

Broadcasting Under Structural Restrictions

Yudai Egami ✉

Kyushu University, Fukuoka, Japan

Tesshu Hanaka ✉ 

Kyushu University, Fukuoka, Japan

Michael Lampis ✉ 

Université Paris-Dauphine, PSL University,
CNRS UMR7243, LAMSADE, Paris, France

Edouard Nemery ✉ 

Université Paris-Dauphine, PSL University,
CNRS UMR7243, LAMSADE, Paris, France

Manolis Vasilakis ✉ 

Université Paris-Dauphine, PSL University,
CNRS UMR7243, LAMSADE, Paris, France

Tatsuya Gima ✉ 

Hokkaido University, Sapporo, Japan

Yasuaki Kobayashi ✉ 

Hokkaido University, Sapporo, Japan

Valia Mitsou ✉

Université Paris Cité, IRIF, CNRS, 75205, Paris,
France

Yota Otachi ✉ 

Nagoya University, Japan

Daniel Vaz ✉ 

LIGM, Université Gustave Eiffel, CNRS, ESIEE
Paris, 77454 Marne-la-Vallée, France

Abstract

In the TELEPHONE BROADCAST problem we are given a graph $G = (V, E)$ with a designated source vertex $s \in V$. Our goal is to transmit a message, which is initially known only to s , to all vertices of the graph by using a process where in each round an informed vertex may transmit the message to one of its uninformed neighbors. The optimization objective is to minimize the number of rounds.

Following up on several recent works, we investigate the structurally parameterized complexity of TELEPHONE BROADCAST. In particular, we first strengthen existing NP-hardness results by showing that the problem remains NP-complete on graphs of bounded tree-depth and also on cactus graphs which are one vertex deletion away from being path forests. Motivated by this (severe) hardness, we study several other parameterizations of the problem and obtain FPT algorithms parameterized by vertex integrity (generalizing a recent FPT algorithm parameterized by vertex cover by Fomin, Fraigniaud, and Golovach [TCS 2024]) and by distance to clique, as well as FPT approximation algorithms parameterized by clique-cover and cluster vertex deletion. Furthermore, we obtain structural results that relate the length of the optimal broadcast protocol of a graph G with its pathwidth and tree-depth. By presenting a substantial improvement over the best previously known bound for pathwidth (Aminian, Kamali, Seyed-Javadi, and Sumedha [ICALP 2025]) we exponentially improve the approximation ratio achievable in polynomial time on graphs of bounded pathwidth from $\mathcal{O}(4^{\text{pw}})$ to $\mathcal{O}(\text{pw})$.

2012 ACM Subject Classification Theory of computation \rightarrow Parameterized complexity and exact algorithms

Keywords and phrases Parameterized Complexity, Structural Graph Parameters, Telephone Broadcast

Digital Object Identifier 10.4230/LIPIcs.MFCS.2025.42

Related Version Full Version: <https://arxiv.org/abs/2504.13669>

Funding This work is partially supported by ANR project ANR-21-CE48-0022 (S-EX-AP-PE-AL).
Tatsuya Gima: JSPS KAKENHI Grant Numbers JP24K23847, JP25K03077.

Tesshu Hanaka: JSPS KAKENHI Grant Numbers JP21H05852, JP21K17707, JP22H00513, JP23H04388, JP25K03077, and JST, CRONOS, Japan Grant Number JPMJCS24K2.

Yasuaki Kobayashi: JSPS KAKENHI Grant Numbers JP20H00595, JP23K28034, JP24H00686, and JP24H00697

Yota Otachi: JSPS KAKENHI Grant Numbers JP21K11752, JP22H00513, JP24H00697, JP25K03076, JP25K03077.



© Yudai Egami, Tatsuya Gima, Tesshu Hanaka, Yasuaki Kobayashi, Michael Lampis, Valia Mitsou, Edouard Nemery, Yota Otachi, Manolis Vasilakis, and Daniel Vaz;
licensed under Creative Commons License CC-BY 4.0

50th International Symposium on Mathematical Foundations of Computer Science (MFCS 2025).

Editors: Paweł Gawrychowski, Filip Mazowiecki, and Michał Skrzypczak; Article No. 42; pp. 42:1–42:18



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Broadcasting in graphs is a well-studied field, with one of the most basic models being the *telephone model* [25]: at the start of the broadcasting process exactly one vertex of the graph (the *source* vertex) is informed (i.e., knows the message), and at the beginning of each round each informed vertex can send the message to at most one of its uninformed neighbors which receives the message by the end of that round. To fix some notation, let G be a connected graph on n vertices and $s \in V(G)$ denote the source vertex. Then, $b(G, s)$ denotes the minimum number of rounds needed to disseminate the message to all vertices of G under the previously described model. As an example, if G is a path and s is one of the degree-1 vertices, then $b(G, s) = n - 1$, while if G is a complete graph, then $b(G, s) = \lceil \log n \rceil$. Notice that for any graph G it holds that $\lceil \log n \rceil \leq b(G, s) \leq n - 1$, and due to the previous examples there exist graphs for which those bounds are tight.

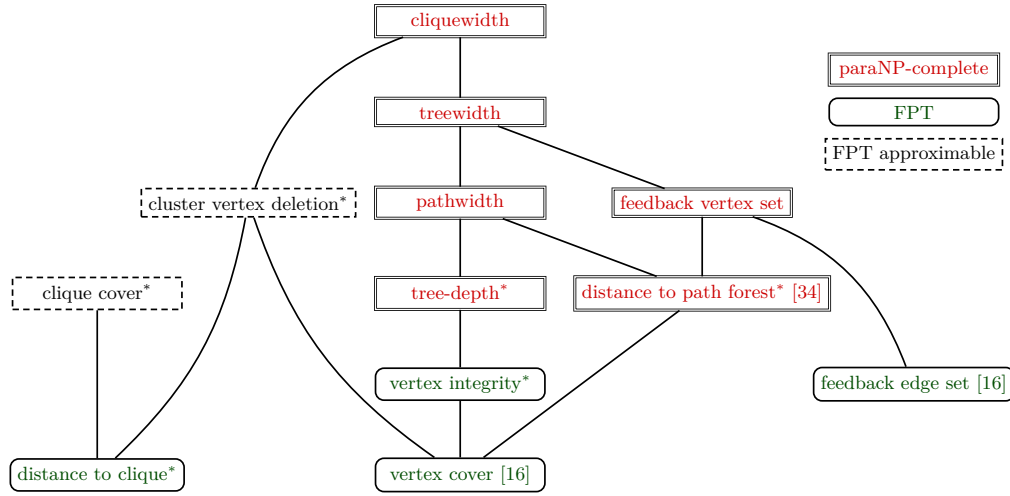
In this paper we study several structural parameterizations¹ of the TELEPHONE BROADCAST problem which asks, given a graph G , a source vertex $s \in V(G)$, and an integer t , whether $b(G, s) \leq t$. In 1981, Slater, Cockayne, and Hedetniemi [33] introduced TELEPHONE BROADCAST and showed that it is NP-complete in general but polynomial-time solvable on trees. Among several results obtained in the intensive study of the problem, the most relevant to us are the very recent ones due to Fomin et al. [16], Tale [34], and Aminian et al. [2], described below.

Fomin, Fraigniaud, and Golovach [16] initiated the study of the parameterized complexity of TELEPHONE BROADCAST. They developed a $3^n n^{\mathcal{O}(1)}$ -time algorithm, where n is the number of vertices, as well as showed that the problem is fixed-parameter tractable parameterized by feedback edge set number, by vertex cover number, and by the number of vertices minus the number of rounds; notice that their first result, along with the fact that $n \leq 2^t$, further implies that the problem is FPT by the number of rounds and leads to an algorithm of double-exponential parametric dependence. Among other open questions they posed, they asked the complexity of TELEPHONE BROADCAST parameterized by other structural parameters and explicitly mentioned treewidth and feedback vertex set number.

Answering the previous question, Tale [34] showed that TELEPHONE BROADCAST is NP-complete on graphs of feedback vertex set number 1 (and thus treewidth 2). As a matter of fact, the proof shows NP-completeness on graphs of vertex deletion distance 2 to path forest (and thus pathwidth 3). Regarding the double-exponential dependence parameterized by the number of rounds t , as discussed in [34], a result of Papadimitriou and Yannakakis [30] implies that TELEPHONE BROADCAST parameterized by t does not admit an algorithm with parameter dependence better than double-exponential (under the ETH). Therefore, the bound obtained by the observation $n \leq 2^t$ applied to the $3^n n^{\mathcal{O}(1)}$ -time algorithm of [16] is essentially optimal.

