Problems in NP Can Admit Double-Exponential Lower Bounds When Parameterized by Treewidth or Vertex Cover

Florent Foucaud
Université Clermont Auvergne, CNRS, Mines Saint-Étienne, Clermont Auvergne INP, LIMOS, 63000 Clermont-Ferrand, France

Esther Galby
Department of Computer Science and Engineering, Chalmers University of Technology, Gothenburg, Sweden
University of Gothenburg, Sweden

Liana Khazaliya
Technische Universität Wien, Austria

Shaohua Li
CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

Fionn Mc Inerney
Technische Universität Wien, Austria

Roohani Sharma
University of Bergen, Norway

Prafullkumar Tale
Indian Institute of Science Education and Research Pune, India

Abstract

Treewidth serves as an important parameter that, when bounded, yields tractability for a wide class of problems. For example, graph problems expressible in Monadic Second Order (MSO) logic and Quantified SAT or, more generally, Quantified CSP, are fixed-parameter tractable parameterized by the treewidth of the input's (primal) graph plus the length of the MSO-formula [Courcelle, Information & Computation 1990] and the quantifier rank [Chen, ECAI 2004], respectively. The algorithms generated by these (meta-)results have running times whose dependence on treewidth is a tower of exponents. A conditional lower bound by Fichte, Hecher, and Pfandler [LICS 2020] shows that, for Quantified SAT, the height of this tower is equal to the number of quantifier alternations. These types of lower bounds, which show that at least *double-exponential* factors in the running time are *necessary*, exhibit the extraordinary level of computational hardness for such problems, and are rare in the current literature: there are only a handful of such lower bounds (for treewidth and vertex cover parameterizations) and all of them are for problems that are #NP-complete, Σ_2^p-complete, Π_2^p-complete, or complete for even higher levels of the polynomial hierarchy.

Our results demonstrate, for the first time, that it is not necessary to go higher up in the polynomial hierarchy to achieve double-exponential lower bounds: we derive double-exponential lower bounds in the treewidth (tw) and the vertex cover number (vc), for natural, important, and well-studied NP-*complete* graph problems. Specifically, we design a *technique* to obtain such lower bounds and show its *versatility* by applying it to three different problems: Metric Dimension, Strong Metric Dimension, and Geodetic Set. We prove that these problems do not admit $2^{2^{o(\mathrm{tw})}} \cdot n^{\mathcal{O}(1)}$-time algorithms, even on bounded diameter graphs, unless the ETH fails (here, n is the number of vertices in the graph). In fact, for Strong Metric Dimension, the double-exponential lower bound holds even for the vertex cover number. We further complement all our lower bounds with matching (and sometimes non-trivial) upper bounds.

For the conditional lower bounds, we design and use a novel, yet simple technique based on *Sperner families* of sets. We believe that the amenability of our technique will lead to obtaining such lower bounds for many other problems in NP.

51st International Colloquium on Automata, Languages, and Programming (ICALP 2024).
Editors: Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson;
Article No. 66; pp. 66:1–66:19
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

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

Keywords and phrases Parameterized Complexity, ETH-based Lower Bounds, Double-Exponential Lower Bounds, Kernelization, Vertex Cover, Treewidth, Diameter, Metric Dimension, Strong Metric Dimension, Geodetic Sets

Digital Object Identifier 10.4230/LIPIcs.ICALP.2024.66

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2307.08149`

Funding *Florent Foucaud*: ANR project GRALMECO (ANR-21-CE48-0004), French government IDEX-ISITE initiative 16-IDEX-0001 (CAP 20-25), International Research Center "Innovation Transportation and Production Systems" of the I-SITE CAP 20-25.
Liana Khazaliya: Vienna Science and Technology Fund (WWTF) [10.47379/ICT22029]; Austrian Science Fund (FWF) [Y1329]; European Union's Horizon 2020 COFUND programme [LogiCS@TUWien, grant agreement No. 101034440].
Shaohua Li: European Research Council (ERC) consolidator grant No. 725978 SYSTEMATIC-GRAPH.
Fionn Mc Inerney: Austrian Science Fund (FWF, project Y1329).

1 Introduction

Many interesting computational problems turn out to be intractable. In these cases, identifying parameters under which the problems become tractable is desirable. In the area of parameterized complexity, treewidth is a cornerstone parameter since a large class of problems become tractable on graphs of bounded treewidth.

Courcelle's celebrated theorem [13] states that the class of graph problems expressible in Monadic Second-Order Logic (MSOL) of constant size is fixed-parameter tractable (FPT) when parameterized by the treewidth of the graph. That is, such problems admit algorithms whose running time is of the form $f(\texttt{tw}) \cdot \mathsf{poly}(n)$, where $\texttt{tw}$ is the treewidth of the input, n is the size of the input, and f is a function that depends only on $\texttt{tw}$. Similarly, a result by Chen [12] shows that the Quantified SAT (Q-SAT) problem can also be solved in time $f(\texttt{tw}) \cdot \mathsf{poly}(n)$, where $\texttt{tw}$ is the treewidth of the primal graph of the input formula and f is a function that depends only on $\texttt{tw}$ and the number of quantifier alternations in the input formula. Q-SAT is a generalization of SAT that allows universal and existential quantifications over the variables. Note that Q-SAT with k quantifier alternations is Π_k^p-complete or Σ_k^p-complete. Unfortunately, in both of the aforementioned results, the function f is a tower of exponents whose height depends roughly on the size of the MSOL and input formulas, respectively. For Q-SAT, the height of this tower equals the number of quantifier alternations in the Q-SAT instance [12].

Over the years, the focus shifted to making such FPT algorithms as efficient as possible. Thus, a natural question is to ask when this higher-exponential dependence on treewidth is necessary. There is a rich literature that provides (conditional) lower bounds on this dependency for many problems, and these bounds are commonly of the form $2^{o(\texttt{tw})}$ or, in some unusual cases, $2^{o(\texttt{tw} \log \texttt{tw})}$ (e.g., [15, 50]) and even $2^{o(\mathsf{poly}(\texttt{tw}))}$ (e.g., [14, 55]). Most notably, these lower bounds are far from the tower of exponents upper bounds given by the (meta-)results discussed above. In this work, we develop a simple technique that allows to prove double-exponential dependence on the treewidth $\texttt{tw}$ and the vertex cover number $\texttt{vc}$, two of the most fundamental graph parameters. Notably, these are the first such results for problems in NP, and we believe that the amenability of our technique will lead to many more similar results for other problems in NP.

Indeed, after a preprint of this paper appeared on arxiv, our technique was also used to prove double-exponential dependence on vc for an NP-complete machine learning problem [11] and double-exponential dependence on the solution size and tw for NP-complete *identification problems* like Test Cover and Locating-Dominating Set [10].

Double-exponential lower bounds: treewidth and vertex cover parameterizations. Fichte, Hecher, and Pfander [25] recently proved that, assuming the Exponential Time Hypothesis[1] (ETH), Q-SAT with k quantifier alternations cannot be solved in time significantly better than a tower of exponents of height k in the treewidth. This exemplifies an interesting but expected trait of this problem: its complexity, in terms of the height of the exponential tower in tw, increases with each quantifier alternation. It strengthened the result that appeared in [48], where conditional double-exponential lower bounds for ∃∀SAT and ∀∃SAT were given. The results in [48] also yield a double-exponential lower bound in vc of the primal graph for both problems. Besides these results, there are only a handful of other problems known to require higher-exponential dependence in the treewidth of the input graph (or the primal graph of the input formula). Specifically, the Π_2^p-complete k-Choosability problem and the Σ_3^p-complete k-Choosability Deletion problem admit a double-exponential and a triple-exponential lower bound in treewidth [52], respectively. Recently, the Σ_p^2-complete problems Cycle HitPack and H-HitPack, for a fixed graph H, were shown to admit tight algorithms that are double-exponential in the treewidth [26]. Further, the Σ_p^2-complete problem Core Stability was shown to admit a tight double-exponential lower bound in the treewidth, even on graphs of bounded degree [32]. Lastly, the #NP-complete counting problem Projected Model Counting admits a double-exponential lower bound in tw [23, 24]. For other double-exponential lower bounds, see [1, 16, 27, 32, 37, 41, 43, 44, 47, 51, 56, 59].

All the double- (or higher) exponential lower bounds in treewidth mentioned so far are for problems that are #NP-complete, Σ_2^p-complete, Π_2^p-complete, or complete for even higher levels of the polynomial hierarchy. To quote [52]: "*Π_2^p-completeness of these problems already gives sufficient explanation why double- [...] exponential dependence on treewidth is needed. [...] the quantifier alternations in the problem definitions are the common underlying reasons for being in the higher levels of the polynomial hierarchy and for requiring unusually large dependence on treewidth.*"

As mentioned above, we develop a technique that allows to demonstrate, for the first time, that it is not necessary to go to higher levels of the polynomial hierarchy to achieve double-exponential lower bounds in the treewidth or the vertex cover number of the graph.

> We prove that three natural and well-studied NP-complete problems admit double-exponential lower bounds in tw or vc, under the ETH. These are the first problems in NP known to admit such lower bounds.[2]

[1] The Exponential Time Hypothesis roughly states that n-variable 3-SAT cannot be solved in time $2^{o(n)}$.

[2] While it may be possible to artificially engineer a graph problem or graph representation of a problem in NP that admits such lower bounds (although, to the best of our knowledge, this has not been done), we emphasize that this is not the case for these three natural and well-established graph problems in NP.

NP-complete metric-based graph problems. We study three *metric-based graph problems*. These problems are Metric Dimension [34, 58], Strong Metric Dimension [57], and Geodetic Set [33], and they arise from network design and network monitoring. Apart from serving as examples for double-exponential dependence on treewidth and the amenability of our technique, these problems are of interest in their own right, and possess a rich literature both in the algorithms and discrete mathematics communities (see, e.g., [2, 3, 4, 5, 7, 8, 17, 19, 20, 21, 31, 38, 45, 53] and the references below). Their non-local nature has posed interesting algorithmic challenges and our results, as we explain later, supplement the already vast literature on the structural parameterizations of these problems. Below we define the three above-mentioned problems formally, and particularly focus on Metric Dimension as it is the most popular and well-studied of the three.

