcal{D}$, and let $(I_i, \bar{I}_i)$ be its corresponding twin intervals in $\mathcal{I}$.

- Suppose $D_i \in \mathfrak{D}_1^t$. Now, if S contains the inner arc of D_i , the contribution of D_i to $\mathcal{N}_S(t)$ would be 0. This is equivalent to including the right interval $\bar{I}_i$ in S' , and thus the twin interval $(I_i, \bar{I}_i)$ does not contribute $\Gamma_{S'}(t')$ for $t' = t/6n$. Informally, the contribution of the dome D_i to the left-hand side of inequality $\mathcal{N}_S(t) \leq t$ and the contribution of $(I_i, \bar{I}_i)$ to the left-hand side of the inequality $\Gamma_{S'}(t') \leq r(t')$ will be both 0. Similarly, if S contains the outer arc of D_i , the contribution of D_i to $\mathcal{N}_S(t)$ would be 1. This is equivalent to including the left interval I_i in the TwIS instance, and $(I_i, \bar{I}_i)$ contribute 1 item to $\Gamma_{S'}(t')$. Informally, the contribution of the dome D_i to the left-hand side of inequality $\mathcal{N}_S(t) \leq t$ and the contribution of $(I_i, \bar{I}_i)$ to the left-hand side of $\Gamma_{S'}(t/(6n')) \leq r(t')$ will be both 1.




Figure 5 An illustration of the construction of the TELEPHONE BROADCASTING instance corresponding to a DOSEPR instance.

- Suppose $D_i \in \mathfrak{D}_3^t$. Assume S contains the inner (respectively the outer) arc of D_i . In this case, in addition to the fixed 1 unit of the contribution of D_i to $\mathcal{N}_S(t)$ (captured in $c(t)$), it further contributes 1 (respectively 0) unit to $\mathcal{N}_S(t)$. This is equivalent to including the right interval $\bar{I}_i$ (respectively the left interval I_i) in S' . In this case, the number of selected intervals in S' that intersect t is increased by 1 (respectively 0). Intuitively, both left-hand sides of $\mathcal{N}_S(t) \leq t$ and $\Gamma_{S'}(t/(6n')) \leq r(t')$ increase by 1 (respectively 0).

We note that $n = |D|$ is polynomial in both n' and m' , and thus $m = 164(n'm')^2$ is polynomial n . Therefore, we can conclude the following lemma, which establishes the NP-hardness of DOSEPR.

► **Lemma 15.** *Answering instances (D, m) of the DOSEPR problem is NP-hard even if m is polynomial in $|D|$.*

4.3 Hardness of Telephone Broadguess in Snowflake Graphs

We refer to the decision variant of the broadcasting problem as the TELEPHONE BROADCASTING problem with instances (G, s, ρ) , where G is the input graph, s is the source, and ρ is a positive integer. The decision question asks whether it is possible to complete broadcasting in G from s within ρ rounds. This section presents our final reduction, from DOSEPR to TELEPHONE BROADCASTING in snowflake graphs.

Construction. Given an instance $\mathcal{D} = (D, m)$ of DOSEPR we construct an instance $\mathcal{T} = (G, r, \rho = 2m)$ of TELEPHONE BROADCASTING, where G is a snowflake graph with center r . For a given $x \in [\rho]$, let $\tilde{x} = \rho - x$. We construct the graph G by creating a reduced caterpillar (see Definition 3) for each dome in $\mathcal{D}$ as follows.

- The reduced caterpillar of a singleton dome (a, b) is a path of length $2\tilde{a} + 2\tilde{b}$ with endpoints p and q (here, p, q, z are special vertices in the reduced caterpillar, where z is any arbitrary vertex).
- The reduced caterpillar of a regular dome (a, b, c, d) is formed by a path of length $2\tilde{b} + 2\tilde{d} + 1$ with endpoints p and q , which are two special vertices of the reduced caterpillar. The other special vertex z of the reduced caterpillar is the vertex at distance $2\tilde{b}$ of p (and thus distance $2\tilde{d}$ of q). In addition to the vertices on the path between p and q , the reduced caterpillar includes $\tilde{a} - \tilde{b}$ (which equals $\tilde{c} - \tilde{d}$) other vertices which find z as their sole neighbor.

