

Faster Parameterized Broadcasting

Édouard Bonnet   

CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR 5668, Lyon, France

Carl Feghali¹   

CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR 5668, Lyon, France

Manolis Vasilakis   

Université Paris Dauphine – PSL, CNRS UMR7243, LAMSADE, Paris, France

Abstract

Given a connected graph G and a source $s \in V(G)$, what is the smallest number of rounds necessary for all vertices of G to receive a message initially only held by s , where at each round every informed vertex passes the message to one of its neighbors? This problem is called TELEPHONE BROADCAST and is surprisingly hard: it remains NP-hard on cycles intersecting at a single shared vertex, in particular, graphs of pathwidth 2 with a linear feedback vertex set of size 1, as well as on graphs with treedepth at most 6 [Egami et al.; MFCS '25]. Vertex cover number, vertex integrity, and distance to clique are among the few parameters for which TELEPHONE BROADCAST is fixed-parameter tractable. There is a $2^{\mathcal{O}(\text{vc}^3)} n^{\mathcal{O}(1)}$ -time algorithm parameterized by vertex cover number vc [Fomin, Fraigniaud, Golovach; TCS '24], a double-exponential algorithm parameterized by vertex integrity vi , and a $2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$ -time algorithm parameterized by distance to clique k [Egami et al.; MFCS '25].

In this paper, we give improved parameterized algorithms for TELEPHONE BROADCAST with running times $2^{\mathcal{O}(\text{vc} \log \text{vc})} n^{\mathcal{O}(1)}$, $2^{\mathcal{O}(\text{vi}^2 \log \text{vi})} n^{\mathcal{O}(1)}$, and $2^{\mathcal{O}(k \log k)} n^{\mathcal{O}(1)}$. The main ingredient that makes our algorithms faster is a Turing reduction to edge-weighted b -MATCHING.

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

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

Funding *Édouard Bonnet*: Supported by the ANR projects TWIN-WIDTH (ANR-21-CE48-0014-01). *Manolis Vasilakis*: Supported by the ANR project ANR-21-CE48-0022 (S-EX-AP-PE-AL) and the GDR ROD Mobility Scholarship.

1 Introduction

Broadcasting is one of the most basic primitives in communication networks: a message initially held by one source has to reach every vertex of the network. In the classical *telephone model* [26], the network is represented by a connected graph G , time is divided into synchronous rounds, and in each round every informed vertex may transmit the message to at most one uninformed neighbor. For a source $s \in V(G)$, the *broadcast time* $b(G, s)$ is the minimum number of rounds needed to inform all vertices. If $|V(G)| = n$, then $\lceil \log_2 n \rceil \leq b(G, s) \leq n - 1$, and both bounds are tight: complete graphs realize the lower bound, while a path whose source is an endpoint realizes the upper bound. Equivalently, a broadcast protocol can be represented by a spanning tree rooted at s together with an ordering of the children of every vertex. TELEPHONE BROADCAST then asks, given a connected graph G , a source s , and an integer t , whether $b(G, s) \leq t$.

¹ We are deeply saddened to report that our colleague and friend, Carl Feghali, recently passed away. We greatly enjoyed working on this project with him.

The problem is old, natural, and algorithmically challenging. Slater, Cockayne, and Hedetniemi [35] initiated the systematic study of the problem, proving that it is NP-complete in general but polynomial-time solvable on trees. The surrounding literature is extensive; see, for instance, the survey of Hedetniemi, Hedetniemi, and Liestman [26], the survey of Fraigniaud and Lazard [19], and the book of Hromkovic et al. [27]. The computational difficulty already appears on quite restricted graph classes: Jansen and Müller [28] proved NP-hardness on chordal graphs and grid graphs, among others.

The parameterized-complexity study of TELEPHONE BROADCAST was initiated by Fomin, Fraigniaud, and Golovach [18]. They gave a $3^n n^{\mathcal{O}(1)}$ -time exact algorithm and proved fixed-parameter tractability for feedback edge set number, vertex cover number, and the below-guarantee parameter $n - t$. Since a protocol of length t can inform at most 2^t vertices, the exact algorithm also gives an FPT algorithm parameterized by t , but with double-exponential dependence on t . This is essentially the right order of magnitude under the Exponential-Time Hypothesis (ETH), by a consequence of the restricted spanning-tree hardness of Papadimitriou and Yannakakis [31], as discussed by Tale [36].

The recent structural picture is quite sharp on the negative side. Tale proved that TELEPHONE BROADCAST remains NP-complete on graphs of feedback vertex set number 1, and hence on graphs of treewidth 2 [36]; in fact, the same work gives hardness for graphs at vertex-deletion distance 2 from a linear forest. Aminian, Kamali, Seyed-Javadi, and Sumedha [2] proved NP-completeness on cactus graphs of pathwidth 2, while also obtaining approximation results for cactus graphs and bounded-pathwidth graphs. Egami et al. [13] conducted a thorough study of the problem under various parameterizations and, apart from improving (and significantly simplifying) the hardness result of [2], pushed the hardness further to graphs of bounded treedepth, and complemented this with positive results for vertex integrity, distance to clique, the combined parameter $t + \text{tw}$, and FPT approximation for dense parameters such as clique cover and cluster vertex deletion. Bringolf, Harutyunyan, Kamali, and Seyed-Javadi [6] recently continued this line by studying hardness and approximation for structured graph classes.

Our Contribution. We improve the parameter dependence of the known positive results for vertex cover number, vertex integrity, and distance to clique. For vertex cover number vc , in Theorem 3 we obtain an algorithm solving TELEPHONE BROADCAST in time $\text{vc}^{\mathcal{O}(\text{vc})} n^{\mathcal{O}(1)}$, improving over the $2^{\mathcal{O}(\text{vc}^3)} n^{\mathcal{O}(1)}$ -time algorithm of Fomin, Fraigniaud, and Golovach [18]. For vertex integrity, our algorithm runs in $\text{vi}^{\mathcal{O}(\text{vi}^2)} n^{\mathcal{O}(1)}$ time (Theorem 6), improving the double-exponential dependence of the algorithm of Egami et al. [13]. Finally, we consider graphs whose vertex-deletion distance to clique is k , and in Theorem 13 we obtain a $k^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$ -time algorithm, improving the previous $2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$ -time bound of Egami et al. [13].

On a high level, all three of our algorithms use the same organizing idea. We first guess a small *template* of the broadcast protocol, recording only the transmissions that interact nontrivially with the relevant modulator. For vertex cover, the template contains the cover and the few independent-set vertices that relay the message back into it. For vertex integrity, the template contracts the few components outside the modulator that are entered several times or that later send the message back. For distance to clique, the template keeps only the clique vertices that interact with the modulator. All other vertices or components are left anonymous. Once the template is fixed, the remaining task is to realize the abstract slots and assign all leftover leaves to available transmissions; we formulate this as an edge-weighted b -MATCHING instance, which is polynomial-time solvable [34]. We note that a similar approach already appears in [18] (and subsequently in [13]). However, a refinement

of it leads to significantly faster and simpler algorithms.