Metric Dimension
Input: A graph G and a positive integer k.
Question: Does there exist $S \subseteq V(G)$ such that $|S| \leq k$ and, for any pair of vertices $u, v \in V(G)$, there exists a vertex $w \in S$ with $d(w, u) \neq d(w, v)$?

The Metric Dimension problem dates back to the 70s [34, 58]. As in geolocation problems, the aim is to distinguish the vertices of a graph via their distances to a solution set. Metric Dimension was first shown to be NP-complete in general graphs in Garey and Johnson's book [30, GT61], and this was later extended to many restricted graph classes [18, 22, 28], including graphs of diameter 2 [28] and graphs of pathwidth 24 [49]. In a seminal paper, Metric Dimension was proven to be W[2]-hard parameterized by the solution size k, even in subcubic bipartite graphs [35]. This drove the subsequent meticulous study of the problem under structural parameterizations.

In particular, the complexity of Metric Dimension parameterized by treewidth remained an intriguing open problem for a long time. Recently, it was shown that Metric Dimension is para-NP-hard parameterized by pathwidth ($\texttt{pw}$) [49] (an earlier result [6] showed that it is W[1]-hard for pathwidth). A subsequent paper showed that the problem is W[1]-hard parameterized by the combined parameter feedback vertex set number ($\texttt{fvs}$) plus pathwidth of the graph [29].

We conclude this part with the definitions of the remaining two problems, both of which are known to be NP-Complete [9, 54]. Geodetic Set is also W[1]-hard parameterized by the solution size, feedback vertex set number, and pathwidth, combined [39].

Strong Metric Dimension
Input: A graph G and a positive integer k.
Question: Does there exist $S \subseteq V(G)$ such that $|S| \leq k$ and, for any pair of vertices $u, v \in V(G)$, there exists a vertex $w \in S$ such that either u lies on some shortest path between v and w, or v lies on some shortest path between u and w?

Geodetic Set
Input: A graph G and a positive integer k.
Question: Does there exist $S \subseteq V(G)$ such that $|S| \leq k$ and, for any vertex $u \in V(G)$, there are two vertices $s_1, s_2 \in S$ such that a shortest path from s_1 to s_2 contains u?

Our technical contributions. As Metric Dimension and Geodetic Set are NP-complete on bounded diameter graphs *or* on bounded treewidth graphs, we study their parameterized complexity with $\texttt{tw} + \texttt{diam}$ as the parameter and prove the following results.

1. METRIC DIMENSION and GEODETIC SET do not admit algorithms running in time $2^{f(\texttt{diam})^{o(\texttt{tw})}} \cdot n^{\mathcal{O}(1)}$, for any computable function f, unless the ETH fails.
2. STRONG METRIC DIMENSION does not even admit an algorithm with a running time of $2^{2^{o(\texttt{vc})}} \cdot n^{\mathcal{O}(1)}$, unless the ETH fails. This also implies the problem does not admit a kernelization algorithm that outputs an instance with $2^{o(\texttt{vc})}$ *vertices*, unless the ETH fails.

The above lower bounds for $\texttt{tw} + \texttt{diam}$, in particular, imply that METRIC DIMENSION and GEODETIC SET on graphs of bounded diameter cannot admit $2^{2^{o(\texttt{tw})}} \cdot n^{\mathcal{O}(1)}$-time algorithms, unless the ETH fails. The reduction for METRIC DIMENSION (sketched in Section 2.3) also works for $\texttt{fvs}$ and $\texttt{td}$, and the one for GEODETIC SET (Section 3) also works for $\texttt{td}$.

We show that all our lower bounds are tight by providing algorithms (kernelization algorithms, respectively) with matching running times (guarantees, respectively).

1. METRIC DIMENSION and GEODETIC SET admit algorithms running in time $2^{\texttt{diam}^{\mathcal{O}(\texttt{tw})}} \cdot n^{\mathcal{O}(1)}$.
2. STRONG METRIC DIMENSION admits an algorithm running in time $2^{2^{\mathcal{O}(\texttt{vc})}} \cdot n^{\mathcal{O}(1)}$ and a kernel with $2^{\mathcal{O}(\texttt{vc})}$ vertices.

The (kernelization) algorithm for the $\texttt{vc}$ parameterization is very simple, whereas the algorithms for the $\texttt{tw} + \texttt{diam}$ parameter are highly non-trivial and require showing interesting locality properties in the instance. Further, for our $\texttt{tw} + \texttt{diam}$ parameterized algorithms, the (double-exponential) dependency of treewidth in the running time is unusual (and rightly so, as exhibited by our lower bounds), as most natural graph problems in NP for which a dedicated algorithm (i.e., *not* relying on Courcelle's theorem) parameterized by treewidth is known, can be solved in time $2^{\mathcal{O}(\texttt{tw})} \cdot n^{\mathcal{O}(1)}$, $2^{\mathcal{O}(\texttt{tw} \cdot \log(\texttt{tw}))} \cdot n^{\mathcal{O}(1)}$ or $2^{\mathcal{O}(\mathsf{poly}(\texttt{tw}))} \cdot n^{\mathcal{O}(1)}$.

Finally, our reductions rely on a novel, yet simple technique based on *Sperner families* of sets that allows to encode particular SAT relations across large sets of variables and clauses into relatively small vertex-separators. As mentioned before, we believe that this technique is the key to obtaining such lower bound results for other problems in NP. In particular, as witnessed by our results, our technique has the additional features that it even allows to prove such lower bounds in very restricted cases, such as bounded diameter graphs, and is not specific to any one structural parameter, as it also works for, e.g., the feedback vertex set number and treedepth.

Due to space constraints, we cannot discuss all of our results in depth, and refer the reader to the full version for full proofs, formal details, and more related work. Nonetheless, we elaborate on our technique and present an overview of the results for METRIC DIMENSION in the next section. Then, in Section 3, we present a formal proof for the lower bound for GEODETIC SET. Finally, we conclude the paper in Section 4.

2 Technical Overview

In this section, we present an overview of our lower bound techniques. We first exhibit our technique to obtain the double-exponential lower bounds in its most general setting. Then, we continue with the problem-specific tools we developed that are required for the reductions.

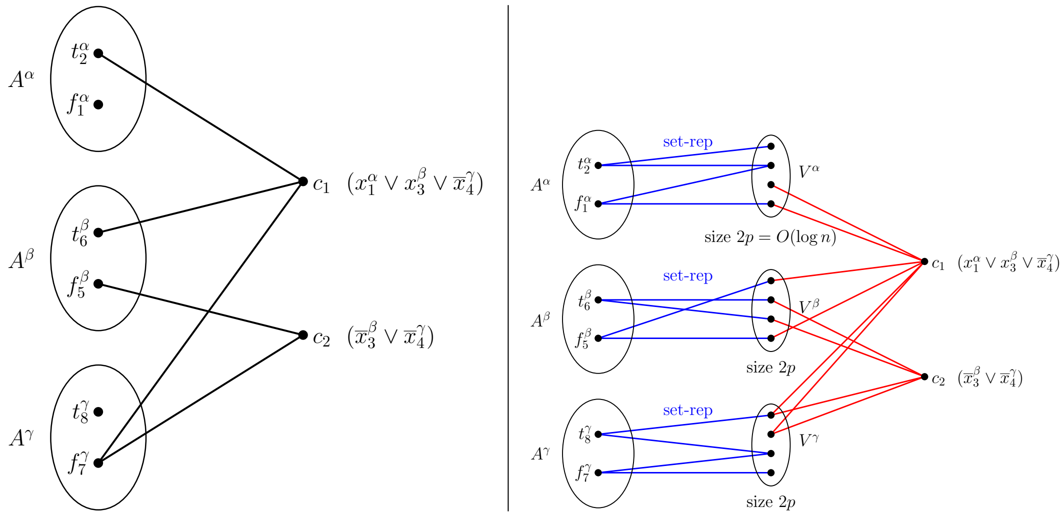

Figure 1 Graph representations of 3-Partitioned-3-SAT. (Left) incidence graph representation. (Right) representation with small separators using our technique. Note, for example, that x_1^α appears as a positive literal in the clause C_1. Thus, on the left, t_2^α is the only literal vertex in A^α incident to c_1, while on the right, t_2^α is the only literal vertex in A^α that does not share a common neighbor with c_1 in V^α. The edges from c_2 to each vertex in V^α are omitted for clarity.

2.1 General Technique for Double-Exponential Lower Bounds

The first integral part of our technique is to reduce from a variant of 3-SAT known as 3-Partitioned-3-SAT that was introduced in [46]. In this problem, the input is a formula ψ in 3-CNF form, together with a partition of the set of its variables into three disjoint sets X^α, X^β, X^γ, with $|X^\alpha| = |X^\beta| = |X^\gamma| = n$, and such that no clause contains more than one variable from each of X^α, X^β, and X^γ. The objective is to determine whether ψ is satisfiable. Unless the ETH fails, 3-Partitioned-3-SAT does not admit an algorithm running in time $2^{o(n)}$ [46, Theorem 3].

Typical reductions from satisfiability problems to graph problems usually entail representing the satisfiability problem by its incidence graph, in which each variable is represented by two vertices corresponding to its positive and negative literals. In this representation, a clause vertex is adjacent to a literal vertex if and only if it contains that literal in ψ (see Figure 1 (left) for an illustration). However, this naive approach does not lead to any structural parameters of the incidence graph being of bounded size. The *core idea* of our technique is to instead represent the relationships between clause and literal vertices via edges from these two sets of vertices to "small" separators (three separators in the case of 3-Partitioned-3-SAT) that encode these relationships.

Formally, this is achieved as follows. For a positive integer p, define $\mathcal{F}_p$ as the collection of subsets of $[2p]$ that contains exactly p integers. We critically use the fact that no set in $\mathcal{F}_p$ is contained in any other set in $\mathcal{F}_p$ (such a collection of sets are called a *Sperner family*). Let ℓ be a positive integer such that $\ell \leq \binom{2p}{p}$. We define $\mathsf{set\text{-}rep} : [\ell] \mapsto \mathcal{F}_p$ as a one-to-one function by arbitrarily assigning a set in $\mathcal{F}_p$ to an integer in $[\ell]$. By the asymptotic estimation of the central binomial coefficient, $\binom{2p}{p} \sim \frac{4^p}{\sqrt{\pi \cdot p}}$ [36]. To get the upper bound of p, we scale down the asymptotic function and have $\ell \leq \frac{4^p}{2^p} = 2^p$. Thus, $p = \mathcal{O}(\log \ell)$.