10:12 On the Complexity of Telephone Broadcasting

For any reduced caterpillar C constructed above (from either a singleton or a regular dome), we label the vertex at distance $\tilde{a}$ from p as the *first gate* of C , denoted as u , and the one at distance $\tilde{b}$ from q as the *second gate* of C and denote it with v . We form the equivalent instance of the TELEPHONE BROADCAST problem by forming a graph G with a center vertex r (source) that is connected to all gates of all reduced caterpillars formed by the domes of D . By Definition 3, G will be a snowflake graph with center r . Figure 5 provides an illustration of this reduction.

Correctness. Before proving the correctness of our reduction, we prove the following lemmas for broadcasting in the reduced caterpillars associated with singleton and regular domes, respectively.

► **Lemma 16.** *Consider a reduced caterpillar C of a singleton dome $D = (a, b)$ with gates u and v . One can complete broadcasting in C within ρ rounds if and only if u is informed no later than time a and v is informed no later than time b .*

► **Lemma 17.** *Consider a reduced caterpillar C of a regular dome $D = (a, b, c, d)$ with gates u and v . One can complete broadcasting in C within ρ rounds if and only if one of the following happens:*

- (i) u is informed by time b and v is informed by time c .
- (ii) u is informed by time a and v is informed by time d .


Proof. First, we show that if either (i) or (ii) hold, then broadcasting completes within ρ rounds.

- Suppose (i) holds, that is, u and v are informed by times b and c , respectively. We describe a broadcast scheme S that completes within ρ rounds. See Figure 6a for an illustration.

In S , vertex u first informs the neighbor on its path towards p , and then it informs the neighbors on its path towards z . This means that, by round ρ , all $\tilde{b}$ vertices on the path between u and p will be informed. Similarly, $\tilde{b} - 1$ vertices on the path between u and z are informed (this excludes z). On the other hand, v first informs the neighbor on its path towards z and then the neighbor on its path towards q . This means that, by round ρ , all $\tilde{d}$ vertices on the path between v and q will be informed; this is because the number of rounds between the time that v is informed and ρ is $\tilde{c}$, and since v first informs the neighbor towards z , $\tilde{c} - 1$ rounds remain for informing vertices between v and q . This number of rounds is sufficient because the length of the path between v and q is $\tilde{d} \leq \tilde{c} - 1$. This way, z will receive the message at time $c + \tilde{d}$, and assuming z informs its uninformed leaves in arbitrary order, the broadcasting in S completes within $c + \tilde{d} + \tilde{c} - \tilde{d} = \rho$ rounds.

- Next, assume (ii) holds, that is, u and v are respectively informed by rounds a and d . See Figure 6b for an illustration. In this case, u first informs its neighbor on the path towards z and then its other neighbor on the path towards p . This means that by round $a + \tilde{b} + 1 \leq a + \tilde{a} = \rho$, all vertices on the path between a and p will be informed. Similarly, by round a , all the $\tilde{b}$ vertices on the path between u and z (including z) will be informed. Thus, all neighbors of z will be informed by $a + \tilde{b} + \tilde{a} - \tilde{b} = \rho$.

On the other hand, v first informs its neighbor on the path toward q and then its neighbor on the path towards z (excluding z). Thus, all vertices on the path between v and q will be informed by $d + \tilde{d} = \rho$ while vertices on the path between v and z (excluding z) will be informed by time $d + 1 - (\tilde{d} - 1) = \rho$.



■ **Figure 6** Two possibilities for informing vertices of a reduced caterpillar (associated with a regular dome). The gates are u and v , where vertex u (respectively v) is responsible for informing the vertices highlighted in blue (respectively red).

Next, assume that there is a broadcast scheme S that completes within ρ rounds. We want to show either (i) or (ii) holds. First, note that in the broadcast tree of S , vertices u and v cannot be in an ancestor-descendent relationship. Otherwise, if u is an ancestor of v (respectively v is an ancestor of u), then q (respectively p) receives the message no earlier than round $\tilde{b} + 2\tilde{d} > \rho$ (respectively $2\tilde{b} + \tilde{d} > \rho$). Second, we note that u (respectively v) cannot receive the message later than round b (respectively d); otherwise p (respectively q) will receive the message after round ρ . Moreover, since z has $\tilde{a} - \tilde{b}$ (which equals $\tilde{c} - \tilde{d}$) neighbor that all are leaves, it must receive the message by round $\rho - (\tilde{a} - \tilde{b}) = a + \tilde{b}$ (which equals $\rho - (\tilde{c} - \tilde{d}) = c + \tilde{d}$) to complete broadcasting within ρ rounds. This means that either u is informed by round a or v is informed by round c (otherwise, z will be informed too late). We conclude that either (i) u is informed by a and v is informed no later than d or (ii) u is informed by b and v is informed by c . ◀

We are now ready to prove the main result of this section.