Further Related Work. Several classical lines of research surround TELEPHONE BROADCAST. One asks for exact algorithms on special graph classes. Besides trees [35], polynomial-time algorithms are known for subclasses of cactus graphs [7], unicyclic graphs [24], fully connected trees [20], and generalized windmill graphs, equivalently, graphs of twin-cover number 1 [1, 10]. Another classical direction is extremal: one asks how sparse a graph can be while still achieving the information-theoretic optimum $\lceil \log_2 n \rceil$; Grigni and Peleg [21] gave tight bounds of this flavor.

The problem is also well studied from the viewpoint of approximation. Ravi gave an $\mathcal{O}(\log^2 n / \log \log n)$ -approximation [32], Kortsarz and Peleg [30] gave an algorithm producing a protocol of length at most $2b(G, s) + \mathcal{O}(\sqrt{n})$, and Bar-Noy, Guha, Naor, and Schieber [3] obtained an $\mathcal{O}(\log n)$ -approximation. Elkin and Kortsarz [14] gave a combinatorial logarithmic approximation for the directed version, and later obtained the $\mathcal{O}(\log n / \log \log n)$ guarantee for telephone multicast [16]. Better approximation ratios are known for specific graph classes [4, 5, 23, 30], and inapproximability results are known as well [14, 33].

Finally, many variants of the telephone model have been studied. These include line- and path-communication models [17, 30] and multicast versions, where only a specified terminal set must be informed. Most of the approximation algorithms above extend to multicast, and directed or otherwise modified multicast variants have also been considered [15, 25].

2 Preliminaries

2.1 Basic notation and broadcasting

Basic notation. We use standard graph notation [11] and assume familiarity with the basic notions of parameterized complexity [9]. All graphs are finite, simple, and undirected, unless explicitly stated otherwise. For integers $x, y \in \mathbb{Z}$, let $[x, y] := \{z \in \mathbb{Z} \mid x \leq z \leq y\}$, and, for a positive integer x , let $[x] := [1, x]$. For a graph $G = (V, E)$ and a set $S \subseteq V$, we write $G[S]$ for the subgraph induced by S , and $G - S$ for $G[V \setminus S]$. We denote the set of connected components of G by $\text{cc}(G)$.

Telephone broadcast. Let G be a connected graph and let $s \in V(G)$ be the source. A *broadcast protocol* of length t is a spanning tree T of G (called its *broadcast tree*), rooted at s , together with a time-stamp function $\lambda: E(T) \rightarrow [t]$. Orient every edge of T away from s . Say, uv is oriented from the parent u to the child v . Then $\lambda(uv)$ is the round in which u informs v . Let $\tau_\lambda(s) = 0$, and let $\tau_\lambda(v)$ be the label of the unique edge of T entering v for every $v \neq s$.

The pair (T, λ) is *feasible* if $\tau_\lambda(u) < \lambda(uv)$ for every edge oriented from u to v , and if the outgoing edges of each vertex have pairwise distinct labels. These are exactly the telephone constraints: a vertex transmits only after being informed, and to at most one neighbor in each round. Conversely, every feasible pair (T, λ) defines a broadcast protocol. The *broadcast time* $b(G, s)$ is the minimum length of a broadcast protocol for source s in G .

2.2 Templates


Our algorithms repeatedly guess a bounded part of a broadcast protocol. We use one common language for these guesses. A *template* records the named vertices, a bounded number of

abstract placeholders, and the transmissions between them. The remaining vertices or components are selected later, by a b -MATCHING instance.

► **Definition 1 (Template).** A template is a tuple $\mathcal{T} = (H, s, S, Q, \lambda)$, where S and Q are disjoint finite sets, $s \in S$, $H = (V_H, A)$ is a loopless directed multigraph with $S \cup Q \subseteq V_H$, and $\lambda: A \rightarrow \mathbb{N}$. The vertices of S are called named vertices; the vertices of Q are called slots. The vertices of $V_H \setminus (S \cup Q)$, if any, are called auxiliary vertices and have no incoming arcs. Parallel arcs in H are allowed.

When the remaining data are clear from the context, we sometimes refer to (H, λ) itself as the template.

Meaning of the labels. An arc $e = (u, v) \in A$ with $\lambda(e) = \rho$ represents a transmission in round ρ from the object represented by u to the object represented by v , unless u is auxiliary. Arcs leaving auxiliary vertices are interpreted by the algorithm that introduced them. The labels determine the first labeled entry into every object that has an incoming arc. Namely, for an object $z \in S \cup Q$ we write $\tau_\lambda(z) := \min\{\lambda(e) \mid e \in A \text{ enters } z\}$, and we set $\tau_\lambda(s) = 0$ for the source. The consistency conditions below ensure that every vertex of $S \setminus \{s\}$ and every slot $q \in Q$ has an incoming arc, so τ_λ gives its reception time. Auxiliary vertices are only bookkeeping devices; they are not realized by vertices or components of the input graph, and no reception time is associated with them. In the distance-to-clique algorithm we use one such vertex ϕ to denote a transmission supplied later from inside the clique.



■ **Figure 1** An example of a template.

Consistency. When a template (H, s, S, Q, λ) is checked against an input graph G , the named vertices are actual vertices of G , so $S \subseteq V(G)$, and $s \in S$ is the source. We say that the template is *consistent* with respect to G if it satisfies the following common local checks:

- s has no incoming arc;
- each $u \in S \setminus \{s\}$ has exactly one incoming arc;
- all outgoing arcs of $u \in S$ have distinct labels, and each label is greater than $\tau_\lambda(u)$;
- every arc with both endpoints in S is an edge of G ;
- every slot has at least one incoming arc;
- for every $q \in Q$, each outgoing arc e from q satisfies $\tau_\lambda(q) < \lambda(e)$.

Thus every represented transmission leaving an object is labeled later than the first labeled entry into that object. Each algorithm may impose additional structural requirements, for example that labels belong to particular sets of rounds.

Slots and realizations. Slots are placeholders. In this paper, each slot is realized either by one unnamed vertex or by one connected component of $G - S$; the algorithm at hand specifies which type is allowed. Thus a slot need not have a unique incoming arc, and a template need not be a tree or even be acyclic. Arcs incident with a slot record only boundary transmissions. The internal behavior of the realized object and the choice of the actual endpoint inside it are checked when the slot is realized.

Formally, a realization of $\mathcal{T}$ maps each slot to a concrete object in $G - S$, with different slots mapped to disjoint objects. It is valid if all represented boundary transmissions not involving auxiliary vertices can be implemented along edges of G , and if every realized slot admits a local protocol respecting the time-stamps of its incident arcs. Auxiliary vertices, if present, are not realized.

Extensions. An *extension* of a realized template adds the transmissions not represented in $\mathcal{T}$, e.g., transmissions to independent-set vertices that do not inform any vertices in the vertex cover, or transmissions inside a clique. A template is *extendable* if some realization and some extension form a valid broadcast protocol within the required number of rounds. Our algorithms enumerate a bounded family of templates and test extendability by a polynomial-time matching subroutine.

Use in the algorithms. For vertex cover number, the template is a rooted tree: S is the vertex cover and the slots are the independent-set vertices that inform vertices of the cover, each realized by a distinct independent-set vertex. For vertex integrity, S is the modulator and the slots are the components of $G - S$ that either inform vertices of the modulator or receive the message from the modulator multiple times, each realized by a distinct connected component; here the template is a directed multigraph and may have several arcs entering the same slot. For distance to clique, S is the modulator to the clique and the slots are the clique vertices informing the modulator. The unique auxiliary vertex ϕ records slots whose first entry is supplied later from inside the clique.