In a very recent work, Aminian, Kamali, Seyed-Javadi, and Sumedha [2] studied TELEPHONE BROADCAST on cactus graphs and on graphs of bounded pathwidth. They showed that the problem is NP-complete on cactus graphs of pathwidth 2, while they presented a factor-2 approximation algorithm for cactus graphs. Furthermore, they established a connection between the pathwidth pw of the input graph G and $b(G, s)$, by proving that for all $s \in V(G)$ it holds that $b(G, s) = \Omega(n^{4^{-(\text{pw}+1)}})$. A consequence of this result is that the polynomial-time approximation algorithm of Elkin and Kortsarz [14] achieves an approximation ratio of $\mathcal{O}(4^{\text{pw}})$ for graphs of pathwidth pw , improving over the factor $\mathcal{O}(\frac{\log n}{\log \log n})$ achieved in general graphs.

¹ Throughout the paper we assume that the reader is familiar with the basics of parameterized complexity,



■ **Figure 1** Parameterized complexity of TELEPHONE BROADCAST with respect to structural graph parameters. The connection between two parameters represents that the upper parameter p is bounded by some computable function $f(\cdot)$ of the lower parameter q , i.e., $p \leq f(q)$. The parameters marked with * are studied in this paper. The double, rounded, and dotted rectangles indicate paraNP-complete, fixed-parameter tractable, and parameterized approximable cases, respectively. The exact parameterized complexity of the parameterizations in the last case remains unsettled.

Our contribution. Our aim in this paper is to thoroughly study TELEPHONE BROADCAST under the perspective of parameterized complexity; for an overview of our results and the relationship between the mentioned parameters we refer to Figure 1. First we show in Section 3 that the problem remains NP-complete even for very restricted families of graphs, namely (i) cactus graphs of vertex deletion distance 1 to path forest (Theorem 1), and (ii) graphs of constant tree-depth (Theorem 5). The reduction employed in Theorem 1 follows along the lines of the one presented by Tale [34], starting from a variation of 3-PARTITION. However, by selecting a more appropriate variant of 3-PARTITION we are able to significantly simplify both the reduction itself and the constructed graph, thus improving over the initial result by Tale [34] and the subsequent one by Aminian et al. [2] (in the sense that graphs with vertex deletion distance 1 from a path forest, as given in our reduction, are a strict subclass of the snowflake graphs given in [2]). More importantly, by modifying said reduction, we further prove that the problem remains NP-hard for graphs of *constant tree-depth*. Note that tree-depth is a parameter not covered by the previous reductions, which heavily relied on constructions of long paths (which inherently have large tree-depth).

Our negative results extend the intractability of the problem from bounded pathwidth to bounded tree-depth instances. It is therefore natural to ask which (more restrictive) parameterizations *do* render the problem tractable. We present several positive results in this direction in Section 4. In particular, in Theorem 6 we show that the problem is FPT parameterized by vertex integrity, which nicely complements the paraNP-hardness in case of parameterization by tree-depth. To obtain this result, we generalize some of the ideas employed by Fomin et al. [16] to show fixed-parameter tractability by vertex cover number. Next, in Theorem 9 we consider graphs which are k vertex deletions away from being a clique and we develop a $2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$ -time algorithm. Moving on, we consider the

as given in standard textbooks [8].

parameterization by both t and the treewidth tw of the input graph, and in Theorem 12 we develop a single-exponential $2^{\mathcal{O}(t \cdot \text{tw})} n^{\mathcal{O}(1)}$ -time algorithm based on dynamic programming (contrast this with the fact that the problem is intractable when parameterized solely by treewidth and requires a double-exponential dependence when parameterized solely by t). As our last algorithmic results, in Section 5 we develop FPT approximation algorithms when parameterized by clique cover number and by cluster vertex deletion number. With these results we clearly trace the limits of tractability for sparse parameters (tree-depth, vertex integrity) but also improve our understanding of some dense graph parameters.

The $2^{\mathcal{O}(t \cdot \text{tw})} n^{\mathcal{O}(1)}$ algorithm is of particular note because it identifies a small island of tractability for TELEPHONE BROADCAST, namely, it implies that TELEPHONE BROADCAST is polynomial-time solvable for graphs of bounded treewidth (XP) whenever we are interested in a *fast* broadcast schedule ($t = \mathcal{O}(\log n)$). This is interesting, because in general the problem seems quite intractable for most parameters and because all hardness reductions we mentioned construct instances which require a large (polynomial in n) number of rounds. This further motivates the question of which graph families may admit a schedule with *few* rounds, which, as we mention below, has been extensively studied and also has algorithmic implications in terms of approximation. Closing the paper, we improve our understanding of this (structural) question by more carefully studying the relationship between the minimum number of necessary rounds and the pathwidth and tree-depth of a graph. In particular, in Section 6 we show as a main result (Corollary 21) that $b(G, s) = \Omega(n^{(2\text{pw}+2)^{-1}})$. This bound exponentially improves over the corresponding bound of Aminian et al. [2] (who showed that $b(G, s) = \Omega(n^{4^{-(\text{pw}+1)}})$). As an immediate application (Corollary 22), by applying an algorithm of Elkin and Kortsarz [14], we obtain a polynomial-time algorithm with approximation ratio $\mathcal{O}(\text{pw})$ (compared to the ratio of $\mathcal{O}(4^{\text{pw}})$ given in [2]).

Related work. There has been an extensive line of research with regards to broadcasting under the telephone model in graphs. The computational complexity of TELEPHONE BROADCAST has been studied for multiple classes of graphs, with the problem being polynomial-time solvable for trees [33], a subclass of cactus graphs [6], fully connected trees [19], generalized windmill graphs (which is equivalent to graphs of twin-cover number 1) [1, 9], and unicyclic graphs [23], with the latter result having been generalized to the previously mentioned FPT algorithm parameterized by feedback edge set number by Fomin et al. [16]. On the other hand, Jansen and Müller [28] showed that TELEPHONE BROADCAST remains NP-complete on several restricted graph classes such as chordal and grid graphs. We also refer to some relevant surveys in [17, 26]. In another line of research, Grigni and Peleg [20] have provided bounds on the minimum number of edges and the minimum maxdegree for graphs G with the property that $b(G, s) = \lceil \log n \rceil$ for all $s \in V(G)$.

The problem is also well-studied in terms of its polynomial-time approximability. Ravi [31] presented a $\mathcal{O}(\frac{\log^2 n}{\log \log n})$ -factor approximation algorithm, while Kortsarz and Peleg [29] presented (among other results) an algorithm that computes a broadcast protocol with at most $2b(G, s) + \mathcal{O}(\sqrt{n})$ rounds. Subsequently, Bar-Noy, Guha, Naor, and Schieber [3] improved the approximation ratio to $\mathcal{O}(\log n)$, while Elkin and Kortsarz [12] provided a *combinatorial* approximation algorithm of the same approximation ratio, as opposed to the algorithms of [3, 31] that involve solving LPs, that additionally generalizes to directed graphs. The current best approximation ratio of $\mathcal{O}(\frac{\log n}{\log \log n})$ is due to Elkin and Kortsarz [14]; as a matter of fact, the ratio of their algorithm is $\mathcal{O}(\frac{\log n}{\log b(G, s)})$, thus whenever $b(G, s) = \Omega(n^\delta)$ for some constant $\delta > 0$, they obtain a *constant* approximation ratio $\mathcal{O}(1/\delta)$. Better approximation ratios are known for specific graph classes [4, 5, 22, 29], while there are also some inapproximability results [12, 32].

There are several related problems in the literature (see e.g. [35] or the survey [25]), and other models of communication have been considered, such as the *open-line* and *open-path* model [15, 29]. An important generalization of TELEPHONE BROADCAST that has been studied is the TELEPHONE MULTICAST problem, where the objective is to minimize the number of rounds needed to inform a specific *subset* of vertices of the input graph. Most of the previously mentioned approximation algorithms ([3, 12, 14, 31]) generalize to this problem resulting in improved approximation ratios, while directed versions (and variations) of it have also been studied [13, 24].

2 Preliminaries

Throughout the paper we use standard graph notations [10], and we assume familiarity with the basic notions of parameterized complexity [8]. For $x, y \in \mathbb{Z}$, let $[x, y] = \{z \in \mathbb{Z} : x \leq z \leq y\}$, while $[x] = [1, x]$.