Let ψ be an instance of 3-Partitioned-3-SAT on $3n$ variables, and let p be the smallest integer such that $2n \leq \binom{2p}{p}$. In particular, $p = \mathcal{O}(\log n)$. Define $\mathsf{set\text{-}rep} : [2n] \mapsto \mathcal{F}_p$ as above. Rename the variables in X^α to x_i^α for all $i \in [n]$. For each variable x_i^α, add two vertices t_{2i}^α

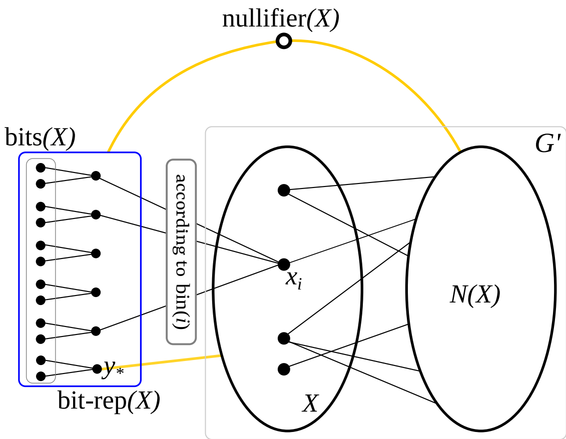

Figure 2 Set Identifying Gadget. The blue box represents $\mathsf{bit\text{-}rep}(X)$ and the yellow lines represent that $\mathsf{nullifier}(X)$ is adjacent to each vertex in $\mathsf{bit\text{-}rep}(X) \cup N(X)$, and $y_\star$ is adjacent to each vertex in X. Also, G' is not necessarily restricted to the graph induced by the vertices in $X \cup N(X)$.

and f^α_{2i-1} corresponding to the positive and negative literals of x^α_i, respectively. Let $A^\alpha = \{t^\alpha_{2i}, f^\alpha_{2i-1} | \ i \in [n]\}$. Add a *validation portal* with $2p$ vertices, denoted by $V^\alpha = \{v^\alpha_1, \dots, v^\alpha_{2p}\}$. For each $i \in [n]$, add the edge $t^\alpha_{2i} v^\alpha_{p'}$ for each $p' \in \mathsf{set\text{-}rep}(2i)$. Similarly, for each $i \in [n]$, add the edge $f^\alpha_{2i-1} v^\alpha_{p'}$ for each $p' \in \mathsf{set\text{-}rep}(2i-1)$. Repeat the above steps for β and γ.

Now, for each clause C_j ($j \in [m]$) in ψ, add a clause vertex c_j. Let $\delta \in \{\alpha, \beta, \gamma\}$. For all $i \in [n]$ and $j \in [m]$, if the variable x^δ_i appears as a positive (negative, respectively) literal in the clause C_j in ψ, then add the edge $c_j v^\delta_{p'}$ for each $p' \in [2p] \setminus \mathsf{set\text{-}rep}(2i)$ ($p' \in [2p] \setminus \mathsf{set\text{-}rep}(2i-1)$, respectively). For all $j \in [m]$, if no variable from X^δ appears in C_j in ψ, then make c_j adjacent to all the vertices in V^δ. See Figure 1 (right) for an illustration.

As a clause contains at most one variable from X^δ in ψ, c_j and t^δ_{2i} (f^δ_{2i-1}, respectively) do not share a common neighbor in V^δ if and only if the clause C_j contains x^δ_i as a positive (negative, respectively) literal in ψ. For the reductions, we use this representation of the relationship between clause and literal vertices. Since $p = \mathcal{O}(\log n)$, this ensures that $\mathtt{tw}(G) = \mathcal{O}(\log n)$, which we exploit along with the fact that, unless the ETH fails, 3-Partitioned-3-SAT does not admit an algorithm running in time $2^{o(n)}$.

2.2 Basic Tools for Lower Bounds

For brevity, we focus on Metric Dimension and explain our problem-specific tools in this context. We use two such simple tools: the *bit representation gadget* and the *set representation gadget*. The set representation gadget is the problem-specific implementation of the above technique, and it uses the bit representation gadget.

Before going further, we need to define some terms related to Metric Dimension. The set S defined in the problem statement of Metric Dimension is called a *resolving set* of G. A subset of vertices $S' \subseteq V(G)$ *resolves* a pair of vertices $u, v \in V(G)$ if there exists a vertex $w \in S'$ such that $d(w,u) \neq d(w,v)$. Lastly, a vertex $u \in V(G)$ is *distinguished* by a subset of vertices $S' \subseteq V(G)$ if, for any $v \in V(G) \setminus \{u\}$, there exists a vertex $w \in S'$ such that $d(w,u) \neq d(w,v)$.

Bit Representation Gadget to Identify Sets. Suppose we are given a graph G' and a subset $X \subseteq V(G')$ of its vertices. Further, suppose that we want to add a vertex set X^+ to G' to obtain a new graph G with the following properties. We want that each vertex in $X \cup X^+$ is distinguished by vertices in X^+ that must be in any resolving set S of G, and that no vertex in X^+ can resolve any "critical pair" of vertices in G. Roughly, a pair of vertices is critical if it forces certain "types" of vertices to be in any resolving set S of G, and the selection of the specific vertices of those types depends on the solution to the problem being reduced from (which, in our case, is 3-Partitioned-3-SAT [46]). We refer to the graph induced by the vertices of X^+, along with the edges connecting X^+ to G', as the *Set Identifying Gadget* for the set X. Given a graph G' and a non-empty subset $X \subseteq V(G')$ of its vertices, to construct such a graph G, we add vertices and edges to G' as follows (see Figure 2):

- The vertex set X^+ that we are aiming to add is the union of a set $\mathsf{bit\text{-}rep}(X)$ and a special vertex denoted by $\mathsf{nullifier}(X)$.
- First, let $X = \{x_i \mid i \in [|X|]\}$, and set $q := \lceil \log(|X| + 2) \rceil + 1$. We select this value for q to *(1)* uniquely represent each integer in $[|X|]$ by its bit-representation in binary (note that we start from 1 and not 0), *(2)* ensure that the only vertex whose bit-representation contains all 1's is $\mathsf{nullifier}(X)$, and *(3)* reserve one spot for an additional vertex $y_\star$.
- For every $i \in [q]$, add three vertices y_i^a, y_i, y_i^b, and add the path (y_i^a, y_i, y_i^b).
- Add 3 vertices $y_\star^a, y_\star, y_\star^b$ and the path $(y_\star^a, y_\star, y_\star^b)$. Add edges to make $\{y_i \mid i \in [q]\} \cup \{y_\star\}$ a clique. Make $y_\star$ adjacent to each vertex in X. Let $\mathsf{bit\text{-}rep}(X) = \{y_i, y_i^a, y_i^b \mid i \in [q]\} \cup \{y_\star, y_\star^a, y_\star^b\}$ and denote its subset by $\mathsf{bits}(X) = \{y_i^a, y_i^b \mid i \in [q]\} \cup \{y_\star^a, y_\star^b\}$.
- For every integer $j \in [|X|]$, let $\mathsf{bin}(j)$ denote the binary representation of j using q bits. Connect x_j with y_i if the i^{th} bit (going from left to right) in $\mathsf{bin}(j)$ is 1.
- Add a vertex, denoted by $\mathsf{nullifier}(X)$, and connect it to each vertex in $\{y_i \mid i \in [q]\} \cup \{y_\star\}$.
- For every vertex $u \in V(G) \setminus (X \cup X^+)$ such that u is adjacent to some vertex in X, add an edge between u and $\mathsf{nullifier}(X)$. We add this vertex to ensure that vertices in $\mathsf{bit\text{-}rep}(X)$ do not resolve critical pairs in $V(G)$.

Set Representation Gadget. We define $\mathsf{set\text{-}rep} : [\ell] \mapsto \mathcal{F}_p$ as in Section 2.1, and recall that $p = \mathcal{O}(\log \ell)$. Suppose we have a "large" collection of vertices, say $A = \{a_1, a_2, \dots, a_\ell\}$, and a "large" collection of critical pairs $C = \{\langle c_1^\circ, c_1^\star \rangle, \langle c_2^\circ, c_2^\star \rangle, \dots, \langle c_m^\circ, c_m^\star \rangle\}$. Moreover, we are given an injective function $\phi : [m] \mapsto [\ell]$. The objective is to design a gadget such that only $a_{\phi(q)} \in A$ can resolve a critical pair $\langle c_q^\circ, c_q^\star \rangle \in C$ for any $q \in [m]$, while keeping the treewidth of this part of the graph of order $\mathcal{O}(\log(|A|))$. With this in mind, we do the following.

- Add vertices and edges to identify the set A and to add critical pairs in C (for each critical pair in C, both vertices share the same bit-representation in the Set Identifying Gadget for C).
- Add a *validation portal*, a clique on $2p$ vertices, denoted by $V = \{v_1, v_2, \dots, v_{2p}\}$, and vertices and edges to identify it.
- For every $i \in [\ell]$ and for every $p' \in \mathsf{set\text{-}rep}(i)$, add the edge $(a_i, v_{p'})$.
- For every critical pair $\langle c_q^\circ, c_q^\star \rangle$, make c_q° adjacent to every vertex in V, and add every edge of the form $(c_q^\star, v_{p'})$ for $p' \in [2p] \setminus \mathsf{set\text{-}rep}(\phi(q))$. Note that the vertices in V that are indexed using integers in $\mathsf{set\text{-}rep}(\phi(q))$ are *not* adjacent with $c_q^\star$.

See Figure 3 for an illustration. Now, consider a critical pair $\langle c_q^\circ, c_q^\star \rangle$ and suppose $i = \phi(q)$.