3 Vertex Cover Number

We start with the parameterization by vertex cover number, denoted by vc . Fomin, Fraignaud, and Golovach [18] proved that TELEPHONE BROADCAST is fixed-parameter tractable parameterized by vc , with running time $2^{\mathcal{O}(\text{vc}^3)} n^{\mathcal{O}(1)}$. We improve this dependence to $\text{vc}^{\mathcal{O}(\text{vc})} n^{\mathcal{O}(1)}$. Our algorithm keeps the main structural insight of their proof: there is an optimal protocol in which the vertex cover is informed very early. We restate this lemma here and later use it to guess only the part of the broadcast tree that is relevant to the vertex cover.

► **Lemma 2** ([18, Lemma 4]). *For every connected graph G , vertex $s \in V(G)$, and vertex cover S of G , there is an optimal broadcast protocol for (G, s) such that every vertex of S receives the message by round $2|S| - 1$.*

We can now present the improved algorithm parameterized by vertex cover number.


► **Theorem 3.** TELEPHONE BROADCAST can be solved in time $\text{vc}^{\mathcal{O}(\text{vc})} n^{\mathcal{O}(1)}$, where vc is the vertex cover number of the input graph.

Proof. Let (G, s, t) be an instance of TELEPHONE BROADCAST, where $G = (V, E)$ is connected and $|V| = n$. Assume that $t < n$, since otherwise the answer is trivially yes. We

first compute in time $2^{\mathcal{O}(\text{vc})}n^{\mathcal{O}(1)}$ a vertex cover S_0 of $G - s$ of size at most vc [8, 22]. We set $S := S_0 \cup \{s\}$, $k := |S| \leq \text{vc} + 1$, and let $I := V \setminus S$ and $L := 2k - 1$. Note that I is an independent set of G , since S is a vertex cover of G . By Lemma 2, there is an optimal broadcast protocol for G with source s such that every vertex of S receives the message by round L . In the following, we search for such an optimal protocol.

Fix such a protocol, and let T be its broadcast tree. We first describe the bounded part of T that we guess. A vertex of I is *marked* if it has at least one child in S . Since I is independent, every marked vertex has its parent in S and all its children in S . Moreover, every vertex of $S \setminus \{s\}$ has a unique parent. Thus, if we assign to every marked vertex one of its children in S , we obtain an injection from the set of marked vertices to $S \setminus \{s\}$. Hence there are at most $k - 1$ marked vertices.

The vertices of S together with the marked vertices form a rooted subtree T^* of T : on the path from any vertex of S to the root, every vertex of I that appears has a child in S , and is therefore marked. All vertices of this subtree receive the message by round L : this is clear for the vertices of S , and each marked vertex receives the message before informing one of its children in S .



■ **Figure 2** Left: a vertex-cover template with $S = \{s, u_1, u_2, u_3, u_4, u_5\}$ and marked slots $Q = \{q_1, q_2\}$. Right: a broadcast tree extending the template, where q_1, q_2 are realized by $i_1, i_2 \in I$, and the unlabeled leaves are unmarked vertices of I .

We enumerate all possible *templates* for such a subtree; see Figure 2 for an illustration. For each $w \in \{0, \dots, k - 1\}$, we introduce a set Q of w *marked slots*. A marked slot is a placeholder for a marked vertex of I in the unknown broadcast tree; it will later be realized by a distinct vertex of I . We enumerate all trees H on vertex set $S \cup Q$, rooted at s , such that every marked slot of Q has its parent in S and a non-empty set of children contained in S .


We orient H away from s and enumerate a label $\lambda(e) \in [L]$ for every arc e of H . The resulting labeled tree, with named vertices S and slots Q , is a template $\mathcal{T} = (H, s, S, Q, \lambda)$ in the sense of Definition 1. We discard the guess unless the consistency conditions of Section 2.2 hold, with the additional constraint that for every $q \in Q$, all its outgoing arcs receive a distinct label.

▷ **Claim 4.** The number of templates considered is $k^{\mathcal{O}(k)}$.

Proof. Fix $w \in \{0, \dots, k - 1\}$ and let $m = k + w \leq 2k$. There are at most $m^m \leq (2k)^{2k}$ trees on the vertex set $S \cup Q$, and the root s determines their orientation. For each such

tree, there are at most $L^{m-1} \leq (2k)^{2k}$ choices of the labels, because $L = 2k - 1 \leq 2k$. The value of w has only k choices, and the consistency checks only discard guesses. Hence the total number of templates is at most $k \cdot (2k)^{4k} = k^{\mathcal{O}(k)}$. $\triangleleft$

Fix a consistent template (H, λ) . For notational convenience, write $\tau = \tau_\lambda$. For $q \in Q$, let $R(q)$ be the set of vertices of S adjacent to q in the underlying tree H . Thus $R(q)$ consists of the parent and the children of q in the template. An actual vertex of I can realize q only if it is adjacent in G to every vertex of $R(q)$. Let $I_q := \{v \in I \mid R(q) \subseteq N_G(v)\}$ denote the set of such vertices. Figure 3 illustrates in particular that, with the notation of the figures, $i_1 \in I_{q_1}$.



■ **Figure 3** A graph admitting the right-hand side of Figure 2 as a broadcast tree, with vertex cover S and independent set I . Vertex $i_1 \in I$ realizes the marked slot $q_1 \in Q$, which is permitted since $R(q_1) = \{s, u_3, u_4\}$ (in blue) is a subset of $N(i_1) = \{s, u_2, u_3, u_4\}$ (with incident thicker edges).

Recall that in b -MATCHING, given an edge-weighted graph and a map β from its vertex set to the set of positive integers, we seek a maximum-weight set M of edges such that each vertex u is incident with at most $\beta(u)$ edges of M (regardless of their weight). We call $\beta(u)$ the capacity of u . We create such an instance as follows. Its vertex set comprises $S \cup I$, and one additional vertex w_q for every marked slot $q \in Q$. The capacities are

$$\beta(v) = 1 \quad \text{for every } v \in I, \quad \beta(w_q) = 1 \quad \text{for every } q \in Q,$$


and, for every $u \in S$,

$$\text{children}(u) = \{x \in V(H) \mid x \text{ is a child of } u \text{ in } H\}, \quad \beta(u) = t - \tau(u) - |\text{children}(u)|.$$

If $\beta(u) < 0$ for some $u \in S$, we discard the template. If $\beta(u) = 0$ for some $u \in S$, then we delete the vertex u from the instance (as it has no capacity). We add an edge of weight 2 between w_q and every vertex of I_q . We also add an edge of weight 1 between $u \in S$ and $v \in I$ whenever $uv \in E(G)$. We say that the template is *extendable* if the resulting b -MATCHING instance has a feasible b -matching of weight at least $|I| + |Q|$. This can be checked in polynomial time, since b -MATCHING is polynomial-time solvable [34].

▷ **Claim 5.** $b(G, s) \leq t$ if and only if at least one consistent template is extendable.

Proof. Suppose first that $b(G, s) \leq t$. By the discussion above, there is a protocol in at most t rounds in which every vertex of S receives the message by round L . Take its template (H, λ) , with labels given by the reception times in the protocol. We construct a feasible b -matching M . For every marked slot $q \in Q$, let $v_q \in I$ be the actual vertex realizing it, and put the edge $w_q v_q$ into M . This edge exists because $v_q \in I_q$. Every vertex of I that is not marked is a leaf of the broadcast tree, because I is independent and all its potential children lie in S . If such a leaf $v \in I$ has parent $u \in S$, put the edge uv into M .