All graphs considered are undirected without loops. Given a graph $G = (V, E)$ and a subset of its vertices $S \subseteq V$, $G[S]$ denotes the subgraph induced by S while $G - S$ denotes $G[V \setminus S]$. We denote the set of connected components of G by $\text{cc}(G)$. For two vertices $u, v \in V$, the distance between u and v in G , denoted by $d(u, v)$, is the length of a u - v shortest path. For $C \in \text{cc}(G)$, $d_C(u, v)$ denotes the distance between u and v in the connected component C . The diameter of a connected graph G , denoted by $\text{diam}(G)$, is the maximum among the lengths of its shortest paths, i.e., $\max_{u, v \in V} d(u, v)$.

For a connected graph $G = (V, E)$ and a source $s \in V(G)$, $b(G, s)$ denotes the minimum number of rounds for broadcasting a message from s to all the other vertices.

A *broadcast protocol*, which represents how a message is broadcast from the originator s , is a spanning tree T rooted at s together with a time-stamp function $\tau: E(T) \rightarrow [t]$, informally indicating when an edge is used to transmit a message. More formally, we require that τ must be a proper coloring of the edges of T with the additional property that for each $v \in V \setminus \{s\}$ the edge connecting v to its parent must have minimum value assigned by τ over all edges incident on v . We say that u forwards the message to v at time t_{uv} if u is the parent of v in the tree and $\tau(uv) = t_{uv}$.

Finally, we give a formal definition of TELEPHONE BROADCAST as follows.

TELEPHONE BROADCAST

Instance: Connected graph $G = (V, E)$, $s \in V$, and $t \in \mathbb{Z}^+$.
Goal: Determine whether $b(G, s) \leq t$.

3 NP-hardness Results

In this section we prove that TELEPHONE BROADCAST remains NP-complete even on very restricted families of graphs, namely on (i) cactus graphs that are distance-1 to path forest (Theorem 1) and (ii) graphs of constant tree-depth (Theorem 5). Our work builds upon previous work by Tale [34], who showed that the problem is NP-complete on graphs of distance-2 to path forest. We note that Theorem 1 improves and significantly simplifies both the initial result by Tale [34] as well as a recent result by Aminian et al. [2] who, via a long series of reductions, proved that TELEPHONE BROADCAST remains NP-complete on a special type of cactus graphs coined *snowflake graphs*.²

² Roughly speaking, snowflake graphs are those that are distance-1 to caterpillar forest, so they are a strict superclass of graphs which are distance-1 from path forests.

Since TELEPHONE BROADCAST clearly belongs to NP, it suffices to show the NP-hardness for each case. The starting point of both of our reductions is a restricted version of NUMERICAL MATCHING WITH TARGET SUMS. Given three multisets A, B, C each containing n positive integers such that $\sum_{a \in A} a + \sum_{b \in B} b = \sum_{c \in C} c$, NUMERICAL MATCHING WITH TARGET SUMS asks to determine whether it is possible to partition the $3n$ given numbers into n triples S_1, \dots, S_n such that each triple contains one $c_i \in C$, one $a_j \in A$, and one $b_k \in B$, with $c_i = a_j + b_k$. NUMERICAL MATCHING WITH TARGET SUMS is known to be strongly NP-complete [18, SP17], even when all input numbers are distinct [27, Appendix], that is, when $A \cup B \cup C$ is a set of cardinality $3n$. Starting from such an instance, by first doubling all elements of $A \cup B \cup C$ and then adding 1 to those of $B \cup C$, one can show that the following restricted version of the problem is NP-complete as well.

RESTRICTED NUMERICAL MATCHING WITH TARGET SUMS

- Instance:** Disjoint sets A, B, C , each containing n positive integers in unary, such that $\sum_{a \in A} a + \sum_{b \in B} b = \sum_{c \in C} c$, and the elements of A are even, while those of $B \cup C$ are odd.
- Goal:** Determine whether it is possible to partition the $3n$ given numbers into n triples S_1, \dots, S_n , such that each triple contains one $c_i \in C$, one $a_j \in A$, and one $b_k \in B$, with $c_i = a_j + b_k$.

The ideas behind both reductions of Theorems 1 and 5 are quite similar. We start by introducing the source vertex s , and setting t to be the maximum target sum. Then, for every target sum, we introduce a gadget on roughly that many vertices and connect both of its endpoints to s . On a high level, if the target sum c_i is due to $a_j + b_k$, at rounds $t - a_j$ and $t - b_k$ the source s is supposed to broadcast towards the said gadget; notice that it is therefore crucial that $A \cap B = \emptyset$. To restrict the broadcasting of s during the rest of the rounds, namely on round $h \in [t] \setminus \{t - x : x \in A \cup B\}$, we attach either a long path or a large star to s that essentially forces it to broadcast towards it at round h . Implementing the previous with the gadgets being long paths is sufficient for Theorem 1 to go through. For Theorem 5 however, we need to proceed in two steps and first reduce to the natural generalization of TELEPHONE BROADCAST where there are multiple source vertices before eventually reducing to TELEPHONE BROADCAST.

3.1 Cactus Graphs

We start with the hardness result for cactus graphs that are distance-1 to path forest.

► **Theorem 1.** TELEPHONE BROADCAST is NP-complete for cactus graphs that are distance-1 to path forest.

Proof. Let $\mathcal{I} = (A, B, C)$ be an instance of RESTRICTED NUMERICAL MATCHING WITH TARGET SUMS, where $A = \{a_i : i \in [n]\}$, $B = \{b_i : i \in [n]\}$, and $C = \{c_i : i \in [n]\}$. We will construct an equivalent instance $\mathcal{J} = (G, s, t)$ of TELEPHONE BROADCAST, where $t = \max_{c \in C} c$. Note that $t > \max_{x \in A \cup B} x$.

We construct the graph G as follows:

- Introduce a vertex s ;
- For every $i \in [n]$, introduce a path \mathcal{P}_i on $c_i + 2$ vertices with endpoints u_i and v_i , and add edges $\{s, u_i\}$ and $\{s, v_i\}$;
- For every $h \in [t] \setminus \{t - x : x \in A \cup B\}$, introduce a path \mathcal{Q}_h on $t - h + 1$ vertices, and add an edge from one of its endpoints to s . Name the other endpoint q_h (for $h = t$, we name the single vertex of \mathcal{Q}_t as q_t instead).

This concludes the construction of G . Notice that G is a cactus graph, while $G - s$ is a collection of paths. Before proving the equivalence of the two instances, we first compute the size of graph G .

▷ **Claim 2.** It holds that $|V(G)| = 1 + \frac{t(t+1)}{2}$.

Proof. Notice that G is comprised of vertex s , as well as the vertices present in \mathcal{P}_i and \mathcal{Q}_h , for $i \in [n]$ and $h \in [t] \setminus \{t - x : x \in A \cup B\}$. Consequently,

$$\begin{aligned}
 |V(G)| &= 1 + \sum_{i=1}^n (c_i + 2) + \sum_{h \in [t] \setminus \{t-x: x \in A \cup B\}} (t - h + 1) \\
 &= 1 + 2n + \sum_{i=1}^n c_i + (t - 2n) + \sum_{h \in [t] \setminus \{t-x: x \in A \cup B\}} (t - h) \\
 &= 1 + t + \sum_{i=1}^n c_i + \sum_{h \in [t]} (t - h) - \sum_{h \in \{t-x: x \in A \cup B\}} (t - h) \\
 &= 1 + \sum_{i=1}^n c_i + \sum_{h \in [t]} h - \sum_{x \in A \cup B} x \\
 &= 1 + \frac{t(t+1)}{2},
 \end{aligned}$$

and the statement follows. \triangleleft

▷ **Claim 3.** If \mathcal{I} is a yes-instance of RESTRICTED NUMERICAL MATCHING WITH TARGET SUMS, then \mathcal{J} is a yes-instance of TELEPHONE BROADCAST.

Proof. Let (S_1, \dots, S_n) be a partition of $A \cup B \cup C$ that certifies that \mathcal{I} is a yes-instance. Rename appropriately the elements of $A \cup B \cup C$ so that $S_i = \{a_i, b_i, c_i\}$ for all $i \in [n]$. Let $h \in [t]$ denote the current round, and consider the following broadcast protocol.