► **Theorem 18.** *TELEPHONE BROADCASTING problem is NP-complete for snowflake graphs.*

Proof. We use the construction mentioned in Section 4.3 to reduce any instance $\mathcal{D} = (D, m)$ of DOSEPR to an instance $\mathcal{T} = (G, r, \rho)$. Note that each dome D_i in $\mathcal{D}$ is bijected to a reduced caterpillar C_i in G . The two possibilities for selecting an arc from a regular dome in D_i translate to two possibilities for informing the gates of C_i .

Assume there is a valid solution S for $\mathcal{D}$ in DOSEPR. We will show there is a broadcasting scheme S' for $\mathcal{T}$ that completes within ρ rounds. For that, we sort all endpoints of all arcs included in S in non-decreasing order. Consider any dome $D_i \in D$, and let e_1 and e_2 be the endpoints of the arc selected from D_i in S . Assume e_1 and e_2 have ranks i_1 and i_2 in the sorted order. In the broadcast scheme S' , the source r informs the gates of reduced caterpillar C_i associated with D_i at rounds i_1 and i_2 . Given that the ranks of all arc endpoints in S are distinct, r informs at most one neighbor at each given round.

Next, we show that S' completes within round ρ . Given that S is a valid solution for $\mathcal{D}$, for any $t \in m$, we have $\mathcal{N}_S(t) \leq t$, where $\mathcal{N}_S(t)$ is the number of arc endpoints in S that are at most t . On the other hand, for any e that is the endpoint of an arc in S , we have $\text{rank}(e) \leq \mathcal{N}_S(e)$ (be definition, $\text{rank}(e)$ is upper bounded by the number of arc endpoints in S that are at most e , while $\mathcal{N}_S(e)$ is exactly the number of endpoints that are at most e). We

conclude that $\text{rank}(e) \leq e$. Therefore, the gates of any reduced caterpillar C_i in G associated with a singleton dome $D_i = (a, b)$ in D receive the message in S' by rounds (a, b) . Similarly, the gates of any reduced caterpillar C_i associated with a regular dome $D_i = (a, b, c, d)$ receive the message in S' by rounds (a, d) (if the arc $(a, d) \in S$) or by rounds (b, c) (if the arc (b, c) is in S). Therefore, by Lemma 16 and 17, broadcasting in S' completes by round ρ .

For the other side of the reduction, suppose there is a broadcast scheme S' for $\mathcal{T}$ that completes within ρ rounds. We will explain how to construct a valid solution S for $\mathcal{D}$. Consider any reduced caterpillar C_i in G associated with a dome $D_i \in D$. Let (α_i, β_i) be the rounds that center r informs the gateways of C_i . Now, if $D_i = (a, b)$ is a singleton dome, by Lemma 16, it must be that the gates of C_i are informed by rounds (a, b) , that is $\alpha \leq a$ and $\beta \leq b$. In this case, the single arc of D_i is included in S . Next, if $D_i = (a, b, c, d)$ is a regular dome, by Lemma 17, it must be that the gates of C_i are informed by either rounds (a, d) or (b, c) , that is either $(\alpha \leq a \text{ and } \beta \leq d)$ or $(\alpha \leq b \text{ and } \beta \leq c)$ must hold. We will include arc (a, d) in S in the former case and (b, c) in the latter case. In other words, any arc endpoint x in S is associated with a gate that is informed within round x in S' . To show that S is a valid solution for $\mathcal{D}$, we show that for any $t \in [m]$, we have $\mathcal{N}_S(t) \leq t$. Consider the set of endpoints in D that contribute to $\mathcal{N}_S(t)$. As mentioned above, any of these endpoints is associated with a gate that is informed by r within round t . Given that r informs up to t gates by round t , it must be that $\mathcal{N}_S(t) \leq t$. We conclude that S is a valid instance of $\mathcal{D}$.