■ **Figure 4** The edge-weighted b -MATCHING instance associated with the template and the graph of Figures 2 and 3. Capacities are in red and edge weights are in black. Implicit capacities and weights have value 1. The bold edges represent the b -matching for the broadcast tree of Figure 2.

By construction, M contains (at most) one edge incident to each vertex of I and each vertex w_q . So their capacity is respected. After receiving the message at round $\tau(u)$, the vertex $u \in S$ has $t - \tau(u)$ transmissions available before the deadline. The template already accounts for exactly $|\text{children}(u)|$ of them. Hence the number of additional leaf children of u in I is at most $\beta(u)$, and the capacity of u is respected. Finally, the weight of M is

$$2|Q| + (|I| - |Q|) = |I| + |Q|,$$

so the template is extendable.

Conversely, suppose that a consistent template is extendable. Let M be a feasible b -matching of weight at least $|I| + |Q|$. Since every vertex of I has capacity 1, each vertex of I is incident with at most one edge of M . Such an edge has weight 2 only if it is incident with some w_q , and otherwise has weight 1. Let m_2 be the number of weight-2 edges in M . Then $m_2 \leq |Q|$, because each w_q has capacity 1. The remaining matched vertices of I contribute weight at most 1 each, and there are at most $|I| - m_2$ of them. Therefore every feasible b -matching has weight at most

$$2m_2 + (|I| - m_2) = |I| + m_2 \leq |I| + |Q|.$$

It follows that M has weight exactly $|I| + |Q|$, and equality in the above upper bound implies that $m_2 = |Q|$, every vertex of I is saturated, and every vertex w_q is saturated.

We construct a broadcast protocol as follows. Start with the guessed template. For each edge $w_q v \in M$, replace $q \in Q$ by the actual vertex v . This is possible because $v \in I_q$, so v is adjacent in G to the parent and all children prescribed for q in the template. The capacity-1 constraints on the vertices of I ensure that these vertices are distinct and are not used again. Every remaining vertex of I is saturated by a unique edge $uv \in M$ with $u \in S$; make each such vertex v a leaf child of u . This gives a spanning tree of G rooted at s .

We schedule every template arc from a parent to a vertex x at round $\tau(x)$. For each $u \in S$, the number of added leaf children assigned to u is at most $\beta(u)$, and there are exactly $\beta(u)$ rounds among $\tau(u) + 1, \dots, t$ not already occupied by template transmissions from u . Schedule these leaf transmissions arbitrarily in those free rounds. All transmissions are along

edges of G , no vertex transmits before receiving the message, and no vertex transmits to two vertices in the same round. Therefore every vertex receives the message by round t . $\triangleleft$

We enumerate $k^{\mathcal{O}(k)}$ templates, and for each of them solve a polynomial-size b -MATCHING instance. Since $k \leq \text{vc} + 1$, the total running time is $\text{vc}^{\mathcal{O}(\text{vc})} n^{\mathcal{O}(1)}$. $\blacktriangleleft$

4 Vertex Integrity

In this section we consider the parameterization by *vertex integrity*, denoted by vi . The vertex integrity of a graph is the smallest value k such that there exists a vertex set S whose deletion leaves only connected components of size at most $k - |S|$. Vertex integrity generalizes vertex cover number in a straightforward way, since it is upper-bounded by the vertex cover number plus one.

Egami et al. [13] proved that TELEPHONE BROADCAST is fixed-parameter tractable parameterized by vertex integrity. The dependence on the parameter in their proof is double-exponential;² we improve it to $\text{vi}^{\mathcal{O}(\text{vi}^2)} n^{\mathcal{O}(1)}$.

At a high level, the algorithm guesses only the part of the protocol where a component of $G - S$ interacts nontrivially with the modulator S . There are few components that send the message back to S , and these components, together with the vertices of S , can be assumed to receive the message only at the very start or at the very end of the protocol [13, Lemma 8]. Almost every other component receives the message from S only once; the only exceptions occur near the end of the protocol [13, Lemma 7]. The template records exactly these bounded interactions, while all remaining components are assigned later by a MAXIMUM-WEIGHT MATCHING instance (equivalently, an edge-weighted b -MATCHING instance with unit capacities).

► **Theorem 6.** TELEPHONE BROADCAST can be solved in time $\text{vi}^{\mathcal{O}(\text{vi}^2)} n^{\mathcal{O}(1)}$, where vi is the vertex integrity of the input graph.

Proof. Let (G, s, t) be an instance of TELEPHONE BROADCAST, where $G = (V, E)$ is connected and $|V| = n$. Assume that $t < n$, since otherwise the answer is trivially yes. Using the algorithm of Drange et al. [12], we compute in time $\text{vi}^{\mathcal{O}(\text{vi})} n^{\mathcal{O}(1)}$ a set $S_0 \subseteq V$ such that every component of $G - (S_0 \cup \{s\})$ has at most $\text{vi} - |S_0|$ vertices. Let $S := S_0 \cup \{s\}$, $\mathcal{D} := \text{cc}(G - S)$, $d := \max_{D \in \mathcal{D}} |V(D)|$, and $k := |S| + d \leq \text{vi} + 1$.

Notation. For a protocol with broadcast tree T , an *entry* into a component $D \in \mathcal{D}$ is an edge of T directed from S to D . A component is a *relay component* if some edge of T is directed from it to S , and otherwise it is a *leaf component*. A leaf component is *exceptional* if it has at least two entries from S . Relay and exceptional components are called *special*.

Let $\Lambda := [1, \min\{t, (2k - 3)k\}] \cup [\max\{1, t - 2k + 1\}, t]$ and $J := [\max\{1, t - k + 2\}, t]$. Consequently, we have that $|\Lambda| = \mathcal{O}(k^2)$ and $|J| \leq k - 1$. Using the results of [13] we obtain the following.

▷ **Claim 7.** If $b(G, s) \leq t$, then there is a broadcast protocol P of length at most t , with broadcast tree T , such that the following hold.

² Their algorithm considers two cases. In the first one, the stated double-exponential dependence can already be improved to $2^{\mathcal{O}(\text{vi}^3)} n^{\mathcal{O}(1)}$ by using their Theorem 12 as a subroutine. The second case, however, still takes time $2^{\mathcal{O}(\text{vi}^5)} n^{\mathcal{O}(1)}$.

1. Every vertex of $S \setminus \{s\}$ and every vertex in a relay component of P is informed in a round of Λ .
2. Every exceptional component of P has all its entries in rounds of J .

Proof. Let $\ell = b(G, s) \leq t$. Following the notation of [13], we say that a protocol is *S-lazy* if, among all optimal broadcast protocols, it minimizes the total number of message forwardings from S . By [13, Lemma 8], there is an *S-lazy* optimal protocol P_0 of length ℓ in which every vertex of $S \setminus \{s\}$ and every vertex in a relay component is informed either during the first $(2k - 3)k$ rounds or during the last $2k$ rounds. Moreover, [13, Lemma 7] implies that every leaf component of P_0 with at least two entries from S has all its entries after round $\ell - k + 1$.