- If $h \in [t] \setminus \{t - x : x \in A \cup B\}$, then set s to forward the message along the path \mathcal{Q}_h . In the subsequent rounds, the path from s to q_h keeps propelling the message towards q_h .
- Otherwise, it holds that $h = t - x$ for some $x \in A \cup B$. If $x = a_i$, then set s to forward the message to u_i , else if $x = b_i$, then towards v_i instead. In the subsequent rounds, the message keeps getting propelled along the path \mathcal{P}_i towards its other endpoint.

This completes the description of the protocol.

We now prove that every vertex of G gets the message by the end of round t . Regarding the vertices belonging to \mathcal{Q}_h , notice that since s forwards the message at round h , and \mathcal{Q}_h contains $t - h + 1$ vertices, all vertices of \mathcal{Q}_h receive the message by the end of round t ; as a matter of fact, q_h receives the message exactly on the last round t .

It remains to argue for the vertices of \mathcal{P}_i . By the protocol, if $S_i = \{a_i, b_i, c_i\}$, one endpoint of \mathcal{P}_i gets the message at round $t - a_i$, while the other at round $t - b_i$. These two vertices forward the message towards each other to convey it to all the vertices in the path connecting them. In that case, the number of vertices of \mathcal{P}_i that receive the message by round t is at least $t - (t - a_i) + 1 + t - (t - b_i) + 1 = a_i + b_i + 2 = c_i + 2$, thus all of its vertices receive the message by the end of round t . \triangleleft

▷ **Claim 4.** If \mathcal{J} is a yes-instance of TELEPHONE BROADCAST, then \mathcal{I} is a yes-instance of RESTRICTED NUMERICAL MATCHING WITH TARGET SUMS.

Proof. Assume there exists a broadcast protocol such that every vertex of G receives the message by the end of round t . We claim that the maximum number of vertices that have received the message after $h \in [t]$ rounds of transmission is $1 + \sum_{i=1}^h i = 1 + \frac{h(h+1)}{2}$. To see this, notice that every vertex apart from s has degree at most 2, thus it can broadcast the message at most once.³ In that case, the number of vertices receiving the message at round h is at most those that received the message at round $h-1$ plus 1 due to s .

Since $|V(G)| = 1 + \frac{t(t+1)}{2}$ due to Claim 2, it follows that this is indeed the case, thus every vertex apart from s and those that are informed in the last round broadcasts the message exactly once. Consequently, every vertex of degree 1 receives the message at the last round, thus s broadcasts the message towards the path Q_h at round h , for all $h \in [t] \setminus \{t-x : x \in A \cup B\}$.

At round $h \in \{t-x : x \in A \cup B\}$, s broadcasts the message towards some path \mathcal{P}_i , and this broadcast results in a total of $x+1$ vertices of path \mathcal{P}_i having received the message by round t . Consequently, it follows that for every path \mathcal{P}_i , s broadcasts towards u_i and v_i at rounds $t-x_i^1$ and $t-x_i^2$, where $t-(t-x_i^1)+1+t-(t-x_i^2)+1 = x_i^1+x_i^2+2 = c_i+2$. Since c_i is odd, that is also the case for exactly one among x_i^1 and x_i^2 , thus $\{x_i^1, x_i^2\}$ contains an element from both A and B . In that case, (S_1, \dots, S_n) certifies that \mathcal{I} is a yes-instance, since it is a partition of $A \cup B \cup C$ where $S_i = \{c_i, x_i^1, x_i^2\}$, each S_i contains a single element out of A , B , and C , while $c_i = x_i^1 + x_i^2$. This completes the proof. \triangleleft

By Claims 3 and 4, Theorem 1 follows. \blacktriangleleft

3.2 Bounded Tree-depth

Now we show the hardness for graphs of tree-depth at most 6. To this end, we introduce a generalization of TELEPHONE BROADCAST as an intermediate problem, where instead of a unique vertex as the initial source, we have a set of vertices. We call the variant TELEPHONE BROADCAST WITH MULTIPLE SOURCES and formalize it as follows, where $b(G, S)$ is defined analogously to $b(G, s)$.

TELEPHONE BROADCAST WITH MULTIPLE SOURCES

Instance: Connected graph $G = (V, E)$, $S \subseteq V$, and $t \in \mathbb{Z}^+$.
Goal: Determine whether $b(G, S) \leq t$.

Our proof thus consists of the following two steps. First, we reduce RESTRICTED NUMERICAL MATCHING WITH TARGET SUMS to TELEPHONE BROADCAST WITH MULTIPLE SOURCES. On a high level, the reduction is similar to the one of Theorem 1, albeit with a few supplementary steps to bound the tree-depth. The main difference is that we use star-like structures equipped with multiple sources instead of paths. Next, we present a reduction from TELEPHONE BROADCAST WITH MULTIPLE SOURCES to TELEPHONE BROADCAST that does not increase tree-depth too much; as a matter of fact, if this reduction is applied to a graph output by the first step, then the tree-depth does not increase at all. We again use star-like structures here.

We highlight that instead of independently presenting the two reductions in distinct theorems, we choose to present them in one, as doing so allows us to present a tighter bound on the tree-depth of the final graph.

► **Theorem 5.** (\star) TELEPHONE BROADCAST is NP-complete for graphs of tree-depth at most 6.

³ We assume that no vertex broadcasts the message towards its already informed neighbors.

4 Fixed-parameter Algorithms

4.1 FPT Algorithm Parameterized by Vertex Integrity

Let $G = (V, E)$ be a graph. The *vertex integrity* of G , denoted $\text{vi}(G)$, is the minimum integer k such that there is a vertex set $S \subseteq V$ with $|S| + \max_{C \in \text{cc}(G-S)} |V(C)| \leq k$. It is known that such a set can be computed by a fixed-parameter tractable algorithm parameterized by k [11].

In this section, we present a fixed-parameter tractable algorithm for TELEPHONE BROADCAST parameterized by vertex integrity.

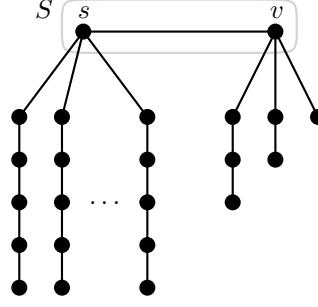
► **Theorem 6.** *TELEPHONE BROADCAST is fixed-parameter tractable parameterized by vertex integrity.*

Overview. Our algorithm can be seen as a generalization of the fixed-parameter tractable algorithm parameterized by vertex cover number due to Fomin et al. [16]. Since the ideas are similar, let us present a brief overview of their algorithm. Let $G = (V, E)$ be a graph with a vertex cover S of size k . They first show that there is an optimal broadcast protocol such that all vertices in the vertex cover S receive the message in the first $2k - 1$ rounds. Since a broadcast protocol with $2k - 1$ rounds involves at most 2^{2k-1} vertices and the vertices in $V \setminus S$ can be partitioned into 2^k classes of *twins*, one can guess which vertices are involved in the first $2k - 1$ rounds. After testing the feasibility of the guessed vertices, it only remains to check whether the completely informed vertex cover S can send the message to the independent set $V \setminus S$ within a given number of rounds. This problem can be solved in polynomial time by a reduction to the BIPARTITE MATCHING problem.

Our algorithm follows the same general strategy, however there are two obstacles one needs to overcome in order to generalize from vertex cover number to vertex integrity. Let $S \subseteq V$ be a set satisfying $|S| + \max_{C \in \text{cc}(G-S)} |V(C)| \leq k$. The first and main obstacle is showing the existence of an optimal broadcast protocol in which the vertices in S receive the message in the first few rounds. Unfortunately, there are instances that do not admit such optimal broadcast protocols (see Figure 2); what we show instead is that there is an optimal broadcast protocol in which the vertices in S receive the message in the *first* and *last* few rounds. The second obstacle is that after the removal of S , we have a set of small connected components instead of an independent set. To deal with this, we show that most of the components of $G - S$ receive the message from S only once and do not send the message back to S . This property allows us to handle these components of $G - S$ as single vertices forming an independent set in some sense.

Some definitions. In the following, we fix a graph $G = (V, E)$ and a source vertex $s \in V$. We also fix a vertex set $S \subseteq V$ that minimizes the sum $|S| + \max_{C \in \text{cc}(G-S)} |V(C)|$ under the condition that $s \in S$. Such S can be found by first finding a set S' minimizing $|S'| + \max_{C \in \text{cc}((G-S)-S')} |V(C)|$, and then setting $S = S' \cup \{s\}$. Let $k = |S| + \max_{C \in \text{cc}(G-S)} |V(C)|$. Note that $k \leq \text{vi}(G) + 1$.