- By the construction, $N(a_i) \cap N(c_q^\circ) \neq \emptyset$, whereas $N(a_i) \cap N(c_q^\star) = \emptyset$. Hence, a_i resolves the critical pair $\langle c_q^\circ, c_q^\star \rangle$ as $d(a_i, c_q^\circ) = 2$ and $d(a_j, c_q^\star) > 2$.
- For any other vertex in A, say a_j, $\mathsf{set\text{-}rep}(j) \setminus \mathsf{set\text{-}rep}(i)$ is a non-empty set. So, there are paths from a_j to c_q° and a_j to $c_q^\star$ through vertices in V with indices in $\mathsf{set\text{-}rep}(j) \setminus \mathsf{set\text{-}rep}(i)$. This implies that $d(a_j, c_q^\circ) = d(a_j, c_q^\star) = 2$ and a_j cannot resolve the pair $\langle c_q^\circ, c_q^\star \rangle$.

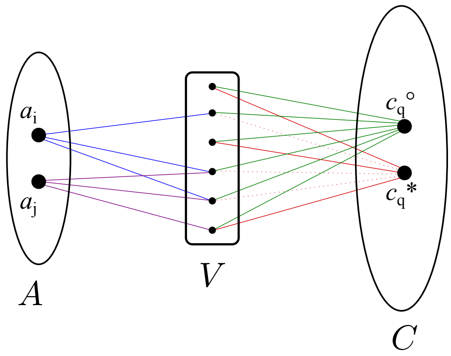

Figure 3 Set Representation Gadget. Let $\phi(q) = i$, i.e., only a_i in A can resolve the critical pair $\langle c_q^\circ, c_q^\star \rangle$. Let the vertices in V be indexed from top to bottom and let $\mathsf{set\text{-}rep}(i) = \{2, 4, 5\}$. By construction, the only vertices in V that $c_q^\star$ is *not* adjacent to are v_2, v_4, and v_5 (this is highlighted by red-dotted edges). Thus, $\text{dist}(a_i, c_q^\circ) = 2$ and $\text{dist}(a_i, c_q^\star) > 2$, and hence, a_i resolves $\langle c_q^\circ, c_q^\star \rangle$. For any other vertex in A, say a_j, $\mathsf{set\text{-}rep}(j) \setminus \mathsf{set\text{-}rep}(i)$ is non-empty, and thus, a_j cannot resolve $\langle c_q^\circ, c_q^\star \rangle$.

2.3 Sketch of the Lower Bound Proof for Metric Dimension

With these tools in hand, we present an overview of the reduction from 3-Partitioned-3-SAT used to prove Theorem 1, which we restate here for convenience.

▶ **Theorem 1.** *Unless the ETH fails, Metric Dimension does not admit an algorithm running in time $2^{f(\mathtt{diam})^{o(\mathtt{tw})}} \cdot n^{\mathcal{O}(1)}$ for any computable function $f : \mathbb{N} \mapsto \mathbb{N}$.*

The reduction in the proof of Theorem 1 takes as input an instance ψ of 3-Partitioned-3-SAT on $3n$ variables and returns (G, k) as an instance of Metric Dimension such that $\mathtt{tw}(G) = \mathcal{O}(\log(n))$ and $\mathtt{diam}(G) = \mathcal{O}(1)$. In the following, we mention a crude outline of the reduction, omitting some technical details.

2.3.1 Reduction

- We rename the variables in X^α to x_i^α for $i \in [n]$. For every variable x_i^α, we add a critical pair $\langle x_i^{\alpha,\circ}, x_i^{\alpha,\star} \rangle$ of vertices. We denote $X^\alpha = \{x_i^{\alpha,\circ}, x_i^{\alpha,\star} \mid i \in [n]\}$.
- For each variable x_i^α, we add the vertices t_{2i}^α and f_{2i-1}^α. Let $A^\alpha = \{t_{2i}^\alpha, f_{2i-1}^\alpha | \ i \in [n]\}$.
- For every $i \in [n]$, we add the edges $(x_i^{\alpha,\circ}, t_{2i}^\alpha)$ and $(x_i^{\alpha,\circ}, f_{2i-1}^\alpha)$ which will ensure that any resolving set contains at least one vertex in $\{t_{2i}^\alpha, f_{2i-1}^\alpha, x_i^{\alpha,\circ}, x_i^{\alpha,\star}\}$ for every $i \in [n]$.
- Let p be the smallest integer such that $2n \leq \binom{2p}{p}$. In particular, $p = \mathcal{O}(\log n)$. Define $\mathsf{set\text{-}rep} : [2n] \mapsto \mathcal{F}_p$ as in Section 2.1.
- We add a *validation portal*, a clique on $2p$ vertices, denoted by $V^\alpha = \{v_1^\alpha, v_2^\alpha, \ldots, v_{2p}^\alpha\}$.
- For each $i \in [n]$, we add the edge $(t_{2i}^\alpha, v_{p'}^\alpha)$ for every $p' \in \mathsf{set\text{-}rep}(2i)$. Similarly, for each $i \in [n]$, we add the edge $(f_{2i-1}^\alpha, v_{p'}^\alpha)$ for every $p' \in \mathsf{set\text{-}rep}(2i-1)$.

  We repeat the above steps to construct $X^\beta, A^\beta, V^\beta$, $X^\gamma, A^\gamma, V^\gamma$.
- For every clause C_q in ψ, we introduce a pair $\langle c_q^\circ, c_q^\star \rangle$ of vertices. Let C be the collection of vertices in such pairs.
- We add edges across C and the portals as follows. Consider a clause C_q in ψ and the corresponding critical pair $\langle c_q^\circ, c_q^\star \rangle$ in C. Let $\delta \in \{\alpha, \beta, \gamma\}$. As ψ is an instance of 3-Partitioned-3-SAT, at most one variable in X^δ appears in C_q, say x_i^δ for some $i \in [n]$. We add all edges of the form $(v_{p'}^\delta, c_q^\circ)$ for every $p' \in [2p]$. If x_i^δ appears as a positive literal

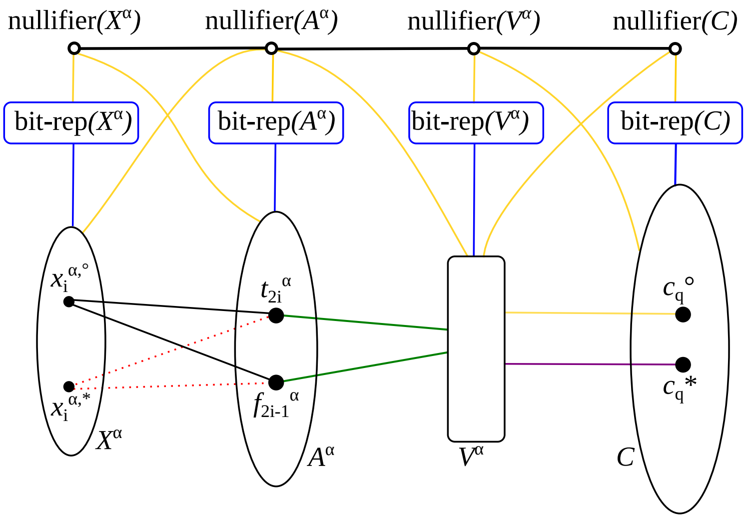

Figure 4 Reduction for proof of Theorem 1. Yellow lines represent that vertex is connected to every vertex in the set the edge goes to. Green edges denote adjacencies with respect to set-rep, e.g., t_{2i}^α is adjacent to $v_j \in V^\alpha$ if $j \in$ set-rep$(2i)$. Purple lines also indicate adjacencies with respect to set-rep, but in a complementary way, i.e., if $x_i \in c_q$, then, for every $p' \in [2p] \setminus$ set-rep$(2i)$, we have $(v_{p'}^\alpha, c_q^\star) \in E(G)$, and if $\overline{x}_i \in c_q$, then, for all $p' \in [2p] \setminus$ set-rep$(2i-1)$, we have $(v_{p'}^\alpha, c_q^\star) \in E(G)$.

in C_q, then we add the edge $(v_{p'}^\delta, c_q^\star)$ for every $p' \in [2p] \setminus$ set-rep$(2i)$ (which corresponds to t_{2i}^δ). If x_i^δ appears as a negative literal in C_q, then we add the edge $(v_{p'}^\delta, c_q^\star)$ for every $p' \in [2p] \setminus$ set-rep$(2i-1)$ (which corresponds to f_{2i-1}^δ). Note that if x_i^δ appears as a positive (negative, respectively) literal in C_q, then the vertices in V^δ whose indices are in set-rep$(2i)$ (set-rep$(2i-1)$, respectively) are *not adjacent* to $c_q^\star$. If no variable in X^δ appears in C_q, then we make each vertex in V^δ adjacent to both c_q° and $c_q^\star$.

For all the sets mentioned above, we add vertices and edges to identify them as shown in Figure 4 (for each critical pair, both vertices share the same bit-representation in their Set Identifying Gadget). This concludes the construction of G. The reduction returns (G, k) as an instance of Metric Dimension for some appropriate value of k.

2.3.2 Correctness of the Reduction

We give an informal description of the proof of correctness of the reverse direction here. Fix $\delta \in \{\alpha, \beta, \gamma\}$. For all $i \in [n]$, the only vertices that can resolve the critical pair $\langle x_i^{\delta,\circ}, x_i^{\delta,\star}\rangle$ are the vertices in $\{x_i^{\delta,\circ}, x_i^{\delta,\star}\} \cup \{t_{2i}^\delta, f_{2i-1}^\delta\}$. This fact and the budget k ensure that any resolving set of G contains exactly one vertex from $\{t_{2i}^\delta, f_{2i-1}^\delta\} \cup \{x_i^{\delta,\circ}, x_i^{\delta,\star}\}$ for all $i \in [n]$. This naturally corresponds to an assignment of the variable x_i^δ if a vertex from $\{t_{2i}^\delta, f_{2i-1}^\delta\}$ is in the resolving set. However, if a vertex from $\{x_i^{\delta,\circ}, x_i^{\delta,\star}\}$ is in the resolving set, then we can see this as giving an arbitrary assignment to the variable x_i^δ. Suppose the clause C_q contains the variable x_i^δ as a positive literal. By the construction, every vertex in V^δ that is adjacent to t_{2i}^δ is *not* adjacent to $c_q^\star$. However, c_q° is adjacent to every vertex in V^δ. Hence, $d(t_{2i}^\delta, c_q^\circ) = 2$, whereas $d(t_{2i}^\delta, c_q^\star) > 2$. Thus, t_{2i}^δ resolves the critical pair $\langle c_q^\circ, c_q^\star\rangle$. Consider any other vertex in A^δ, say t_{2j}^δ. Since set-rep$(2i)$ is not a subset of set-rep$(2j)$ (as both have the same cardinality), there is at least one integer, say p', in set-rep$(2j) \setminus$ set-rep$(2i)$. The vertex $v_{p'}^\delta \in V^\delta$ is adjacent to t_{2j}^δ, c_q°, and $c_q^\star$. Hence, t_{2j}^δ cannot resolve the critical pair $\langle c_q^\circ, c_q^\star\rangle$ as both these vertices are at distance 2 from it. Also, as ψ is an instance of

3-Partitioned-3-SAT, C_q contains at most one variable in X^δ, which is x_i^δ in this case. This also helps to encode the fact that at most one vertex from A^δ should be able to resolve the critical pair $\langle c_q^\circ, c_q^\star \rangle$. Since vertices in X^δ cannot resolve critical pairs $\langle c_q^\circ, c_q^\star \rangle$ in C, then finding a resolving set in G corresponds to finding a satisfying assignment for ψ.