One can verify that the size of the graph G in the TELEPHONE BROADCASTING instance is polynomial on $n = |D|$ and m . Therefore, given that the DOSEPR is NP-hard by Lemma 15, we conclude that TELEPHONE BROADCASTING is NP-hard. TELEPHONE BROADCASTING in general graphs is known to be in NP [28]. Together, these results establish that the NP-completeness of TELEPHONE BROADCASTING in snowflake graphs. ◀

5 Constant-Factor Approximation for Bounded Pathwidth

Graphs of pathwidth 1 are caterpillars [24], which are special types of trees, for which the TELEPHONE BROADCASTING problem is solvable in linear time [11]. On the other hand, for graphs with a pathwidth larger than 1, the TELEPHONE BROADCASTING problem is NP-hard, as established by our result for the hardness of the problem in snowflake graphs (Theorem 18). Recall that snowflake graphs have pathwidth 2 (Observation 5).

In this section, we establish the existence of a constant-factor approximation for TELEPHONE BROADCASTING on graphs with bounded pathwidth. Recall that we use $\mathbf{br}^*(G, s)$ to denote the optimal broadcasting time for an instance (G, s) . We will demonstrate that the algorithm of Elkin and Kortsarz [9], which has an approximation factor of $\mathcal{O}\left(\frac{\log n}{\log \mathbf{br}^*(G, s)}\right)$ for any graph G , achieves a constant factor approximation for graphs of bounded pathwidth. For general graphs, this algorithm has an approximation factor of $\mathcal{O}(\log n / \log \log n)$ (because $\mathbf{br}^*(G, s) \geq \log n$), which is the best known approximation factor.

For any graph G of pathwidth w , we will show that $\mathbf{br}^*(G, s) = \Omega(n^{4^{-(w+1)}})$, which establishes that the algorithm of Elkin and Kortsarz [9] has a constant factor approximation for graphs of bounded pathwidth w .

To find a lower bound for $\mathbf{br}^*(G, s)$ where G is a graph of constant pathwidth, we repeatedly remove a vertex from G and show that the broadcasting in the remainder of G is not much slower compared to G . For that, we will use the following lemma, which holds for any instance of TELEPHONE BROADCASTING (but we only use it for graphs of bounded pathwidth).

► **Lemma 19.** *Consider any instance (G, s) of TELEPHONE BROADCASTING, and let v be any arbitrary vertex in G . Let $H_1, \dots, H_m$ be the connected components resulting from removing v from G ($m \geq 1$). For any $i \in [m]$, choose an arbitrary vertex $s_i \in H_i$. Then, the following inequality holds: $\sum_{i \in [m]} \mathbf{br}^*(H_i, s_i) \leq \mathbf{br}^*(G, s)(2\mathbf{br}^*(G, s) + 1)$.*

Sketch of the proof. Suppose T^* is an optimal broadcast tree for (G, S) . Removing v from T^* creates disjoint subtrees τ_i . We “glue” these trees arbitrarily to get a spanning tree for each connected component H_i and show that these trees provide the broadcast guarantees stated in the lemma (details in [1]). ◀

We are now ready to state the main result of this section.

► **Theorem 20.** *There is a polynomial-time algorithm for TELEPHONE BROADCASTING in graphs with constant pathwidth w , achieving an approximation ratio of $\mathcal{O}(4^w)$.*

Sketch of the proof. As mentioned above, it suffices to prove a lower bound of $\Omega(n^{4^{-(w+1)}})$ for $\mathbf{br}^*(G, s)$. For that, we assume G is a “ w -path” in the sense that any two vertices that share a bag are neighbors. If G is not a w -path, we will add the missing edges to get a w -path without increasing the pathwidth. Clearly, the addition of new edges does not decrease the broadcast time of G , and since we are looking for a lower bound for the broadcast time, it suffices to focus on w -paths.

We assume a *standard path decomposition* of the input graph G , in which no bag is a subset of another. In such decompositions, we define the notion of *span* of a vertex v as the number of bags where v appears. The span of a decomposition is then the maximum span over the spans of all its vertices.