We say that C and C' in $\text{cc}(G-S)$ have the same *type* if G admits an automorphism η such that $\eta(V(C)) = V(C')$, $\eta(V(C')) = V(C)$, and $\eta(v) = v$ for each $v \in V \setminus (V(C) \cup V(C'))$. By the equivalence relation of having the same type, we partition $\text{cc}(G-S)$ into the equivalence classes $\mathcal{C}_1, \dots, \mathcal{C}_p$. We can upper-bound p , the number of types, by a function of k ; e.g., $p < 2^{k^2}$, which is an upper bound of the number of labeled graphs formed by at most k vertices, i.e., by the vertices in S and a component $C \in \text{cc}(G-S)$. Observe that, in a broadcast protocol, two components of $G-S$ with the same type can be used interchangeably.



■ **Figure 2** The set $S = \{s, v\}$ is the unique set satisfying that $|S| + \max_{C \in \text{cc}(G-S)} |V(C)| \leq 7$ ($= \text{vi}(G)$). In all optimal broadcast protocols, v is the last vertex that receives the message from s as v needs only three rounds after it receives the message, while other neighbors need four rounds.

Let P be a broadcast protocol for (G, s) . We denote by $r(P)$ the number of rounds in P . For $C \in \text{cc}(G - S)$, let $a_P(C)$ denote the *arrival time* at C in P ; that is, $a_P(C)$ is the first round when some vertex in S forwards the message to some vertex in C . We call $C \in \text{cc}(G - S)$ a *leaf component* of P if no vertex of C forwards the message to vertices of S . We call the vertices in the leaf components *unimportant* and the rest of the vertices *important*. That is, the important vertices are those belonging to S as well as in the non-leaf components of $G - S$.

4.1.1 Optimal Broadcast Protocols with Nice Properties

We say that an optimal broadcast protocol P for (G, s) is *S-lazy* if, among all optimal broadcast protocols, P minimizes the total number of message-forwardings from S .

► **Lemma 7.** (\star) *If P is an S-lazy optimal broadcast protocol for (G, s) , then every leaf component C with $a_P(C) \leq r(P) - k + 1$ receives the message from S exactly once.*

► **Lemma 8.** (\star) *There is an S-lazy optimal broadcast protocol P such that all important vertices receive the message only in the first $(2k - 3)k$ rounds and the last $2k$ rounds.*

4.1.2 The Algorithm Parameterized by Vertex Integrity

Now we present a fixed-parameter tractable algorithm parameterized by vertex integrity for finding an optimal broadcast protocol for (G, s) . We first guess the number of rounds t in an optimal broadcast protocol. This is possible as $t < |V|$. If $t < 2k^2$, then we apply the fixed-parameter algorithm parameterized by t [16]. In the following, we assume that $t \geq 2k^2$.

We find an optimal broadcast protocol P for (G, s) with t rounds such that

- every leaf component receives the message from S exactly once, and
- the important vertices receive the message only in the first $(2k - 3)k$ rounds and the last $2k$ rounds.

Such a protocol exists by Lemmas 7 and 8. As in the proof of Lemma 8, we also assume that no vertex takes a break in P .

To describe the algorithm, we partition the set of rounds $[t]$ in P into three subsets of consecutive rounds $R_1 = [1, (2k - 3)k]$, $R_2 = [(2k - 3)k + 1, t - 2k]$, and $R_3 = [t - 2k + 1, t]$. Note that $|R_2| = (t - 2k) - ((2k - 3)k + 1) + 1 = t - 2k^2 + k \geq k$ as $t \geq 2k^2$.

The algorithm works as follows.

0. Guess $S_1 \subseteq S$ and $\mathcal{C}_1, \mathcal{C}_3 \subseteq \text{cc}(G - S)$ as follows:
 - $s \in S_1$ and $S_1 \setminus \{s\}$ is the set of vertices in S that receive the message in R_1 ;
 - $\mathcal{C}_1 = \{C \in \text{cc}(G - S) : a_P(C) \in R_1\}$ with $|\mathcal{C}_1| \leq (2k - 3)k^2$;
 - $\mathcal{C}_3 = \{C \in \text{cc}(G - S) : a_P(C) \in R_3\}$ with $|\mathcal{C}_3| \leq 2k^2$.
1. Check if the combination of S_1 and \mathcal{C}_1 is feasible: that is, the algorithm checks whether there is a *partial* broadcast protocol P_1 for $G[S_1 \cup \bigcup_{C \in \mathcal{C}_1} V(C)]$ with $(2k - 3)k$ rounds that starts at s and informs all vertices in S_1 and at least one vertex of each $C \in \mathcal{C}_1$.
2. Let $\mathcal{C}_2 = \text{cc}(G - S) \setminus (\mathcal{C}_1 \cup \mathcal{C}_3)$. Check if S_1 can forward the message to each of \mathcal{C}_2 in R_2 : that is, the algorithm checks whether there is a *multi-source partial* broadcast protocol P_2 for $G[S_1 \cup \bigcup_{C \in \mathcal{C}_2} V(C)]$ with $t - (2k - 1)k$ rounds that starts at S_1 and informs exactly one vertex (equivalently, at least one vertex) of each $C \in \mathcal{C}_2$.
3. Check if the combination of S_1 and \mathcal{C}_3 is feasible: that is, the algorithm checks whether there is a *multi-source* broadcast protocol P_3 for $G[S \cup \bigcup_{C \in \mathcal{C}_3} V(C)]$ with $2k$ rounds that starts at S_1 .

Correctness. We first show that an optimal broadcast protocol for (G, s) with t rounds exists if and only if correctly guessed $S_1, \mathcal{C}_1, \mathcal{C}_3$ pass all three tests by the algorithm.

If an optimal broadcast protocol for (G, s) with t rounds exists, then there is one with the properties that we assumed in the algorithm. Let P be such a protocol. We construct $S_1, \mathcal{C}_1, \mathcal{C}_3$ from P as defined in the algorithm. Observe that at each round, at most $|S| \leq k$ components in $\text{cc}(G - S)$ may receive the message from S , and thus $|\{C \in \text{cc}(G - S) : a_P(C) \in R_1\}| \leq (2k - 3)k^2$ holds. By the same reasoning, we have $|\{C \in \text{cc}(G - S) : a_P(C) \in R_3\}| \leq 2k^2$. Clearly, $S_1, \mathcal{C}_1, \mathcal{C}_3$ pass all three tests.

Conversely, assume that some $S_1, \mathcal{C}_1, \mathcal{C}_3$ pass all three tests by the algorithm. Let P_1, P_2, P_3 be the protocols found by the algorithm. Let P' be the protocol obtained by concatenating P_1, P_2, P_3 in this ordering. In P' , some vertices of the components in \mathcal{C}_1 and \mathcal{C}_2 may remain uninformed, while all other vertices receive the message. We modify P' to obtain a desired protocol. We first consider a component $C_1 \in \mathcal{C}_1$ that has uninformed vertices in P' . In P' , at least one vertex in C_1 receives the message in P_1 and no vertex in C_1 sends or receives the message in P_2 . Thus, we can modify P' in such a way that the vertices in C_1 send and receive the message locally in C_1 and inform all vertices of C_1 using at most $k - 1$ rounds in P_2 . Next we consider a component $C_2 \in \mathcal{C}_2$ that has uninformed vertices in P' . Similarly to the previous case, we can modify P' locally in C_2 to inform all vertices of C_2 using at most $k - 1$ round in P_3 . Now the obtained protocol is a broadcast protocol for (G, s) with t rounds.

Running time. Now we show that the algorithm can be implemented as a fixed-parameter tractable algorithm parameterized by k .

We first estimate the number of possible candidates for guessing S_1, \mathcal{C}_1 , and \mathcal{C}_3 . For S_1 , we simply try all $2^{|S|-1} = 2^{k-1}$ subsets of S that contain s . For \mathcal{C}_i ($i \in \{1, 3\}$), it suffices to try all possibilities of how many components of each type are included in \mathcal{C}_i . Since the number of types is less than 2^{k^2} , the number of candidates for \mathcal{C}_i can be bounded from above by a function in k ; i.e., $2^{k^2 \cdot (2k-3)k^2}$ for \mathcal{C}_1 and $2^{k^2 \cdot 2k^2}$ for \mathcal{C}_3 . In total, we test at most 2^{2k^5} candidates for the combination of $S_1, \mathcal{C}_1, \mathcal{C}_3$. Thus, it suffices to show that, for a fixed combination of $S_1, \mathcal{C}_1, \mathcal{C}_3$, each of the three tests can be done with a fixed-parameter tractable algorithm parameterized by k .