2.3.3 Lower Bounds Obtained from the Reduction

Let $Z = \{V^\delta \cup X^+ \mid X \in \{X^\delta, A^\delta, V^\delta, C\}, \delta \in \{\alpha, \beta, \gamma\}\}$. Note that $|Z| = \mathcal{O}(\log(n))$ and $G - Z$ is a collection of P_3's and isolated vertices. Hence, $\mathtt{tw}(G)$, $\mathtt{fvs}(G)$, and $\mathtt{td}(G)$ are upper bounded by $\mathcal{O}(\log(n))$. Also, G has constant diameter. Thus, if there is an algorithm for Metric Dimension that runs in time $2^{f(\mathtt{diam})^{o(\mathtt{tw})}}$ (or $2^{f(\mathtt{diam})^{o(\mathtt{fvs})}}$ or $2^{f(\mathtt{diam})^{o(td)}}$), then there is an algorithm solving 3-Partitioned-3-SAT in time $2^{o(n)}$, contradicting the ETH.

2.4 High-Level Description of the Dynamic Programming Algorithm for Metric Dimension

The aim of this subsection is to give an informal description of how we prove the upper bound concerning the parameter $\mathtt{tw} + \mathtt{diam}$ for Metric Dimension. To this end, we give a dynamic programming algorithm on a tree decomposition for Metric Dimension. The algorithm is inspired by the one from [7] for chordal graphs, though there are some non-trivial differences. We will assume that a tree decomposition of the input graph G of width w is given to us. Note that one can compute a tree decomposition of width $w \leq 2\mathtt{tw}(G) + 1$ in time $2^{\mathcal{O}(\mathtt{tw}(G))}n$ [42], and it can be transformed into a *nice tree decomposition* of the same width with $\mathcal{O}(wn)$ bags in time $\mathcal{O}(w^2 n)$ [40]. We now give a high-level overview of the dynamic programming algorithm used to prove the following theorem.

▶ **Theorem 2.** *Metric Dimension admits an algorithm running in time* $2^{\mathtt{diam}^{\mathcal{O}(\mathtt{tw})}} \cdot n^{\mathcal{O}(1)}$.

In [7], as the diameter of the graph was unbounded, it was crucial to restrict the computations for each step of the dynamic programming to vertices "not too far" from the current bag. This was possible due to the metric properties of chordal graphs. In our case, as we consider the diameter of the graph as a parameter, we do not need such restrictions, which makes the proof a little bit simpler.

We now give an intuitive description of the dynamic programming scheme. At each step of the algorithm, we consider a bounded number of *solution types*, depending on the properties of the solution vertices with respect to the current bag. At a given dynamic programming step, we will assume that the current solution resolves all vertex pairs in G_i. Such a vertex pair may be resolved by a vertex from $G - G_i$, or by a vertex in G_i itself.

Any bag X_i of the tree decomposition whose node i lies on a path between two join nodes in T, forms a separator of G: there are no edges between the vertices of $G_i - X_i$ and $G - G_i$. For a vertex v not in X_i, we consider its distance-vector to the vertices of X_i; the distance-vectors induce an equivalence relation on the vertices of $G - X_i$, whose classes we call *X_i-classes*. Consider the two subgraphs G_i and $G - G_i$. Any two solution vertices x, y from $G - G_i$ that are in the same X_i-class, resolve the exact same pairs of vertices from G_i. Thus, for this purpose, it is irrelevant whether x or y will be in a resolving set, and it is sufficient to know that a vertex of their X_i-class will eventually be chosen. In this way, one can check whether a vertex pair from G_i is resolved by a solution vertex of $G - G_i$.

The same idea is used to "remember" the previously computed solution: it is sufficient to remember the X_i-classes of the vertices in the previously computed resolving set, rather than the vertices themselves.

It is slightly more delicate to make sure that vertex pairs in G_i are resolved in the case where such a pair is resolved by a vertex in G_i. Indeed, this must be ensured, in particular when processing a join node i, for vertex pairs belonging to bags in the two sub-trees corresponding to the children i_1, i_2 of i. Such pairs may be resolved by four types of solution vertices: from $G - G_i$, X_i, $G_{i_1} - X_i$, or $G_{i_2} - X_i$. To ensure this, the dynamic programming scheme makes sure that, at each step, for any possible pair C_1, C_2 of X_i-classes, all vertex pairs $\langle u, v\rangle$ consisting of a vertex u of G_i with class C_1 and a vertex v of $G - G_i$ with class C_2 are resolved. The crucial step here is that when a new vertex v is introduced (i.e., added to a bag X_i to form $X_{i'}$), depending on its X_i-class, it must be made sure that it is resolved from all other vertices depending on their X_i-classes, as described above. To ensure that v is distinguished from all other vertices of G_i, we keep track of vertex pairs of $G_i \times (G - G_i)$ that are already resolved by the partial solution, and enforce that, when processing bag $X_{i'}$, for every vertex x of G_i, the pair $\langle x, v\rangle$ is already resolved. As v belongs to the new bag $X_{i'}$, we know its distances to all resolving vertices (indeed, $X_{i'}$-classes of solution vertices can be computed from their X_i-classes), and thus, the information can be updated accurately.

For a bag X_i and a vertex v not in X_i, the number of possible distance vectors to the vertices of X_i is at most $\mathtt{diam}(G)^{|X_i|}$. Thus, a solution for bag X_i will consist of: (i) the subset of vertices of X_i selected in the solution; (ii) a subset of the $\mathtt{diam}(G)^{|X_i|}$ possible vectors to denote the X_i-classes from which the currently computed solution (for G_i) contains at least one vertex in the resolving set; (iii) a subset of the $\mathtt{diam}(G)^{|X_i|}$ possible vectors denoting the X_i-classes from which the future solution needs at least one vertex of $G - G_i$ in the resolving set; (iv) a subset of the $\mathtt{diam}(G)^{|X_i|} \times \mathtt{diam}(G)^{|X_i|}$ possible pairs of vectors representing the X_i-classes of the pairs of vertices in $G_i \times (G - G_i)$ that are already resolved by the partial solution.

3 Geodetic Set: Lower Bound Regarding Diameter plus Treewidth

The aim of this section is to prove the following theorem.

▶ **Theorem 3.** *Unless the ETH fails,* Geodetic Set *does not admit an algorithm running in time* $2^{f(\mathtt{diam})^{o(\mathtt{tw})}} \cdot n^{\mathcal{O}(1)}$ *for any computable function* $f : \mathbb{N} \mapsto \mathbb{N}$.

Here, we present a different reduction from 3-Partitioned-3-SAT to Geodetic Set. The reduction takes as input an instance ψ of 3-Partitioned-3-SAT on $3n$ variables and returns (G, k) as an instance of Geodetic Set such that $\mathtt{tw}(G) = \mathcal{O}(\log(n))$ and $\mathtt{diam}(G) = \mathcal{O}(1)$. We rely on the tool of set representation from Section 2.2, that, for convenience, we reintroduce in the context of Geodetic Set in the next subsection.

3.1 Preliminary Tool: Set Representation

For a positive integer p, define $\mathcal{F}_p$ as the collection of subsets of $[2p]$ that contains exactly p integers. We critically use the fact that no set in $\mathcal{F}_p$ is contained in any other set in $\mathcal{F}_p$ (such a collection of sets is called a *Sperner family*). Let ℓ be a positive integer such that $\ell \le \binom{2p}{p}$. We define $\textsf{set-rep} : [\ell] \mapsto \mathcal{F}_p$ as a one-to-one function by arbitrarily assigning a set in $\mathcal{F}_p$ to an integer in $[\ell]$. By the asymptotic estimation of the central binomial coefficient, $\binom{2p}{p} \sim \frac{4^p}{\sqrt{\pi \cdot p}}$ [36]. To get the upper bound of p, we scale down the asymptotic function and have $\ell \le \frac{4^p}{2^p} = 2^p$. Thus, $p = \mathcal{O}(\log \ell)$.