To obtain P , we shift the time-stamps of the last $2k$ rounds of P_0 to occupy rounds $[t - 2k + 1, t]$ instead, leaving all earlier time-stamps unchanged. This only delays a suffix of the protocol. Hence feasibility is preserved: earlier and shifted time-stamps remain disjoint, and every transmission delayed in this way still occurs after its sender is informed. The resulting protocol has length at most t .

The first $(2k - 3)k$ rounds of P_0 either remain unchanged or are shifted into the last $2k$ rounds with respect to t , while the last $2k$ rounds of P_0 are shifted into the last $2k$ rounds with respect to t . This proves the first item. Since the last $k - 1$ rounds of P_0 are contained in the last $2k$ ones, every entry into a leaf component with at least two entries is shifted into J as well, proving the second item. $\triangleleft$

$\triangleright$ **Claim 8.** Let P be a protocol with broadcast tree T satisfying the two properties of Claim 7. Then P has at most $|S| - 1$ relay components and at most $|S|(k - 1)/2$ exceptional components.

Proof. Each relay component informs at least one vertex of $S \setminus \{s\}$, and distinct relay components inform distinct such vertices, since every vertex has a unique parent in T . Hence there are at most $|S| - 1$ relay components.

By the assumed second property, every entry into an exceptional component occurs in a round of J . Since T is a tree, these entries inform distinct vertices of the component. Each exceptional component has at least two entries, while the vertices of S make at most $|S||J| \leq |S|(k - 1)$ transmissions in these rounds. Thus there are at most $|S|(k - 1)/2$ exceptional components. $\triangleleft$

Thus only $\mathcal{O}(k^2)$ special components need to be represented explicitly in a broadcast protocol satisfying the two properties of Claim 7. Every other component is a leaf component entered exactly once from S .


Templates. We now describe the bounded object that we guess, in the same spirit as the template in the proof of Theorem 3. The vertices of S are kept explicitly. Each special component is contracted to one formal vertex, called a *slot*. A slot is a *relay slot* or an *exceptional slot*, according to the type of the component that it represents.

Formally, this is a template $\mathcal{T} = (H, s, S, Q, \lambda)$ in the sense of Definition 1, where $H = (S \cup Q, A)$. We write $Q_{\text{rel}} \subseteq Q$ for the relay slots and $Q_{\text{exc}} = Q \setminus Q_{\text{rel}}$ for the exceptional slots. The arcs in A record the transmissions involving S and the special components, and are only of the following forms, where we additionally record the set of allowed labels for each type of arc:

$$S \xrightarrow{\Lambda} S, \quad S \xrightarrow{\Lambda} Q_{\text{rel}}, \quad Q_{\text{rel}} \xrightarrow{\Lambda} S, \quad S \xrightarrow{J} Q_{\text{exc}}.$$

An arc incident with a slot records only the endpoint in S and the round of the transmission; the endpoint inside the represented component is chosen later, when the slot is realized.

Unlike the vertex-cover template, this guessed object need not be a tree. A special component may receive several entries from S , and a relay component may also send the message back to S . Figure 5 illustrates the guessed object and how it is realized by actual components of $G - S$.



■ **Figure 5** Left: a vertex-integrity template with $S = \{s, u_1, u_2, u_3\}$ and slots $Q = \{q_{\text{rel}}, q_{\text{exc}}\}$. Right: a broadcast tree extending the template, where $q_{\text{rel}}, q_{\text{exc}}$ are realized by $D_1, D_2 \in \mathcal{D}$, and $D_3, D_4 \in \mathcal{D}$ are leaf components of the protocol.

We enumerate templates as follows. We use at most $|S| - 1$ relay slots and at most $|S|(k - 1)/2$ exceptional slots, as guaranteed by Claim 8. For these slots we enumerate arcs of the allowed forms above, with labels from the relevant sets of rounds. The next claim shows that it is enough to consider guesses with $\mathcal{O}(k^2)$ arcs. Here we call a (vertex-integrity) template *consistent* if it satisfies the common consistency checks of Section 2.2 and the following additional conditions:

- every relay slot has at least one incoming arc from S and at least one outgoing arc to S ;
- every exceptional slot has at least two incoming arcs from S ;
- no slot has more than d incoming arcs.³

▷ **Claim 9.** The number of templates considered is $k^{\mathcal{O}(k^2)}$.

Proof. The number of slots is $\mathcal{O}(k^2)$, and hence every template has $\mathcal{O}(k^2)$ vertices. In a consistent template, relay slots contribute at most $(|S| - 1)d$ incoming arcs from S and at most $|S| - 1$ outgoing arcs to S . Exceptional slots contribute at most $|S||J| \leq |S|(k - 1)$ incoming arcs, because each vertex of S transmits at most once in each round of J . Finally, there are at most $|S| - 1$ arcs with both endpoints in S . Thus the number of arcs is $\mathcal{O}(k^2)$.

For each arc we have $\mathcal{O}(k^2)$ choices for its tail, $\mathcal{O}(k^2)$ choices for its head, and $\mathcal{O}(k^2)$ choices for its label. Therefore the total number of guesses is $k^{\mathcal{O}(k^2)}$, while the consistency checks only discard guesses. ◁

Fix a consistent template. A component $D \in \mathcal{D}$ *realizes* a slot $q \in Q$ if each arc incident with q can be implemented, with its prescribed direction and time-stamp, by choosing an

³ Recall that d is the largest number of vertices among the connected components of $G - S$.

endpoint in D adjacent to the corresponding vertex of S , and by adding transmissions inside D , so that every vertex of D receives the message exactly once by round t , and every vertex of D transmits only after it is informed and at most once per round. For a relay slot we additionally require that all vertices of D receive in rounds of Λ . No further restriction is needed for an exceptional slot: all its entries occur in J , so every valid realization completing by round t receives all vertices of D in rounds of J . Since a slot has $\mathcal{O}(k)$ incident arcs and $|V(D)| \leq d \leq k$, whether D realizes it can be tested in $k^{\mathcal{O}(k)}$ time by exhaustively guessing the corresponding local protocol.


For every component $D \in \mathcal{D}$ and every $v \in V(D)$, compute $b(D, v)$. Since $|V(D)| \leq d \leq k$, all these values can be computed in total time $2^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$ by the $3^n n^{\mathcal{O}(1)}$ -time algorithm of Fomin et al. [18].

Extending a template. Fix a consistent template (H, λ) , and, for notational convenience, write $\tau = \tau_\lambda$. Let Q be its set of slots. The components assigned to its slots account for every interaction with S that is not a single entry into a leaf component. Thus, as for the unmarked independent-set vertices in the proof of Theorem 3, every remaining component only has to be assigned one unused transmission from S .

We construct a MAXIMUM-WEIGHT MATCHING instance. Its vertex set contains one vertex v_D for every component $D \in \mathcal{D}$ and one vertex w_q for every slot $q \in Q$. We add an edge of weight 2 between v_D and w_q exactly when D realizes q . For every $u \in S$ and every round $\rho \in \{\tau(u) + 1, \dots, t\}$ in which no arc of H leaves u with label ρ , create a vertex $x_{u,\rho}$. We add an edge of weight 1 between v_D and $x_{u,\rho}$ if there is a vertex $v \in N_G(u) \cap V(D)$ such that $\rho + b(D, v) \leq t$. The template is *extendable* if this instance has a matching of weight at least

$$|\mathcal{D}| + |Q|.$$

The weight-2 edges pay for covering both a component and a slot, while the weight-1 edges pay only for covering a component by an ordinary one-entry broadcast from S . The instance has polynomial size, since $t < n$, and MAXIMUM-WEIGHT MATCHING is polynomial-time solvable [29, Section 11].