The first test can be done in time depending only on k because the question is decidable and the target graph $G[S_1 \cup \bigcup_{C \in \mathcal{C}_1} V(C)]$ has at most $k + |\mathcal{C}_1|k < 2k^4$ vertices. In the same way, we can see that the third test can be done in time depending only on k .

For the second test, we can use the polynomial-time subroutine that Fomin et al. [16] used for their algorithm parameterized by vertex cover number. The subroutine checks in polynomial time whether there is a multi-source broadcast protocol with a given number of rounds starting at a vertex cover. To use this subroutine, we construct a graph G_2 from $G[S_1 \cup \bigcup_{C \in \mathcal{C}_2} V(C)]$ by replacing each component C_2 in \mathcal{C}_2 with a single vertex that is adjacent to the vertices in $S_1 \cap N(V(C_2))$. Observe that the desired multi-source partial broadcast protocol P_2 for $G[S_1 \cup \bigcup_{C \in \mathcal{C}_2} V(C)]$ exists if and only if there is a multi-source broadcast protocol for G_2 with $t - (2k - 1)k$ rounds starting at S_1 . Since S_1 is a vertex cover of G_2 , we can use the subroutine of Fomin et al. [16].

4.2 FPT Algorithm Parameterized by Distance to Clique

In this section we consider TELEPHONE BROADCAST parameterized by distance to clique k , in which the input graph contains a subset $X \subseteq V$ of size k such that $G - X$ is a clique. We call X a *modulator* of G and let $n = |V| - k$. Computing a (smallest) modulator is equivalent to computing a (minimum) vertex cover and thus can be done in time $2^k n^{\mathcal{O}(1)}$ [7, 21].

We show that, in these instances, the amount of information necessary to encode a solution is bounded, in that we can recreate a solution knowing the following about the vertices of the modulator: on which round each is informed, how many vertices each informs, as well as some details about the vertices that inform them (see Lemma 11 and the description below for details). We solve the problem by enumerating all possibilities for this information, and then recreating a matching solution in polynomial time.

Our approach is based on the following two facts: (i) $b(G, s) \leq k + \lceil \log_2 n \rceil$, and (ii) to encode a solution we only need to know the information described above.

The rest of this section is dedicated to proving the following theorem.

► **Theorem 9.** *TELEPHONE BROADCAST is fixed-parameter tractable parameterized by distance to clique k , and can be solved in time $2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$.*

We start by formalizing and proving the two facts mentioned above:

► **Lemma 10.** *(*) Let $G = (V, E)$ be a connected graph and $X \subseteq V$ with $|X| = k$ such that $G - X$ is a clique of size n . Then $b(G, s) \leq k + \lceil \log_2 n \rceil$ for all $s \in V$.*

The lemma below formalizes the process of recreating a solution, that is, given restricted information about a broadcast protocol, computing a protocol that matches that information. For simplicity of presentation, the lemma lists in compact form the information that is needed about the broadcast protocol, and we provide additional explanation below the lemma.

► **Lemma 11.** *(*) Let (G, s) be an instance of TELEPHONE BROADCAST, X be a modulator of G of size k , and P be a broadcast protocol. We can construct in polynomial time a broadcast protocol P' with the same number of rounds as P , given the following information about P :*

1. t , the number of rounds of P ;
2. for each $v \in X$, the round τ_v in which v is informed;
3. for each $v \in X$, the vertex a_v that informs it (its informer), where a vertex in $V \setminus (X \cup \{s\})$ is identified by its neighborhood in X as well as the set of vertices in X it informs;
4. for each informer a_v , $v \in X$, its informer b_v if it is in X , or $b_v = \phi$ otherwise (representing either an informer in $V \setminus X$ or that $a_v = s$);
5. for each $v \in X$, the number c_v of vertices that it informs, that is, how many times it transmits the message to an uninformed vertex.

For Point 3, we do not need to know the exact vertex of the clique that informs a vertex $v \in X$, as there can be many clique vertices that are functionally equivalent. Instead, we identify vertices in $V \setminus (X \cup \{s\})$ by their neighborhood in X ; if multiple vertices in X share an informer, the first (in an arbitrary order) encodes it by its neighborhood, while the following simply indicate that they have the same informer as a previous vertex. We can thus encode the informer of $u \in X$ using a number i_u from 0 to $2^k + 2k$ as follows. Take an arbitrary ordering of $X = \{v_0, \dots, v_{k-1}\}$. If $i_u < 2^k$, a_u is a (yet unused) vertex with neighborhood in X given by the binary representation of i_u ; if $2^k \leq i_u < 2^k + k$, a_u is the vertex of X with index $i_u - 2^k$ (i.e. $v_{i_u - 2^k}$); if $2^k + k \leq i_u < 2^k + 2k$, a_u is the informer of the vertex of X with index $i_u - 2^k - k$, that is, $a_u = a_w$ with $w = v_{i_u - 2^k - k}$; and if $i_u = 2^k + 2k$, then $a_u = s$. For Point 4, we need to know the informer b_v of each a_v ; however, if b_v is in $V \setminus X$, the actual vertex is irrelevant, and thus we simply notate this case as $b_v = \phi$; we notate it similarly if $a_u = s$, in which case it has no informer.

Let us see how this lemma can be applied to solve the problem.

The algorithm starts by computing a modulator X of size at most k . We remark that a modulator of G is a vertex cover in \bar{G} , and thus the minimum modulator can be computed by employing any FPT algorithm for VERTEX COVER (see e.g. [7, 21]) on \bar{G} .

Then, it simply enumerates the information required by Lemma 11; for each possibility it uses the procedure in the lemma to obtain a protocol. If there is a feasible solution to the problem, Lemma 11 guarantees that the algorithm obtains a valid protocol when it uses the correct information, and thus the algorithm is correct.

As to the running time, it is dominated by the enumeration of all the options for Lemma 11: there are $\lceil \log_2 n \rceil + k$ possibilities for t by Lemma 10, and the same number of possibilities for each τ_v , and each c_v ; for each a_v , there are $2k + 2^k$ options, corresponding to the vertices in X , their informers (if vertices share an informer), or a new vertex from one of each 2^k neighborhood classes; for each b_v , there are $k + 1$ choices. Thus, the total number of possibilities is

$$(\lceil \log_2 n \rceil + k)^{2k+1} \cdot (2k + 2^k)^k \cdot (k + 1)^k = (\lceil \log_2 n \rceil)^{k+1} \cdot k^{\mathcal{O}(k)} \cdot 2^{k^2}.$$

As the procedure in Lemma 11 runs in polynomial time, the running time of the algorithm is FPT and bounded by $(2^k \log n)^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$. We can achieve the running time bound of $2^{\mathcal{O}(k^2)}$ by case distinction: if $\log n \leq 2^k$, then the bound follows by substitution; otherwise, we substitute $k < \log \log n$ and get a running time of $2^{\mathcal{O}((\log \log n)^2)} = 2^{\mathcal{O}(\log n)} = n^{\mathcal{O}(1)}$.

4.3 Algorithm for Parameter Treewidth plus Time

In this section we show, via a dynamic programming algorithm, that TELEPHONE BROADCAST admits a single-exponential FPT algorithm if we parameterize by both treewidth and the target number of rounds t . Let us first give an informal high-level overview of the main ideas.

The strategy of our algorithm will be to compute the proper edge-coloring which assigns a time-stamp to each edge used by an optimal protocol. Intuitively, it is natural that this leads to a complexity of the form $2^{t \cdot \text{tw}}$, because for each vertex v of a bag we need to keep track of the set of colors which have already been used on an edge incident on v . We therefore store a subset of $[t]$ for each vertex of a bag. To facilitate the algorithm we will also keep track of some further information, notably including the time at which a vertex is informed by the protocol (but this only adds a further t^{tw} factor to the complexity). This extra information will allow us to infer the direction in which each edge is supposed to be used (that is, which endpoint transmits the message and which receives it) and verify that the protocol is consistent (that is, each vertex starts transmitting the message after receiving it).