We will apply the existence of such a function in the context of Geodetic Set. Suppose we have a “large” collection of vertices, say $A = \{a_1, a_2, \ldots, a_\ell\}$, and a “large” collection of vertices $C = \{c_1, c_2, \ldots, c_m\}$. Moreover, we are given a function $\phi : [m] \mapsto [\ell]$. The basic

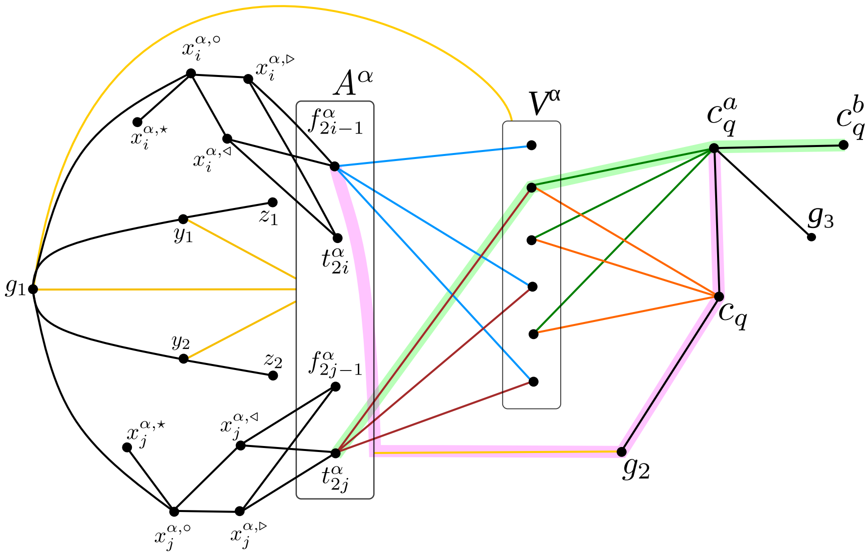

Figure 5 Overview of the reduction. We only draw A^α and V^α here, as A^β, A^γ, V^β, and V^γ are similar. The yellow lines joining g_1, g_2, y_1, and y_2 to sets indicate that the corresponding vertex is adjacent to all the vertices of the corresponding set. Suppose that f_{2i-1}^α and t_{2j}^α are in the geodetic set and $\overline{x}_i$ appears in the clause c_q. The thick green path is a shortest path between t_{2j}^α and c_q^b which does not cover c_q. The thick violet path plus the edge (c_q^a, c_q^b) is a shortest path between f_{2i-1}^α and c_q^b covering c_q.

idea is to design gadgets such that c_q is only covered by the shortest path from $a_{\phi(q)} \in A$ to c_q^b (c_q^b is forced to be chosen in the geodetic set) for any $q \in [m]$, while keeping the treewidth of this part of the graph of order $\mathcal{O}(\log(|A|))$. To do so, we create a "small" intermediate set V (of size $\mathcal{O}(\log(|A|))$) through which will go the shortest paths between vertices in A and C, and we connect a_i to the vertices of V corresponding to the bit-representation of set-rep(i), and c_q (with $i = \phi(q)$) to all the other vertices of V. In this way, the construction will ensure that c_q is covered by a shortest path between $a_{\phi(q)}$ and c_q^b, but is not covered by any other shortest path between a vertex of A and a vertex of C. We give the details in the following subsection.

3.2 Reduction

Consider an instance ψ of 3-Partitioned-3-SAT, with $X^\alpha, X^\beta, X^\gamma$ the partition of the variable set. From ψ, we construct the graph G as follows. We describe the construction of X^α, with the constructions for X^β and X^γ being analogous. See Figure 5 for an illustration. We rename the variables in X^α to x_i^α for $i \in [n]$.

- For every variable x_i^α, we add the vertices t_{2i}^α and f_{2i-1}^α. Formally, $A^\alpha = \{t_{2i}^\alpha, f_{2i-1}^\alpha \mid i \in [n]\}$, and hence, $|A^\alpha| = 2n$.
- For every variable x_i^α, we add four vertices: $x_i^{\alpha,\triangleleft}, x_i^{\alpha,\triangleright}, x_i^{\alpha,\circ}, x_i^{\alpha,\star}$. We make $x_i^{\alpha,\triangleleft}$ and $x_i^{\alpha,\triangleright}$ adjacent to both t_{2i}^α and f_{2i-1}^α. We make $x_i^{\alpha,\circ}$ adjacent to both $x_i^{\alpha,\triangleleft}$ and $x_i^{\alpha,\triangleright}$. We make $x_i^{\alpha,\star}$ adjacent to $x_i^{\alpha,\circ}$.
- We add the vertices y_1, y_2, z_1, z_2. We make y_1 and y_2 adjacent to every vertex of A^α. We make y_i adjacent to z_i for $i \in \{1,2\}$. Note that y_1, y_2, z_1, z_2 are common to X^β and X^γ.
- We add the vertex g_1 and make it adjacent to y_1, y_2, and $x_i^{\alpha,\circ}$ for each $i \in [n]$. Note that g_1 is common to X^β and X^γ. We add edges between g_1 and every vertex of A^α.

- Let p be the smallest positive integer such that $2n \leq \binom{2p}{p}$. In particular, $p = \mathcal{O}(\log n)$. We add a *validation portal*, a clique on $2p$ vertices, denoted by $V^\alpha = \{v_1^\alpha, v_2^\alpha, \dots, v_{2p}^\alpha\}$. For each $\delta \in \{\alpha, \beta, \gamma\}$, we add edges between g_1 and every vertex of V^δ.
- For every clause C_q in ψ, we introduce three vertices: c_q, c_q^a, c_q^b. We add the edges (c_q, c_q^a) and (c_q^a, c_q^b).
- Define $\textsf{set-rep} : [2n] \mapsto \mathcal{F}_p$ as an arbitrary injective function, where $\mathcal{F}_p$ is the Sperner family (and p is as defined two items above). Add the edge $(t_{2i}^\alpha, v_{p'}^\alpha)$ for every $p' \in \textsf{set-rep}(2i)$ and the edge $(f_{2i-1}^\alpha, v_{p'}^\alpha)$ for every $p' \in \textsf{set-rep}(2i-1)$. If the variable x_i^α appears positively in the clause C_q, then we add the edges $(c_q, v_{p'}^\alpha)$ and $(c_q^a, v_{p'}^\alpha)$ for every $p' \in [2p] \setminus \textsf{set-rep}(2i)$. If the variable x_i^α appears negatively in the clause C_q, then we add the edges $(c_q, v_{p'}^\alpha)$ and $(c_q^a, v_{p'}^\alpha)$ for every $p' \in [2p] \setminus \textsf{set-rep}(2i-1)$.
- Add a vertex g_2 and make g_2 adjacent to every vertex of A^α and every vertex of $\{c_q : q \in [m]\}$. Note that g_2 is common to X^β and X^γ.
- Add a vertex g_3 and make it adjacent to every vertex of $\{c_q^a : q \in [m]\}$. Note that g_3 and the vertices of $\{c_q, c_q^a, c_q^b : q \in [m]\}$ are common to X^β and X^γ.

This concludes the construction of G. The reduction returns (G, k) as an instance of Geodetic Set where $k = 6n + m + 2$.

3.3 Correctness of the Reduction

Suppose, given an instance ψ of 3-Partitioned-3-SAT, that the reduction above returns (G, k) as an instance of Geodetic Set.

▶ **Lemma 4.** *If ψ is a satisfiable 3-Partitioned-3-SAT formula, then G admits a geodetic set of size k.*

Proof. Suppose that $\pi : X^\alpha \cup X^\beta \cup X^\gamma \mapsto \{\texttt{True}, \texttt{False}\}$ is a satisfying assignment for ψ. We construct a geodetic set S of size k for G using this assignment.

For every $\delta \in \{\alpha, \beta, \gamma\}$ and $i \in [n]$, if $\pi(x_i^\delta) = \texttt{True}$, then let $t_{2i}^\delta \in S$, and otherwise, $f_{2i-1}^\delta \in S$. We also put z_1, z_2, $x_i^{\delta,\star}$, and c_q^b into S for all $i \in [n]$, $\delta \in \{\alpha, \beta, \gamma\}$, and $q \in [m]$. Note that $|S| = k$.

Now, we show that S is indeed a geodetic set of G. First, y_1, y_2, z_1, z_2, g_1, and all the vertices of $A^\alpha, A^\beta, A^\gamma$ are covered by a shortest path between z_1 and z_2. Then, for each $\delta \in \{\alpha, \beta, \gamma\}$ and $i \in [n]$, $x_i^{\delta,\lhd}$, $x_i^{\delta,\rhd}$, $x_i^{\delta,\circ}$, and $x_i^{\delta,\star}$ are covered by a shortest path between $S \cap \{t_{2i}^\delta, f_{2i-1}^\delta\}$ and $x_i^{\delta,\star}$. The vertex g_3 is covered by any shortest path between c_q^b and $c_{q'}^b$, where C_q and $C_{q'}$ are two clauses of ψ. Suppose that $\pi(x_i^\delta)$, for some $i \in [n]$ and $\delta \in \{\alpha, \beta, \gamma\}$, satisfies some clause C_q. By our construction, if x_i^δ appears positively (negatively, respectively) in C_q, then t_{2i}^δ (f_{2i-1}^δ, respectively) and c_q^b are at distance four since t_{2i}^δ (f_{2i-1}^δ, respectively) and c_q^a have no common neighbor in V^δ. Moreover, there is a shortest path from t_{2i}^δ (f_{2i-1}^δ, respectively) to c_q^b of length four, covering g_2, c_q, c_q^a, and c_q^b; there is also a shortest path from t_{2i}^δ (f_{2i-1}^δ, respectively) to c_q^b of length four, covering $v_j^\delta, v_h^\delta, c_q^a$, and c_q^b, where $v_j^\delta \in V^\delta$ is a vertex adjacent to t_{2i}^δ (f_{2i-1}^δ, respectively) and v_h^δ is any vertex of V^δ that is not adjacent to t_{2i}^δ (f_{2i-1}^δ, respectively). Thus, every vertex of V^δ for $\delta \in \{\alpha, \beta, \gamma\}$ is covered by a shortest path between two vertices of S. Since every clause of ψ is satisfied by π, it follows that every vertex of $\{c_q, c_q^a, c_q^b : q \in [m]\}$ is covered by a shortest path between two vertices of S. As a result, S is a geodetic set of G. ◀

▶ **Lemma 5.** *If G admits a geodetic set of size k, then ψ is a satisfiable 3-Partitioned-3-SAT formula.*

Proof. Suppose that G has a geodetic set S of size at most k. Since they all have degree 1, z_1, z_2, $x_i^{\delta,\star}$, and c_q^b for all $i \in [n]$, $\delta \in \{\alpha, \beta, \gamma\}$, and $q \in [m]$ must be in any geodetic set S of G.

▷ Claim 6. For each $i \in [n]$ and $\delta \in \{\alpha, \beta, \gamma\}$, exactly one of t_{2i}^{δ} and f_{2i-1}^{δ} must be in S.