The proof is established by induction over the size of the input graph G . For that, we consider two cases. First, if the span of the decomposition is “small”, we argue that the diameter of G will be “large”, and then the desired lower bound for $\mathbf{br}^*(G, s)$ holds. On the other hand, when the span is “large”, we consider the vertex v_m that has the maximum span and consider the graph G_{v_m} induced by the vertices that share a bag with v_m . Note that G_{v_m} is a smaller graph compared to G , and we can show that broadcasting in G_{v_m} , starting from any vertex, cannot be much slower than broadcasting in G , starting from s [1]. Moreover, we will use Lemma 19 to show that if we extract v_m from G_{v_m} to get a set $\{H_1, \dots, H_q\}$ of disjoint connected components, total broadcast time in these connected components (starting from arbitrary sources) is not much slower in the absence of v_m . On the other hand, since v_m is removed from G_{v_m} , any of these components H_i has a pathwidth that is at least one unit less than G_{v_m} , and we can use an inductive argument to achieve a lower bound on the broadcast times of any H_i [1]. In summary, by the induction hypothesis, total broadcast time in H_i ’s gives a lower bound for broadcasting in G_{v_m} , which itself gives a lower bound for broadcasting in G . ◀

6 Concluding Remarks

In this paper, we resolved an open problem by proving the NP-completeness of TELEPHONE BROADCASTING for cactus graphs as well as graphs of pathwidth 2. We also established a 2-approximation algorithm for cactus graphs and a constant-factor approximation algorithm for graphs of bounded pathwidth. A possible direction for future work is to improve the approximation factor for graphs of bounded pathwidth or cactus graphs. In particular, it remains an open question whether Polynomial Time Approximation Schemes (PTASs) exist for these graph classes. A major open problem in this domain is determining whether a

constant-factor approximation exists for general graphs. While progress on this question has been slow, the algorithmic ideas developed in this work may be applicable to broadcasting in other families of sparse graphs.

References

- 1 Aida Aminian, Shahin Kamali, Seyed-Mohammad Seyed-Javadi, and Sumedha. On the complexity of telephone broadcasting: From cacti to bounded pathwidth graphs. *arXiv preprint arXiv:2501.12316*, 2025. doi:10.48550/arXiv.2501.12316.
- 2 Amotz Bar-Noy, Sudipto Guha, Joseph Naor, and Baruch Schieber. Message multicasting in heterogeneous networks. *SIAM Journal on Computing*, 30(2):347–358, 2000. doi:10.1137/S0097539798347906.
- 3 Boaz Ben-Moshe, Amit Dvir, Michael Segal, and Arie Tamir. Centdian computation in cactus graphs. *Journal of Graph Algorithms and Applications*, 16(2):199–224, 2012. doi:10.7155/JGAA.00255.
- 4 Puspall Bhabak and Hovhannes A. Harutyunyan. Constant approximation for broadcasting in k-cycle graph. In *Proceedings of first Conference of Algorithms and Discrete Applied Mathematics*, volume 8959, pages 21–32, 2015. doi:10.1007/978-3-319-14974-5_3.
- 5 Maja Čevnik and Janez Žerovnik. Broadcasting on cactus graphs. *Journal of Combinatorial optimization*, 33:292–316, 2017. doi:10.1007/S10878-015-9957-8.
- 6 Peter Damaschke. A linear-time optimal broadcasting algorithm in stars of cliques. *Journal of Graph Algorithms and Applications*, 28(1):385–388, 2024. doi:10.7155/JGAA.V28I1.2981.
- 7 Anne-Laure Ehresmann. Approximation algorithms for broadcasting in flower graphs. Master’s thesis, Concordia University, 2021.
- 8 Michael Elkin and Guy Kortsarz. A combinatorial logarithmic approximation algorithm for the directed telephone broadcast problem. *SIAM Journal on Computing*, 35(3):672, 2005.
- 9 Michael Elkin and Guy Kortsarz. Sublogarithmic approximation for telephone multicast. *Journal of Computer and System Sciences*, 72(4):648–659, 2006. doi:10.1016/J.JCSS.2005.12.002.
- 10 Fedor V Fomin, Pierre Fraigniaud, and Petr A Golovach. Parameterized complexity of broadcasting in graphs. *Theoretical Computer Science*, 997:114508, 2024. doi:10.1016/J.TCS.2024.114508.
- 11 Fraigniaud and Mitjana. Polynomial-time algorithms for minimum-time broadcast in trees. *Theory of Computing Systems*, 35(6):641–665, 2002. doi:10.1007/S00224-002-1047-5.
- 12 Michael R Garey and David S Johnson. *Computers and intractability*, volume 29. wh freeman New York, 2002.
- 13 Saber Gholami, Hovhannes A Harutyunyan, and Edward Maraachlian. Optimal broadcasting in fully connected trees. *Journal of Interconnection Networks*, 23(01):2150037, 2023.
- 14 Michelangelo Grigni and David Peleg. Tight bounds on minimum broadcast networks. *SIAM Journal on Discrete Mathematics*, 4(2):207–222, 1991. doi:10.1137/0404021.
- 15 Hovhannes Harutyunyan, George Laza, and Edward Maraachlian. Broadcasting in necklace graphs. In *Proceedings of the 2nd Canadian Conference on Computer Science and Software Engineering*, pages 253–256, 2009. doi:10.1145/1557626.1557667.
- 16 Hovhannes Harutyunyan and Edward Maraachlian. Linear algorithm for broadcasting in unicyclic graphs. In *Proceedings of the 13th Annual International Conference on Computing and Combinatorics*, pages 372–382, 2007. doi:10.1007/978-3-540-73545-8_37.
- 17 Hovhannes A Harutyunyan and Narek Hovhannisyan. Improved approximation for broadcasting in k-path graphs. In *International Conference on Combinatorial Optimization and Applications*, pages 111–122, 2023. doi:10.1007/978-3-031-49614-1_7.
- 18 Hovhannes A Harutyunyan, Narek Hovhannisyan, and Edward Maraachlian. Broadcasting in chains of rings. In *Proceedings of the Fourteenth International Conference on Ubiquitous and Future Networks*, pages 506–511, 2023. doi:10.1109/ICUFN57995.2023.10201234.