■ **Figure 6** Part of a MAXIMUM-WEIGHT MATCHING instance associated with the template of Figure 5. Edge weights are in black; implicit weights have value 1. Bold edges highlight a solution that corresponds to the realization of Figure 5.

▷ **Claim 10.** If $b(G, s) \leq t$, then an extendable template exists.

Proof. Suppose that $b(G, s) \leq t$ and take the protocol P satisfying the two properties of Claim 7. Introduce a relay slot for every relay component and an exceptional slot for every exceptional component, and form H by recording precisely the transmissions incident with S in the restriction of P to these components and S . The label ranges follow from Claim 7, and the counting arguments above ensure that this template is enumerated. Each special component realizes its corresponding slot. Add the corresponding weight-2 edge to a matching M .

Every component not represented by a slot is a leaf component entered exactly once. If such a component D is entered from $u \in S$ through $v \in V(D)$ at round ρ , then the restriction of P to D implies $\rho + b(D, v) \leq t$. Add the edge between v_D and $x_{u,\rho}$ to M . The chosen edges are a matching: each component and each slot is used once, and two ordinary edges incident with the same $x_{u,\rho}$ would correspond to two transmissions by u in round ρ in P . Hence M is feasible and has weight $2|Q| + (|\mathcal{D}| - |Q|) = |\mathcal{D}| + |Q|$. ◁

▷ **Claim 11.** If an extendable template exists, then $b(G, s) \leq t$.

Proof. Let a consistent template and a matching M of weight at least $|\mathcal{D}| + |Q|$ be given. Let m_2 be the number of weight-2 edges in M . Since each such edge is incident with a distinct slot, $m_2 \leq |Q|$. Since M is a matching, the number of weight-1 edges is at most $|\mathcal{D}| - m_2$. Therefore

$$w(M) \leq 2m_2 + |\mathcal{D}| - m_2 = |\mathcal{D}| + m_2 \leq |\mathcal{D}| + |Q|.$$

Equality must hold throughout. Thus every component is matched, every slot is matched by a weight-2 edge, and every remaining component is matched to a unique vertex $x_{u,\rho}$.

Realize each slot by its matched component, using a witnessing local protocol for that realization. For every remaining component D matched to $x_{u,\rho}$, choose a neighbor $v \in N_G(u) \cap V(D)$ such that $\rho + b(D, v) \leq t$, let u inform v at round ρ , and run an optimal broadcast inside D from v . The inequality defining the edge to $x_{u,\rho}$ guarantees completion by round t .

All transmissions are along edges of G : this follows from consistency for arcs inside S , from the realizations for slot components, and from the definition of the edges to the vertices $x_{u,\rho}$. No vertex of S transmits before it is informed or twice in the same round: consistency handles the template arcs, while each vertex $x_{u,\rho}$ represents a free round of u and is incident with at most one edge of the matching M . The local protocols give the same guarantees inside the components, and different components are disjoint.

Finally, every vertex is informed. Indeed, every vertex of $S \setminus \{s\}$ has the unique incoming transmission prescribed by the template, every slot is realized by a matched component, and every other component receives one ordinary entry from S . Following incoming transmissions backwards strictly decreases time, and therefore reaches the source s . Hence this is a broadcast protocol of length at most t , and $b(G, s) \leq t$. ◁

The algorithm accepts if at least one consistent template is extendable. Enumerating the templates and testing all component-slot compatibilities takes time $k^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)} = \text{vi}^{\mathcal{O}(\text{vi}^2)} n^{\mathcal{O}(1)}$, and each extension test is polynomial. ◀

5 Distance to Clique

A set $S \subseteq V(G)$ is a *modulator to a clique* if $G - S$ is a clique. The *distance to clique* of a graph G is the minimum size of such a set. Equivalently, a modulator to a clique in G is a

vertex cover of the complement $\overline{G}$. Thus a minimum modulator can be computed in time $2^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$ [8, 22]. Egami et al. [13] proved that TELEPHONE BROADCAST is fixed-parameter tractable parameterized by distance to clique k , presenting an algorithm of running time $2^{\mathcal{O}(k^2)}n^{\mathcal{O}(1)}$. In this section we improve this dependence to $k^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$.

The reason this parameter admits a compact description of a protocol is that the clique side is highly symmetric. Once one clique vertex is informed, the message can spread inside the clique without using the modulator again. Thus the only clique vertices that need to be represented explicitly are those that send the message back to the modulator. After adding the source to the modulator, there are at most $|S| - 1$ such vertices. We represent them by slots in a template, use a b -MATCHING instance to choose the actual clique vertices realizing these slots, and then complete the broadcast inside the clique greedily.

We use the following upper bound of Egami et al. [13] to restrict the length of the optimal protocol.

► **Lemma 12** ([13, Lemma 10]). *Let G be a connected graph, and let $S \subseteq V(G)$ be such that $G - S$ is a clique of size $N \geq 1$. Then $b(G, s) \leq |S| + \lceil \log_2 N \rceil$ for every source $s \in V(G)$.*

We now present the improved algorithm parameterized by distance to clique.

► **Theorem 13.** TELEPHONE BROADCAST can be solved in time $k^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$, where k is the distance to clique of the input graph.

Proof. Let (G, s, t) be an instance of TELEPHONE BROADCAST, where $G = (V, E)$ is connected and $|V| = n$. We compute a minimum modulator S_0 to a clique in time $2^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$. Let $S := S_0 \cup \{s\}$, $h := |S| \leq k + 1$, and $C := V \setminus S$ where $N := |C|$. It holds that $s \in S$ and $G[C] = G - S$ is a clique. Assume that $t < h + \lceil \log_2 N \rceil$, since otherwise the answer is trivially yes by Lemma 12.

Template. Consider an optimal protocol P and let T be its broadcast tree. We call a clique vertex *marked* if it has a child in S in T . Assigning to each marked clique vertex one of its children in $S \setminus \{s\}$ gives an injection into $S \setminus \{s\}$, because each vertex has a unique parent in T . Hence there are at most $h - 1$ marked clique vertices.

We enumerate all possible guesses for this bounded part of the protocol. For each $w \in \{0, \dots, h - 1\}$, we introduce a set Q of w *marked slots*. A marked slot is a placeholder for a marked vertex of C ; it will later be realized by a distinct clique vertex. The named vertices of the template are the vertices of S .

We define a template $\mathcal{T} = (H, s, S, Q, \lambda)$, with $H = (S \cup Q \cup \{\phi\}, A)$, by a parent map, where ϕ is an auxiliary vertex. For every $u \in S \setminus \{s\}$, we guess a parent $p(u) \in (S \cup Q) \setminus \{u\}$. For every slot $q \in Q$, we guess $p(q) \in S \cup \{\phi\}$, where ϕ means that the clique vertex realizing q is informed from inside the clique. We add the arc $p(z)z$ for every $z \in (S \setminus \{s\}) \cup Q$. Thus every vertex of $(S \setminus \{s\}) \cup Q$ has a unique incoming arc, while ϕ has none.