Proof. Since S is a geodetic set, for each $i \in [n]$ and $\delta \in \{\alpha, \beta, \gamma\}$ $x_i^{\delta,\lhd}$ and $x_i^{\delta,\rhd}$ must be covered by shortest paths between two vertices of S. If $t_{2i}^{\delta} \in S$ ($f_{2i-1}^{\delta} \in S$, respectively), $x_i^{\delta,\lhd}$ and $x_i^{\delta,\rhd}$ are covered by shortest paths between $t_{2i}^{\delta} \in S$ ($f_{2i-1}^{\delta} \in S$, respectively) and $x_i^{\delta,\star}$. Suppose that, for some $i' \in [n]$ and $\delta' \in \{\alpha, \beta, \gamma\}$, neither of $t_{2i'}^{\delta'}$ and $f_{2i'-1}^{\delta'}$ is in S. Moreover, if neither of $x_{i'}^{\delta',\lhd}$ and $x_{i'}^{\delta',\rhd}$ is in S, then, due to the edges incident with g_1, no vertices in S have a shortest path containing any of these two vertices. Similarly, if only one of $x_{i'}^{\delta',\lhd}$ and $x_{i'}^{\delta',\rhd}$ is in S, then the other is not covered by S. Thus, if neither of $t_{2i'}^{\delta'}$ and $f_{2i'-1}^{\delta'}$ is in S, then both $x_{i'}^{\delta',\lhd}$ and $x_{i'}^{\delta',\rhd}$ must be in S. Since $k - |\{z_1, z_2\} \cup \{x_i^{\delta,\star} : i \in [n], \delta \in \{\alpha, \beta, \gamma\}\} \cup \{c_q^b : q \in [m]\}| = 3n$, we conclude that exactly one of t_{2i}^{δ} and f_{2i-1}^{δ} must be in S for each $i \in [n]$ and $\delta \in \{\alpha, \beta, \gamma\}$. ◁

By Claim 6 and earlier arguments, we now have that $|S| = k$.

▷ Claim 7. For each $q \in [m]$, the vertex c_q is covered either by a shortest path between c_q^b and t_{2i}^{δ}, where the variable x_i^{δ} appears positively in the clause C_q, or by a shortest path between c_q^b and f_{2i-1}^{δ}, where the variable x_i^{δ} appears negatively in the clause C_q. Moreover, c_q is covered by no other type of shortest path between two vertices in S.

Proof. By the construction of G, if the variable x_i^{δ} appears positively in the clause C_q, then there is a shortest path from t_{2i}^{δ} to c_q^b of length four covering g_2, c_q, c_q^a, and c_q^b. If the variable x_i^{δ} appears negatively in the clause C_q, then there is a shortest path from f_{2i-1}^{δ} to c_q^b of length four covering g_2, c_q, c_q^a, and c_q^b.

Next, we show that c_q is not covered by any shortest path between any other two vertices of S. We can check that c_q is not covered by any of the shortest paths between z_1 and z_2, between z_j ($j \in \{1, 2\}$) and $x_i^{\delta,\star}$ ($i \in [n], \delta \in \{\alpha, \beta, \gamma\}$), and between z_j ($j \in \{1, 2\}$) and $S \cap \{t_{2i}^{\delta}, f_{2i-1}^{\delta}\}$ ($i \in [n], \delta \in \{\alpha, \beta, \gamma\}$). Note that any shortest path from z_j ($j \in \{1, 2\}$) to c_q^b ($q \in [m]$) is of length five, covering y_j, some vertex of A^{δ} ($\delta \in \{\alpha, \beta, \gamma\}$), some vertex of V^{δ}, c_q^a, and c_q^b.

We can check that c_q is not covered by any of the shortest paths between $x_i^{\delta,\star}$ and $x_{i'}^{\delta',\star}$ ($i, i' \in [n], \delta, \delta' \in \{\alpha, \beta, \gamma\}$), and between $x_i^{\delta,\star}$ and $S \cap \{t_{2i'}^{\delta'}, f_{2i'-1}^{\delta'}\}$ ($i, i' \in [n], \delta, \delta' \in \{\alpha, \beta, \gamma\}$). Note that any shortest path from $x_i^{\delta,\star}$ ($i \in [n], \delta \in \{\alpha, \beta, \gamma\}$) to c_q^b ($q \in [m]$) is of length five, covering $x_i^{\delta,\circ}$, g_1, some vertex of V^{δ}, c_q^a, and c_q^b.

Note that any shortest path between c_q^b and $c_{q'}^b$ ($q, q' \in [m]$) is of length four, covering c_q^a, g_3, and $c_{q'}^a$.

We can check that c_q is not covered by any shortest paths between $S \cap \{t_{2i}^{\delta}, f_{2i-1}^{\delta}\}$ and $S \cap \{t_{2i'}^{\delta'}, f_{2i'-1}^{\delta'}\}$ ($i, i' \in [n], \delta, \delta' \in \{\alpha, \beta, \gamma\}$).

If the variable x_i^{δ} does not appear positively in the clause C_q, then any shortest path between c_q^b and t_{2i}^{δ} is of length three (because c_q^a and t_{2i}^{δ} have a common neighbour in V^{δ}), covering some vertex of V^{δ} and c_q^a, but not c_q. Similarly, if x_i^{δ} does not appear negatively in C_q, then any shortest path between c_q^b and f_{2i-1}^{δ} is of length three and does not cover c_q.

By the case analysis above, the claim is true. ◁

By Claim 6, exactly one vertex of t_{2i}^{δ} and f_{2i-1}^{δ} belongs to S for each $i \in [n]$ and $\delta \in \{\alpha, \beta, \gamma\}$. We define an assignment π to the variables of ψ as follows. For each $i \in [n]$ and $\delta \in \{\alpha, \beta, \gamma\}$, if $t_{2i}^{\delta} \in S$, then $\pi(x_i^{\delta}) = \texttt{True}$. Otherwise, $\pi(x_i^{\delta}) = \texttt{False}$. Since S is a geodetic set for G, every vertex c_q ($q \in [m]$) is covered by a shortest path between two vertices of S. By Claim 7, every vertex c_q ($q \in [m]$) is covered by a shortest path between $S \cap \{t^\delta_{2i}, f^\delta_{2i-1}\}$ and c^b_q, where the variable x^δ_i appears in the clause C_q. It follows that every clause C_q is satisfied by $\pi(x^\delta_i)$. As a result, ψ is a satisfiable 3-Partitioned-3-SAT formula. ◀

Proof of Theorem 3. First, it is not hard to check that the diameter of G is at most 5. Then, let $X = V^\alpha \cup V^\beta \cup V^\gamma \cup \{g_1, g_2, g_3, y_1, y_2\}$. We can check that every component of $G \setminus X$ has at most six vertices and $|X| = \mathcal{O}(\log n)$. Thus, the treewidth $\mathtt{tw}(G)$ – in fact, even the treedepth $\mathtt{td}(G)$ – of G is bounded by $\mathcal{O}(\log n)$. By the description of the reduction, it takes polynomial time to compute the reduced instance. Hence, if there is an algorithm for Geodetic Set that runs in time $2^{f(\mathtt{diam})^{o(\mathtt{tw})}}$ (or $2^{f(\mathtt{diam})^{o(td)}}$), then there is an algorithm running in time $2^{o(n)}$ for 3-Partitioned-3-SAT, which contradicts the ETH. ◀

4 Conclusion

We have shown (under the ETH) that three natural metric-based graph problems, Metric Dimension, Geodetic Set, and Strong Metric Dimension, exhibit tight (double-) exponential running times for the standard structural parameterizations by treewidth and vertex cover number. This includes tight double-exponential running times for treewidth plus diameter (Metric Dimension and Geodetic Set) and for vertex cover (Strong Metric Dimension).

Such tight double-exponential running times for FPT structural paramaterizations of graph problems had previously been observed only for counting problems and problems complete for classes above NP. Thus, surprisingly, our results show that some natural problems can be in NP and still exhibit such a behavior.

It would be interesting to see whether this phenomenon holds for other graph problems in NP, and for other structural parameterizations. Perhaps one can determine certain properties shared by these metric-based graph problems, that imply such running times, with the goal of generalizing our approach to a broader class of problems. In particular, concerning the general versatile technique that we designed to obtain the double-exponential lower bounds, it would be intriguing to see for which other problems in NP our technique works.

In fact, after this paper appeared online, our technique was successfully applied to an NP-complete problem in machine learning [11] (for vc) as well as NP-complete identification problems [10] (for tw).

References

1 A. Achilleos, M. Lampis, and V. Mitsou. Parameterized modal satisfiability. *Algorithmica*, 64(1):38–55, 2012.

2 J. Ahn, L. Jaffke, O. Kwon, and P. T. Lima. Well-partitioned chordal graphs. *Discrete Mathematics*, 345(10):112985, 2022.

3 F. Barbero, L. Isenmann, and J. Thiebaut. On the distance identifying set meta-problem and applications to the complexity of identifying problems on graphs. *Algorithmica*, 82(8):2243–2266, 2020. doi:10.1007/S00453-020-00674-X.

4 R. Belmonte, F. V. Fomin, P. A. Golovach, and M. S. Ramanujan. Metric dimension of bounded tree-length graphs. *SIAM J. Discrete Math.*, 31(2):1217–1243, 2017.

5 B. Bergougnoux, O. Defrain, and F. Mc Inerney. Enumerating minimal solution sets for metric graph problems. In *Proc. of the 50th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2024)*, Lecture Notes in Computer Science. Springer, 2024.

6 E. Bonnet and N. Purohit. Metric dimension parameterized by treewidth. *Algorithmica*, 83:2606–2633, 2021.

7 N. Bousquet, Q. Deschamps, and A. Parreau. Metric dimension parameterized by treewidth in chordal graphs. In *Proc. of the 49th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2023)*, volume 14093 of *Lecture Notes in Computer Science*, pages 130–142. Springer, 2023.

8 D. Chakraborty, S. Das, F. Foucaud, H. Gahlawat, D. Lajou, and B. Roy. Algorithms and complexity for geodetic sets on planar and chordal graphs. In *31st International Symposium on Algorithms and Computation (ISAAC 2020)*, volume 181 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020.

9 D. Chakraborty, F. Foucaud, H. Gahlawat, S. K. Ghosh, and B. Roy. Hardness and approximation for the geodetic set problem in some graph classes. In *Proc. of the 6th International Conference on Algorithms and Discrete Applied Mathematics (CALDAM 2020)*, volume 12016 of *Lecture Notes in Computer Science*, pages 102–115, Cham, 2020. Springer.