- 19 Hovhannes A Harutyunyan and Edward Maraachlian. On broadcasting in unicyclic graphs. *Journal of Combinatorial Optimization*, 16:307–322, 2008. doi:10.1007/S10878-008-9160-2.
- 20 Sandra M Hedetniemi, Stephen T Hedetniemi, and Arthur L Liestman. A survey of gossiping and broadcasting in communication networks. *Networks*, 18(4):319–349, 1988. doi:10.1002/NET.3230180406.
- 21 Guy Kortsarz and David Peleg. Approximation algorithms for minimum-time broadcast. *SIAM Journal on Discrete Mathematics*, 8(3):401–427, 1995. doi:10.1137/S0895480193245923.
- 22 Arthur L. Liestman and Joseph G. Peters. Broadcast networks of bounded degree. *SIAM Journal on Discrete Mathematics*, 1(4):531–540, 1988. doi:10.1137/0401049.
- 23 Martin Middendorf. Minimum broadcast time is np-complete for 3-regular planar graphs and deadline 2. *Information Processing Letters*, 46(6):281–287, 1993. doi:10.1016/0020-0190(93)90066-I.
- 24 Andrzej Proskurowski and Jan Arne Telle. Classes of graphs with restricted interval models. *Discrete Mathematics and Theoretical Computer Science*, 3(4):167–176, 1999. doi:10.46298/DMTCS.263.
- 25 Ramamoorthi Ravi. Rapid rumor ramification: Approximating the minimum broadcast time. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 202–213, 1994.
- 26 Neil Robertson and Paul D Seymour. Graph minors. i. excluding a forest. *Journal of Combinatorial Theory, Series B*, 35(1):39–61, 1983. doi:10.1016/0095-8956(83)90079-5.
- 27 Weiping Shang, Pengjun Wan, and Xiaodong Hu. Approximation algorithms for minimum broadcast schedule problem in wireless sensor networks. *Frontiers of Mathematics in China*, 5:75–87, 2010.
- 28 Peter J. Slater, Ernest J. Cockayne, and Stephen T. Hedetniemi. Information dissemination in trees. *SIAM Journal on Computing*, 10(4):692–701, 1981. doi:10.1137/0210052.
- 29 Elena Stöhr. Broadcasting in the butterfly network. *Information Processing Letters*, 39(1):41–43, 1991. doi:10.1016/0020-0190(91)90060-U.
- 30 Prafullkumar Tale. Double exponential lower bound for telephone broadcast. *arXiv preprint arXiv:2403.03501*, 2024. doi:10.48550/arXiv.2403.03501.
- 31 Craig A Tovey. A simplified np-complete satisfiability problem. *Discrete Applied Mathematics*, 8(1):85–89, 1984. doi:10.1016/0166-218X(84)90081-7.