We also guess the label of every arc of H from $[t]$, as well as a *load vector* $\ell: S \rightarrow \{0, \dots, t\}$, where $\ell(u)$ is the number of unmarked clique vertices informed directly by u . The marked slots informed by u are already visible in the template. For every $u \in S$, let

$$\text{children}_S(u) := \{u' \in S \setminus \{s\} \mid p(u') = u\} \quad \text{and} \quad \text{children}_Q(u) := \{q \in Q \mid p(q) = u\}.$$

For the rest of the proof, the guessed object is the pair $(\mathcal{T} = (H, s, S, Q, \lambda), \ell)$, where, for notational convenience, we write $\tau = \tau_\lambda$. We call this pair *consistent* if $\mathcal{T}$ satisfies the common consistency checks of Section 2.2 and the following additional conditions:

- every slot has at least one child in S ;

- the outgoing arcs of each slot have pairwise distinct labels;
- for every $u \in S$, $|\text{children}_S(u)| + |\text{children}_Q(u)| + \ell(u) \leq t - \tau(u)$.

The last condition says that, after u is informed, there are enough rounds for the represented outgoing arcs of u and for the $\ell(u)$ additional entries from u into the clique.

▷ **Claim 14.** The number of consistent pairs $(\mathcal{T}, \ell)$ considered is $h^{\mathcal{O}(h)} n^{\mathcal{O}(1)}$.

Proof. Fix $|Q| = w \leq h - 1$. There are at most $(2h)^h$ choices for the parents of the vertices in $S \setminus \{s\}$, and at most $(h + 1)^h$ choices for the parents of the slots. Since $t < h + \lceil \log_2 n \rceil$, the labels and the load vector together have $(h + \log n)^{\mathcal{O}(h)}$ possibilities in total. Hence the number of guesses is

$$h^{\mathcal{O}(h)} (h + \log n)^{\mathcal{O}(h)}.$$

If $\log n \leq h^2$, this is $h^{\mathcal{O}(h)}$. Otherwise $h < \sqrt{\log n}$, and the factor $(h + \log n)^{\mathcal{O}(h)}$ is bounded by $n^{\mathcal{O}(1)}$. $\triangleleft$

Realizing the template. Fix a consistent pair $(\mathcal{T}, \ell)$. For a slot $q \in Q$, let

$$R(q) := \{u \in S \setminus \{s\} \mid p(u) = q\} \cup \begin{cases} \{p(q)\}, & \text{if } p(q) \in S, \\ \emptyset, & \text{if } p(q) = \phi. \end{cases}$$

A clique vertex can realize q only if it is adjacent to every vertex of $R(q)$: it has to inform the children of q in S , and, when $p(q) \in S$, it has to receive the message from $p(q)$. Let $C_q := \{v \in C \mid R(q) \subseteq N_G(v)\}$ denote the set of clique vertices that can realize q .

We use one edge-weighted b -MATCHING instance for two tasks: realizing the slots by distinct clique vertices, and choosing the $\ell(u)$ unmarked clique vertices that are informed directly by each $u \in S$. Its vertex set contains one vertex for every $v \in C$, one vertex w_q for every slot $q \in Q$, and one vertex for every $u \in S$ with $\ell(u) > 0$. The capacities are

$$\beta(v) = 1 \quad (v \in C), \quad \beta(w_q) = 1 \quad (q \in Q), \quad \beta(u) = \ell(u) \quad (u \in S, \ell(u) > 0).$$

We add an edge of weight 2 between w_q and every vertex of C_q . We also add an edge of weight 1 between $u \in S$ and $v \in C$ whenever $uv \in E(G)$ and $\ell(u) > 0$.

A feasible b -matching of weight

$$2|Q| + \sum_{u \in S} \ell(u)$$

is called a *realizing matching* for the pair $(\mathcal{T}, \ell)$. This is the maximum possible weight: the slot vertices contribute at most $2|Q|$, and the vertices of S contribute at most $\sum_{u \in S} \ell(u)$. Hence a realizing matching saturates every slot vertex w_q and uses exactly $\ell(u)$ edges incident with each $u \in S$. Equivalently, it chooses vertices $v_q \in C_q$ for all slots $q \in Q$ and sets $B_u \subseteq N_G(u) \cap C$ with $|B_u| = \ell(u)$ for every $u \in S$, such that all vertices in $\{v_q \mid q \in Q\} \cup \bigcup_{u \in S} B_u$ are distinct. We discard the pair if no realizing matching exists. This test is polynomial-time, since maximum-weight b -MATCHING is polynomial-time solvable [34].

Extending a realization. Fix a consistent pair $(\mathcal{T}, \ell)$ and any realizing matching M . For every slot $q \in Q$, let $v_q \in C$ be the vertex matched to w_q . For every $u \in S$, let

$$B_u := \{v \in C \mid uv \in M\}$$

if $\ell(u) > 0$, and let $B_u := \emptyset$ otherwise. Then $|B_u| = \ell(u)$. Since C is a clique, the greedy completion below depends on M only through the fact that every slot is realized and that $|B_u| = \ell(u)$ for every $u \in S$; the actual identities of the chosen clique vertices are irrelevant.

We first schedule all template arcs not leaving ϕ at their labels: arcs inside S use the corresponding edges of G , an arc uq is realized as uv_q , and an arc qu is realized as v_qu . For every $u \in S$, let

$$F_u := \{\rho \in [\tau(u) + 1, t] \mid \text{no arc of } H \text{ leaves } u \text{ with label } \rho\}.$$

By consistency, $|F_u| \geq \ell(u)$. We schedule the vertices of B_u in the $\ell(u)$ earliest rounds of F_u .

It remains to complete the broadcast inside the clique. All clique vertices except the vertices v_q realizing slots and the vertices in the sets B_u are anonymous. Slots with parent in S and vertices in the sets B_u have fixed incoming transmissions from S . Slots with parent ϕ must instead be informed from inside the clique at their incoming label. Starting with no informed clique vertex, we proceed round by round. At the beginning of round ρ , let N_ρ be the number of informed clique vertices. Let

$$D_\rho := \{q \in Q \mid \text{there is an arc } e = (q, u) \in A \text{ with } u \in S \text{ and } \lambda(e) = \rho\}$$

be the slots that must inform a vertex of S in this round. These slots are already committed to a transmission to S , and therefore are not available for clique-internal transmissions in round ρ . Let

$$E_\rho := \{q \in Q \mid p(q) = \phi \text{ and } \tau(q) = \rho\}$$

be the slots that must be informed from inside the clique in this round. If some slot of D_ρ is not informed at the beginning of the round, or if $N_\rho - |D_\rho| < |E_\rho|$, the test fails. Otherwise, there are $N_\rho - |D_\rho|$ informed clique vertices available for clique-internal transmissions. We use $|E_\rho|$ of them to inform the slots of E_ρ , and the remaining available transmissions to inform as many still uninformed anonymous clique vertices as possible. At the end of the round, we mark as informed all clique vertices reached in this round: the slots of E_ρ , the anonymous clique vertices just chosen greedily, and the clique vertices whose fixed incoming transmission from S occurs in round ρ . The latter are the slots with parent in S and the vertices of the sets B_u scheduled in round ρ .