10 D. Chakraborty, F. Foucaud, D. Majumdar, and P. Tale. Tight (double) exponential bounds for identification problems: Locating-dominating set and test cover, 2024. arXiv:2402.08346.

11 J. Chalopin, V. Chepoi, F. Mc Inerney, and S. Ratel. Non-clashing teaching maps for balls in graphs. *Arxiv:2309.02876*, 2023.

12 H. Chen. Quantified constraint satisfaction and bounded treewidth. In *Proc. of the 16th Eureopean Conference on Artificial Intelligence, ECAI'2004*, pages 161–165. IOS Press, 2004.

13 B. Courcelle. The monadic second-order logic of graphs. I. recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990.

14 M. Cygan, D. Marx, M. Pilipczuk, and M. Pilipczuk. Hitting forbidden subgraphs in graphs of bounded treewidth. *Inf. Comput.*, 256(C):62–82, 2017.

15 M. Cygan, J. Nederlof, M. Pilipczuk, M. Pilipczuk, J. M. M. van Rooij, and J. O. Wojtaszczyk. Solving connectivity problems parameterized by treewidth in single exponential time. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS 2011)*, pages 150–159, 2011.

16 M. Cygan, M. Pilipczuk, and M. Pilipczuk. Known algorithms for edge clique cover are probably optimal. *SIAM J. Comput.*, 45(1):67–83, 2016. doi:10.1137/130947076.

17 B. DasGupta and N. Mobasheri. On optimal approximability results for computing the strong metric dimension. *Discrete Applied Math.*, 221:18–24, 2017.

18 J. Díaz, O. Pottonen, M. J. Serna, and E. J. van Leeuwen. Complexity of metric dimension on planar graphs. *J. Comput. Syst. Sci.*, 83(1):132–158, 2017.

19 M. C. Dourado, F. Protti, D. Rautenbach, and J. L. Szwarcfiter. Some remarks on the geodetic number of a graph. *Discrete Mathematics*, 310(4):832–837, 2010.

20 T. Ekim, A. Erey, P. Heggernes, P. van't Hof, and D. Meister. Computing minimum geodetic sets of proper interval graphs. In *Proc. of the 10th Latin American Symposium on Theoretical Informatics (LATIN 2012)*, volume 7256 of *Lecture Notes in Computer Science*, pages 279–290. Springer, 2012.

21 D. Eppstein. Metric dimension parameterized by max leaf number. *Journal of Graph Algorithms and Applications*, 19(1):313–323, 2015.

22 L. Epstein, A. Levin, and G. J. Woeginger. The (weighted) metric dimension of graphs: Hard and easy cases. *Algorithmica*, 72(4):1130–1171, 2015.

23 J. K. Fichte, M. Hecher, M. Morak, P. Thier, and S. Woltran. Solving projected model counting by utilizing treewidth and its limits. *Artif. Intell.*, 314:103810, 2023.

24 J. K. Fichte, M. Hecher, M. Morak, and S. Woltran. Exploiting treewidth for projected model counting and its limits. In *Theory and Applications of Satisfiability Testing - SAT 2018 - 21st International Conference, SAT 2018, Proc.*, volume 10929 of *Lecture Notes in Computer Science*, pages 165–184. Springer, 2018.

25 J. K. Fichte, M. Hecher, and A. Pfandler. Lower bounds for QBFs of bounded treewidth. In *LICS '20: 35th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 410–424. ACM, 2020.

26 J. Focke, F. Frei, S. Li, D. Marx, P. Schepper, R. Sharma, and K. Wegrzycki. Hitting meets packing: How hard can it be? *CoRR*, abs/2402.14927, 2024. `doi:10.48550/arXiv.2402.14927`.

27 F. V. Fomin, P. A. Golovach, D. Lokshtanov, S. Saurabh, and M. Zehavi. Clique-width III: hamiltonian cycle and the odd case of graph coloring. *ACM Trans. Algorithms*, 15(1):9:1–9:27, 2019.

28 F. Foucaud, G. B. Mertzios, R. Naserasr, A. Parreau, and P. Valicov. Identification, location-domination and metric dimension on interval and permutation graphs. II. Algorithms and complexity. *Algorithmica*, 78(3):914–944, 2017.

29 E. Galby, L. Khazaliya, F. Mc Inerney, R. Sharma, and P. Tale. Metric dimension parameterized by feedback vertex set and other structural parameters. *SIAM J. Discrete Math.*, 37(4):2241–2264, 2023.

30 M. R. Garey and D. S. Johnson. *Computers and Intractability - A guide to NP-completeness*. W.H. Freeman and Company, 1979.

31 T. Gima, T. Hanaka, M. Kiyomi, Y. Kobayashi, and Y. Otachi. Exploring the gap between treedepth and vertex cover through vertex integrity. *Theoretical Computer Science*, 918:60–76, 2022. `doi:10.1016/J.TCS.2022.03.021`.

32 T. Hanaka, H. Köhler, and M. Lampis. Core stability in additively separable hedonic games of low treewidth, 2024. `arXiv:2402.10815`.

33 F. Harary, E. Loukakis, and C. Tsouros. The geodetic number of a graph. *Mathematical and Computer Modelling*, 17(11):89–95, 1993.

34 F. Harary and R. A. Melter. On the metric dimension of a graph. *Ars Combinatoria*, 2:191–195, 1976.

35 S. Hartung and A. Nichterlein. On the parameterized and approximation hardness of metric dimension. In *Proc. of the 28th Conference on Computational Complexity, CCC 2013*, pages 266–276. IEEE Computer Society, 2013.

36 A. Ian. *Combinatorics of Finite Sets*. Oxford University Press, 1987.

37 K. Jansen, KM. Klein, and A. Lassota. The double exponential runtime is tight for 2-stage stochastic ILPs. *Math. Program.*, 197:1145–1172, 2023.

38 M. M. Kanté and L. Nourine. Polynomial time algorithms for computing a minimum hull set in distance-hereditary and chordal graphs. *SIAM J. Discrete Math.*, 30(1):311–326, 2016.

39 L. Kellerhals and T. Koana. Parameterized complexity of geodetic set. *Journal of Graph Algorithms and Applications*, 26(4):401–419, 2022.

40 T. Kloks. *Treewidth, Computations and Approximations*. Springer, 1994.

41 D. Knop, M. Pilipczuk, and M. Wrochna. Tight complexity lower bounds for integer linear programming with few constraints. *ACM Trans. Comput. Theory*, 12(3):19:1–19:19, 2020.

42 T. Korhonen. A single-exponential time 2-approximation algorithm for treewidth. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS 2021)*, pages 184–192, 2022.

43 L. Kowalik, A. Lassota, K. Majewski, M. Pilipczuk, and M. Sokołowski. Detecting points in integer cones of polytopes is double-exponentially hard. In *2024 Symposium on Simplicity in Algorithms (SOSA)*, pages 279–285, 2024.

44 M. Künnemann, F. Mazowiecki, L. Schütze, H. Sinclair-Banks, and K. Węgrzycki. Coverability in VASS Revisited: Improving Rackoff's Bound to Obtain Conditional Optimality. In *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, volume 261 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 131:1–131:20, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

45 D. Kuziak, M. L. Puertas, J. A. Rodríguez-Velázquez, and I. G. Yero. Strong resolving graphs: The realization and the characterization problems. *Discrete Applied Math.*, 236:270–287, 2018.

46 M. Lampis, N. Melissinos, and M. Vasilakis. Parameterized max min feedback vertex set. In *48th International Symposium on Mathematical Foundations of Computer Science, MFCS 2023*, volume 272 of *LIPIcs*, pages 62:1–62:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023.

47 M. Lampis, S. Mengel, and V. Mitsou. QBF as an alternative to courcelle's theorem. In *Theory and Applications of Satisfiability Testing - SAT 2018 - 21st International Conference, SAT 2018*, volume 10929 of *Lecture Notes in Computer Science*, pages 235–252. Springer, 2018.

48 M. Lampis and V. Mitsou. Treewidth with a quantifier alternation revisited. In *12th International Symposium on Parameterized and Exact Computation (IPEC 2017)*, volume 89 of *LIPIcs*, pages 26:1–26:12, 2017.

49 S. Li and M. Pilipczuk. Hardness of metric dimension in graphs of constant treewidth. *Algorithmica*, 84(11):3110–3155, 2022.

50 D. Lokshtanov, D. Marx, and S. Saurabh. Slightly superexponential parameterized problems. *SIAM J. Comput.*, 47(3):675–702, 2018.

51 D. Lokshtanov, S. Saurabh, S. Suri, and J. Xue. An ETH-tight algorithm for multi-team formation. In *41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2021*, volume 213 of *LIPIcs*, pages 28:1–28:9. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

52 D. Marx and V. Mitsou. Double-exponential and triple-exponential bounds for choosability problems parameterized by treewidth. In *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, volume 55 of *LIPIcs*, pages 28:1–28:15, 2016.

53 M. Mezzini. Polynomial time algorithm for computing a minimum geodetic set in outerplanar graphs. *Theoretical Computer Science*, 745:63–74, 2018.

54 O. R. Oellermann and J. Peters-Fransen. The strong metric dimension of graphs and digraphs. *Discrete Applied Mathematics*, 155(3):356–364, 2007.

55 M. Pilipczuk. Problems parameterized by treewidth tractable in single exponential time: A logical approach. In *Mathematical Foundations of Computer Science 2011 - 36th International Symposium, MFCS 2011, Proc.*, volume 6907 of *Lecture Notes in Computer Science*, pages 520–531. Springer, 2011.

56 M. Pilipczuk and M. Sorge. A double exponential lower bound for the distinct vectors problem. *Discret. Math. Theor. Comput. Sci.*, 22(4), 2020.

57 A. Sebő and E. Tannier. On metric generators of graphs. *Mathematics of Operations Research*, 29(2):383–393, 2004.

58 P. J. Slater. Leaves of trees. In *Proc. of the 6th Southeastern Conference on Combinatorics, Graph Theory, and Computing*, pages 549–559. Congressus Numerantium, No. XIV. Utilitas Mathematica, 1975.

59 P. Tale. Double exponential lower bound for telephone broadcast, 2024. `arXiv:2403.03501`.