► **Theorem 12.** (\star) *There is an algorithm which takes as input an n -vertex graph $G = (V, E)$, a vertex $s \in V$, an integer t , and a nice tree decomposition of G of width tw and determines whether $b(G, s) \leq t$ in time $2^{\mathcal{O}(t \cdot \text{tw})} n^{\mathcal{O}(1)}$.*

5 Parameterized Approximation Results

In this section, we present FPT approximation algorithms parameterized by clique cover number and cluster vertex deletion number. Both parameters are more general than distance to clique. The precise parameterized complexity for the two parameters remains unsettled.

5.1 Clique Cover Number

For a graph G , a partition of its vertex set $V(G)$ into cliques is a *clique cover*. The minimum integer p such that G admits a clique cover with p cliques is the *clique cover number*. The clique cover number of a graph is equal to the chromatic number of its complement. In this subsection, we show that the optimization version of TELEPHONE BROADCAST admits an FPT-AS (i.e., an $f(p, \varepsilon) n^{\mathcal{O}(1)}$ -time $(1 + \varepsilon)$ -approximation algorithm for $\varepsilon > 0$) when parameterized by clique cover number p .

► **Theorem 13.** (\star) *Given a connected n -vertex graph G , $s \in V(G)$, and a clique cover Q_1, \dots, Q_p of G , one can find a broadcast protocol with at most $(1 + \varepsilon) \cdot b(G, s)$ rounds in time $3^{2^{p/\varepsilon}} p \cdot n^{\mathcal{O}(1)}$.*

5.2 Cluster Vertex Deletion Number

For a graph G , a set $S \subseteq V(G)$ is a *cluster vertex deletion* if each connected component of $G - S$ is a complete graph. The *cluster vertex deletion number* of G is the size of a minimum cluster vertex deletion set of G .

► **Theorem 14.** (\star) *Given a connected graph G with cluster vertex deletion number cvd and $s \in V(G)$, one can find a broadcast protocol with at most $(2 + \varepsilon) \cdot b(G, s)$ rounds in time $f(\text{cvd}, \varepsilon) \cdot n^{\mathcal{O}(1)}$.*

6 Lower Bounds with Respect to Structural Parameters

Here we present several results showcasing the relationship between the structure of the input graph and the minimum number of rounds required by any broadcast protocol to complete. On an intuitive level, graphs close to being stars or paths require a large (that is, polynomial in the size of the graph) number of rounds to broadcast the message, and the aim of this section is to quantify this intuition and provide corresponding lower bounds with respect to the tree-depth and the pathwidth of the input graph. We start with Lemmas 15–17, which are essential building blocks in proving the main results of this section in Theorems 18 and 20. We note that Theorem 20 improves over an analogous lower bound recently obtained by Aminian et al. [2] and leads to Corollaries 21 and 22.

► **Lemma 15.** *Given a connected graph G and a non-empty set $S \subseteq V(G)$, for all $s \in V(G)$ it holds that $b(G, s) \geq |\text{cc}(G - S)|/|S|$.*

Proof. Let a component of $G - S$ be *active* when there exists a vertex belonging to it that has received the message. It holds that on every round the number of additional components of $G - S$ that become active is at most $|S|$. In that case, if $s \in S$, then at least $|\text{cc}(G - S)|/|S|$

rounds are required in order to make all of the components of $G - S$ active. Notice that the lower bound holds in the case $s \notin S$ as well, since then, although there exists a component which is already active at the start, after the first round of the protocol no additional component of $G - S$ has become active. ◀

► **Lemma 16.** *Given a connected graph G , for all $s \in V(G)$ it holds that $b(G, s) \geq \text{diam}(G)/2$.*

Proof. Consider a protocol of minimum broadcasting time $b(G, s)$ as well as a shortest path P in G such that $\text{diam}(G) = |V(P)| - 1$ with u_1 and u_2 denoting its endpoints. For every $v \in V(G)$ it holds that $d(v, s) \leq b(G, s)$, that is, every vertex of G is at distance at most $b(G, s)$ from s ; if that were not the case v cannot receive the message in time. Finally, since P is a shortest path it holds that $\text{diam}(G) = d(u_1, u_2) \leq d(u_1, s) + d(s, u_2) \leq 2b(G, s)$ and the statement follows. ◀

► **Lemma 17.** (\star) *Given a connected graph G and a non-empty set $S \subseteq V(G)$, for all $s \in V(G)$ it holds that*

$$b(G, s) \geq \sqrt{\frac{\sum_{C \in \text{cc}(G-S)} \text{diam}(C)}{2|S|}}.$$

We are now ready to proceed with the main result of this section concerning tree-depth.

► **Theorem 18.** *Let G be a connected graph with $n > 1$ vertices of tree-depth td . Then, for all $s \in V(G)$ it holds that $b(G, s) \geq n^{1/\text{td}}/\text{td}$.*

Proof. Since G is of tree-depth td , there exists an *elimination tree* T of G of height at most td . This is a rooted tree on the same vertex set, with r denoting the root vertex, where every pair of vertices adjacent in G adheres to the ancestor/descendant relation.

Notice that if $n^{1/\text{td}} \leq 2$, then the statement holds, unless G is an isolated vertex. In the following, assume that $n^{1/\text{td}} > 2$. We will prove that there exists a vertex with at least $n^{1/\text{td}}$ children in T . Assume that this is not the case. Then, it holds that

$$n = |V(T)| \leq \sum_{i=0}^{\text{td}-1} n^{(1/\text{td}) \cdot i} = \frac{n - 1}{n^{1/\text{td}} - 1} < n - 1,$$

which is a contradiction. Consequently, there exists a vertex u with at least $n^{1/\text{td}}$ children in T .

Let $S \subseteq V(G)$ contain the vertices present in the u - r path of T , where $1 \leq |S| \leq \text{td}$. Then, $G - S$ has at least $n^{1/\text{td}}$ connected components, 1 per child of u in T , and the statement follows due to Lemma 15. ◀

Prior to presenting Theorem 20 we need the following definition.

► **Definition 19** ([2, Proof overview of Theorem 5.2]). *A standard path decomposition is a path decomposition in which the vertices in each bag are not a subset of the vertices in the preceding or succeeding bag along the path.*

Observe that in a standard path decomposition there are no two distinct bags with the vertices of the first one being a subset of the vertices of the second. We define the *length* of a path decomposition to be the number of its bags.

► **Theorem 20.** (\star) *Let G be a connected graph with more than one vertex that admits a standard path decomposition of width pw and length ℓ . Then, for all $s \in V(G)$ it holds that $b(G, s) \geq \sqrt{\ell^{1/(\text{pw}+1)}/2\text{pw}}$.*

Observe that one can easily transform a path decomposition into a standard path decomposition of the same width; it suffices to discard any bag whose vertices are a subset of another bag's vertices. Furthermore, notice that in any path decomposition every vertex appears at least once. This, along with Theorem 20, leads to Corollary 21, which in turn, along with the approximation algorithm of Elkin and Kortsarz [14], leads to Corollary 22.

► **Corollary 21.** *Let G be a connected graph with $n > 1$ vertices of pathwidth pw . Then, for all $s \in V(G)$ it holds that*

$$b(G, s) \geq \sqrt{\frac{(n/(\text{pw} + 1))^{1/(\text{pw}+1)}}{2^{\text{pw}}}}.$$

► **Corollary 22.** *There is a polynomial-time algorithm for TELEPHONE BROADCAST achieving an approximation ratio of $\mathcal{O}(\text{pw})$, where pw denotes the pathwidth of the input graph.*

7 Conclusion

In this paper we have given an extensive investigation of the complexity of the TELEPHONE BROADCAST problem with respect to many of the most prominent structural graph parameters. Let us also mention a few concrete open questions. First, with regards to dense graph parameters, we have presented an exact FPT algorithm parameterized by the vertex deletion distance to clique, but only approximation algorithms for the more general parameters cluster vertex deletion number and clique-cover number. Can these be improved to exact FPT algorithms?

Second, with respect to treewidth, even though the problem is NP-hard for graphs of treewidth 2, we showed that an XP algorithm is possible if we seek a fast broadcast schedule, that is, $t = \mathcal{O}(\log n)$, because the problem admits an algorithm with parameter dependence $2^{\mathcal{O}(t \cdot \text{tw})}$ (Theorem 12). Can this be improved to an FPT algorithm parameterized by treewidth, under this restriction on t ? Note that treewidth is the only parameter for which this question is interesting, as for graphs of small pathwidth we have established that the broadcast time must be significantly higher than $\mathcal{O}(\log n)$ (Corollary 21).