Since the test only uses the fixed time-stamps and the numbers $|B_u| = \ell(u)$, it is independent of the particular realizing matching. We call the pair $(\mathcal{T}, \ell)$ *extendable* if it admits a realizing matching and this greedy test succeeds.

▷ **Claim 15.** If $b(G, s) \leq t$, then there exists a consistent pair $(\mathcal{T}, \ell)$ that is extendable.

Proof. Suppose that $b(G, s) \leq t$, and fix a protocol P of length at most t with broadcast tree T . Create one slot for every marked clique vertex of P . For every vertex of $S \setminus \{s\}$, record its parent in T ; if the parent is a marked clique vertex, record the corresponding slot. For every slot, record its parent if that parent lies in S , and record ϕ otherwise. The labels of the incoming arcs are the reception times read directly from P . Finally, set $\ell(u)$ to be the number of unmarked clique vertices informed directly by u in P . This gives one of the pairs enumerated above, and it is consistent.

We now build a realizing matching. For every slot q , match w_q to the actual clique vertex realizing q in P . This edge exists because all adjacencies to vertices of S required by the template occur in P . For every clique vertex that is informed by some $u \in S$ and is not represented by a slot, add the edge between u and that clique vertex. The chosen edges

respect all capacities, because a clique vertex appears only once in the broadcast tree, and each vertex u informs exactly $\ell(u)$ unmarked clique vertices. The weight is $2|Q| + \sum_{u \in S} \ell(u)$, so this is a realizing matching.

It remains to see that the greedy extension test succeeds. Compare it with the restriction of P to the clique. The template fixes the same slot-to- S transmissions as P , and every slot with parent ϕ must be informed in P at its guessed round. For each $u \in S$, the extension schedules exactly as many unmarked clique vertices from u as P does, and it uses the earliest free rounds of u for them. Hence these entries occur no later than the corresponding entries of P . Since unmarked clique vertices are interchangeable, an induction over the rounds shows that, just before each round, the greedy extension has informed at least as many clique vertices as P has, while the same slots are busy sending to S . Thus whenever P can inform the slots with $p(q) = \phi$, the greedy extension can do so as well; after that it uses all remaining clique-internal capacity. Since P informs all clique vertices by round t , the greedy extension also does. Hence the pair $(\mathcal{T}, \ell)$ is extendable. $\triangleleft$

$\triangleright$ **Claim 16.** If there exists a consistent pair $(\mathcal{T}, \ell)$ that is extendable, then $b(G, s) \leq t$.

Proof. Suppose that a consistent pair $(\mathcal{T}, \ell)$ is extendable, and fix a realizing matching and greedy extension witnessing this. The construction above gives a protocol of length at most t . All transmissions are along edges of G : this follows from consistency for arcs inside S , from the definition of the sets C_q for arcs between S and slots, from the matching edges between S and C , and from the fact that C is a clique for all clique-internal transmissions, including those used to satisfy arcs leaving ϕ .

No vertex transmits before being informed or twice in the same round. For vertices of S this follows from consistency and from the choice of the free rounds F_u ; for clique vertices it is enforced by the greedy simulation, which treats slot-to- S transmissions as occupied rounds. Finally, every vertex is informed by construction. Following incoming transmissions backward strictly decreases time, and therefore reaches the source s . Hence $b(G, s) \leq t$. $\triangleleft$

By the two preceding claims, testing all pairs $(\mathcal{T}, \ell)$ decides whether $b(G, s) \leq t$. By Claim 14, the number of pairs is $h^{\mathcal{O}(h)} n^{\mathcal{O}(1)} = k^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$, since $h \leq k + 1$. For each pair we solve one polynomial-size b -MATCHING instance to test whether it admits a realizing matching and, if so, run the polynomial-time greedy extension test. Hence the total running time is $k^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$. $\blacktriangleleft$

6 Conclusion

We gave faster FPT algorithms for TELEPHONE BROADCAST under various parameterizations, leading to faster and simpler algorithms that make use of maximum-weight b -MATCHING. As an interesting open question, we ask whether the problem admits an algorithm running in time $2^{\mathcal{O}(\text{vc})} n^{\mathcal{O}(1)}$ or perhaps whether a $\text{vc}^{\mathcal{O}(\text{vc})}$ dependence can be excluded under the ETH. Given the problem's hardness, another interesting question is which other structural parameterizations render it FPT; is that the case for cutwidth, or perhaps bandwidth?

References

- 1 Akash Ambashankar and Hovhannes A. Harutyunyan. Broadcasting in stars of cliques and path-connected cliques. *Algorithms*, 18(2):76, 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*, volume 334 of *LIPICs*, pages 10:1–10:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPICs.ICALP.2025.10.
- 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 Jeffrey Bringolf, Hovhannes A. Harutyunyan, Shahin Kamali, and Seyed-Mohammad Seyed-Javadi. On hardness and approximation of broadcasting in structured graphs, 2026. arXiv:2510.20026.
 - 7 Maja Cevnik and Janez Zerovnik. Broadcasting on cactus graphs. *J. Comb. Optim.*, 33(1):292–316, 2017. doi:10.1007/S10878-015-9957-8.
 - 8 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.
 - 9 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.
 - 10 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.
 - 11 Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate texts in mathematics*. Springer, 2017. doi:10.1007/978-3-662-53622-3.
 - 12 Pål Grønås Drange, Markus Sortland 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.
 - 13 Yudai Egami, Tatsuya Gima, Tesshu Hanaka, Yasuaki Kobayashi, Michael Lampis, Valia Mitsou, Edouard Nemery, Yota Otachi, Manolis Vasilakis, and Daniel Vaz. Broadcasting under structural restrictions. In *50th International Symposium on Mathematical Foundations of Computer Science, MFCS 2025*, volume 345 of *LIPICs*, pages 42:1–42:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPICs.MFCS.2025.42.
 - 14 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.
 - 15 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.
 - 16 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.
 - 17 Arthur M. Farley. Minimum-time line broadcast networks. *Networks*, 10(1):59–70, 1980. doi:10.1002/NET.3230100106.
 - 18 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.
 - 19 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.
 - 20 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.

- 21 Michelangelo Grigni and David Peleg. Tight bounds on minimum broadcast networks. *SIAM J. Discret. Math.*, 4(2):207–222, 1991. doi:10.1137/0404021.
- 22 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*, LIPIcs, pages 40:1–40:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICs.STACS.2024.40.
- 23 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.
- 24 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.
- 25 Daniel Hathcock, Guy Kortsarz, and R. Ravi. The telephone k-multicast problem. *Algorithmica*, 88(2):25, 2026. doi:10.1007/S00453-026-01377-5.
- 26 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.
- 27 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.
- 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 Bernhard Korte and Jens Vygen. *Combinatorial Optimization: Theory and Algorithms*. Springer, 6 edition, 2018. doi:10.1007/978-3-662-56039-6.
- 30 Guy Kortsarz and David Peleg. Approximation algorithms for minimum-time broadcast. *SIAM J. Discret. Math.*, 8(3):401–427, 1995. doi:10.1137/S0895480193245923.
- 31 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.
- 32 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.
- 33 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.
- 34 Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2003.
- 35 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.
- 36 Prafullkumar Tale. Telephone broadcast on graphs of treewidth two. *Theor. Comput. Sci.*, 1045:115282, 2025. doi:10.1016/J.TCS.2025.115282.