Finally, the computational complexity of TELEPHONE BROADCAST on graph classes would be another interesting research direction. In particular, the complexity on split graphs remains open [28].

References

- 1 Akash Ambashankar and Hovhannes A. Harutyunyan. Broadcasting in stars of cliques and path-connected cliques. *Algorithms*, 18(2), 2025. doi:10.3390/a18020076.
- 2 Aida Aminian, Shahin Kamali, Seyed-Mohammad Seyed-Javadi, and Sumedha. On the complexity of telephone broadcasting: From cacti to bounded pathwidth graphs. In *52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, July 8-11, 2025, Aarhus, Denmark*, LIPIcs. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025.
- 3 Amotz Bar-Noy, Sudipto Guha, Joseph Naor, and Baruch Schieber. Message multicasting in heterogeneous networks. *SIAM J. Comput.*, 30(2):347–358, 2000. doi:10.1137/S0097539798347906.
- 4 Puspall Bhabak and Hovhannes A. Harutyunyan. Constant approximation for broadcasting in k -cycle graph. In *Algorithms and Discrete Applied Mathematics - First International Conference, CALDAM 2015. Proceedings*, volume 8959 of *Lecture Notes in Computer Science*, pages 21–32. Springer, 2015. doi:10.1007/978-3-319-14974-5_3.

- 5 Puspall Bhabak and Hovhannes A. Harutyunyan. Approximation algorithm for the broadcast time in k -path graph. *J. Interconnect. Networks*, 19(4):1950006:1–1950006:22, 2019. doi:10.1142/S0219265919500063.
- 6 Maja Cevnik and Janez Zerovnik. Broadcasting on cactus graphs. *J. Comb. Optim.*, 33(1):292–316, 2017. doi:10.1007/S10878-015-9957-8.
- 7 Jianer Chen, Iyad A. Kanj, and Ge Xia. Improved upper bounds for vertex cover. *Theor. Comput. Sci.*, 411(40-42):3736–3756, 2010. doi:10.1016/J.TCS.2010.06.026.
- 8 Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.
- 9 Peter Damaschke. A linear-time optimal broadcasting algorithm in stars of cliques. *J. Graph Algorithms Appl.*, 28(1):385–388, 2024. doi:10.7155/JGAA.V28I1.2981.
- 10 Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate texts in mathematics*. Springer, 2017. doi:10.1007/978-3-662-53622-3.
- 11 Pål Grønås Drange, Markus S. Dregi, and Pim van ’t Hof. On the computational complexity of vertex integrity and component order connectivity. *Algorithmica*, 76(4):1181–1202, 2016. doi:10.1007/S00453-016-0127-X.
- 12 Michael Elkin and Guy Kortsarz. A combinatorial logarithmic approximation algorithm for the directed telephone broadcast problem. *SIAM J. Comput.*, 35(3):672–689, 2005. doi:10.1137/S0097539704440740.
- 13 Michael Elkin and Guy Kortsarz. An approximation algorithm for the directed telephone multicast problem. *Algorithmica*, 45(4):569–583, 2006. doi:10.1007/S00453-005-1196-4.
- 14 Michael Elkin and Guy Kortsarz. Sublogarithmic approximation for telephone multicast. *J. Comput. Syst. Sci.*, 72(4):648–659, 2006. doi:10.1016/J.JCSS.2005.12.002.
- 15 Arthur M. Farley. Minimum-time line broadcast networks. *Networks*, 10(1):59–70, 1980. doi:10.1002/NET.3230100106.
- 16 Fedor V. Fomin, Pierre Fraigniaud, and Petr A. Golovach. Parameterized complexity of broadcasting in graphs. *Theor. Comput. Sci.*, 997:114508, 2024. doi:10.1016/J.TCS.2024.114508.
- 17 Pierre Fraigniaud and Emmanuel Lazard. Methods and problems of communication in usual networks. *Discret. Appl. Math.*, 53(1-3):79–133, 1994. doi:10.1016/0166-218X(94)90180-5.
- 18 M. R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- 19 Mohammad Saber Gholami, Hovhannes A. Harutyunyan, and Edward Maraachlian. Optimal broadcasting in fully connected trees. *J. Interconnect. Networks*, 23(1):2150037:1–2150037:20, 2023. doi:10.1142/S0219265921500377.
- 20 Michelangelo Grigni and David Peleg. Tight bounds on minimum broadcast networks. *SIAM J. Discret. Math.*, 4(2):207–222, 1991. doi:10.1137/0404021.
- 21 David G. Harris and N. S. Narayanaswamy. A faster algorithm for vertex cover parameterized by solution size. In *41st International Symposium on Theoretical Aspects of Computer Science, STACS 2024*, volume 289 of *LIPICs*, pages 40:1–40:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICs.STACS.2024.40.
- 22 Hovhannes A. Harutyunyan and Narek A. Hovhannisyan. Broadcasting in split graphs. In *Algorithms and Complexity - 13th International Conference, CIAC 2023, Proceedings*, volume 13898 of *Lecture Notes in Computer Science*, pages 278–292. Springer, 2023. doi:10.1007/978-3-031-30448-4_20.
- 23 Hovhannes A. Harutyunyan and Edward Maraachlian. On broadcasting in unicyclic graphs. *J. Comb. Optim.*, 16(3):307–322, 2008. doi:10.1007/S10878-008-9160-2.
- 24 Daniel Hathcock, Guy Kortsarz, and R. Ravi. The telephone k -multicast problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2024*, volume 317 of *LIPICs*, pages 21:1–21:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICs.APPROX/RANDOM.2024.21.

- 25 Sandra Mitchell 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.
- 26 Juraj Hromkovic, Ralf Klasing, Andrzej Pelc, Peter Ruzicka, and Walter Unger. *Dissemination of Information in Communication Networks - Broadcasting, Gossiping, Leader Election, and Fault-Tolerance*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2005. doi:10.1007/B137871.
- 27 Heather Hulett, Todd G. Will, and Gerhard J. Woeginger. Multigraph realizations of degree sequences: Maximization is easy, minimization is hard. *Oper. Res. Lett.*, 36(5):594–596, 2008. doi:10.1016/J.ORL.2008.05.004.
- 28 Klaus Jansen and Haiko Müller. The minimum broadcast time problem for several processor networks. *Theor. Comput. Sci.*, 147(1&2):69–85, 1995. doi:10.1016/0304-3975(94)00230-G.
- 29 Guy Kortsarz and David Peleg. Approximation algorithms for minimum-time broadcast. *SIAM J. Discret. Math.*, 8(3):401–427, 1995. doi:10.1137/S0895480193245923.
- 30 Christos H. Papadimitriou and Mihalis Yannakakis. The complexity of restricted spanning tree problems. *J. ACM*, 29(2):285–309, 1982. doi:10.1145/322307.322309.
- 31 R. Ravi. Rapid rumor ramification: Approximating the minimum broadcast time (extended abstract). In *35th Annual Symposium on Foundations of Computer Science, Santa Fe, New Mexico, USA, 20-22 November 1994*, pages 202–213. IEEE Computer Society, 1994. doi:10.1109/SFCS.1994.365693.
- 32 Christian Schindelhauer. On the inapproximability of broadcasting time. In *Approximation Algorithms for Combinatorial Optimization, Third International Workshop, APPROX 2000, Proceedings*, volume 1913 of *Lecture Notes in Computer Science*, pages 226–237. Springer, 2000. doi:10.1007/3-540-44436-X_23.
- 33 Peter J. Slater, Ernest J. Cockayne, and Stephen T. Hedetniemi. Information dissemination in trees. *SIAM J. Comput.*, 10(4):692–701, 1981. doi:10.1137/0210052.
- 34 Prafullkumar Tale. Telephone broadcast on graphs of treewidth two. *Theor. Comput. Sci.*, 1045:115282, 2025. doi:10.1016/J.TCS.2025.115282.
- 35 Jinghan Xu and Zhiyuan Li. Broadcast graph is NP-complete. In *Algorithms and Discrete Applied Mathematics - 11th International Conference, CALDAM 2025, Proceedings*, volume 15536 of *Lecture Notes in Computer Science*, pages 343–357. Springer, 2025. doi:10.1007/978-3-031-83438-7_29.