

On the hardness of recognizing graphs of small mim-width and its variants

Max Dupré la Tour* Manuel Lafond† Ndiamé Ndiaye‡

Abstract

The mim-width of a graph is a powerful structural parameter that, when bounded by a constant, allows several hard problems to be polynomial-time solvable - with a recent meta-theorem encompassing a large class of problems [SODA2023]. Since its introduction, several variants such as sim-width and omim-width were developed, along with a linear version of these parameters. It was recently shown that mim-width and all these variants are paraNP-hard, a consequence of the NP-hardness of distinguishing between graphs of linear mim-width at most 1211 and graphs of sim-width at least 1216 [ICALP2025]. The complexity of recognizing graphs of small width, particularly those close to 1, remained open, despite their especially attractive algorithmic applications.

In this work, we show that the width recognition problems remain NP-hard even on small widths. Specifically, after introducing the novel parameter Omim-width sandwiched between omim-width and mim-width, we show that: (1) deciding whether a graph has sim-width = 1, omim-width = 1, or Omin-width = 1 is NP-hard, and the same is true for their linear variants; (2) the problems of deciding whether mim-width ≤ 2 or linear mim-width ≤ 2 are both NP-hard. Interestingly, our reductions are relatively simple and are from the Unrooted Quartet Consistency problem, which is of great interest in computational biology but is not commonly used (if ever) in the theory of algorithms.

1 Introduction

Many hard problems on graphs become easier when the underlying graph has a simple structure, and over the years several “width” parameters have been developed to capture such structure. One of these is the mim-width (maximum induced matching width), introduced by Vatschelle [Vat12]. A branch decomposition of a graph G is a ternary tree whose leaves represent the vertices of G , and each of its edges naturally defines a cut of the graph. For any such decomposition, each cut induces a bipartite graph, in which we consider the size

*McGill University: maxduprelatour@gmail.com

†Université de Sherbrooke: manuel.lafond@usherbrooke.ca

‡McGill University: ndiame.ndiaye@mail.mcgill.ca

of a maximum induced matching. The mim-width of G is then the minimum, over all branch decompositions of G , of the maximum size of such an induced matching across all cuts of the decomposition [Vat12; BV13].

It is a powerful parameter that can be constant even on graphs of high clique-width, and several algorithmic problems are polynomial-time solvable on graphs of constant mim-width [JKT20a; JKT20b; Jaf+19], under the assumption that a branch decomposition is given. This is notably witnessed by a recent logic-based meta-theorem that solves many problems in polynomial-time under these assumptions [BDJ23].

Since its inception, several variants of mim-width were developed: the *sim-width* parameter is the same, except that edges of G within both sides of the cut are kept [Kan+17; MY23]; the *one-sided mim-width* (omim-width) keeps only the edges within either side of the cut and takes the minimum [BKR23]; the *large one-sided mim-width* (Omim-width), which we introduce in this work, instead takes the maximum; and the *linear* variant of all these parameters require the branch decomposition to be a caterpillar (see precise definitions below).

Recognizing graphs of mim-width at most k was quickly established to be $W[1]$ -hard in parameter k [SV16], but the question of XP membership was posed several times since the creation of the parameter (in addition to the previous references, this was also raised in [BPT22; BKR23; OST24; BDJ23]). The same remained open for all the aforementioned variants of mim-width, except the the linear sim-width = 1 case shown NP-hard in [Zie+18]. All these questions were finally answered by Bergougnoux, Bonnet, and Duronin [BBD25], where the authors showed that it is NP-hard to distinguish between graphs of linear mim-width at most 1211 and graphs of sim-width at least 1216, implying the paraNP-hardness of mim-width but also of all the above variants.

It is still open whether there is a constant $c < 1211$ such that recognizing graphs of mim-width (and variants) at most c is in P. Determining whether such a c exists is quite relevant, because graphs of small width are those that matter for algorithmic applications. In this work, we reduce this knowledge gap significantly by proving the following results:

Theorem 1. *The problem of deciding whether a graph has sim-width, omim-width, or Omim-width equal to 1 is NP-complete. Moreover, assuming the ETH, it cannot be solved in time $2^{o(n)}$, where n is the number of vertices of the graph. The same holds for the linear variant of these three parameters.*

Theorem 2. *The problem of deciding whether a graph has mim-width or linear mim-width at most 2 is NP-complete. Moreover, assuming the ETH, it cannot be solved in time $2^{o(n)}$, where n is the number of vertices of the graph.*

Note that the ETH refers to the Exponential Time Hypothesis [IP01]. We also point out that the new parameter Omim-width is a lower bound on mim-width, and an upper bound on omim-width, which in turn is an upper bound on sim-width.

Our reductions are relatively simple and arguably less involved than in [BBD25], which require a series of reductions with complex gadgets. Another unique as-

pect of our reductions is that they are all from the Unrooted Quartet Consistency (UQC) problem, shown NP-hard in [Ste92]. The input to UQC is a set of trees on four leaves, and one must decide whether there is a single tree that contains all of them.

This is not a common problem to reduce from, as UQC is often more seen as an “end-user” problem, in the sense that it is mostly popular for its bioinformatics applications. Its hardness has mainly been used to justify heuristics and approximations, and we are not aware of UQC being used to show the hardness of other problems (except for some that obviously contain UQC as a special case [BS01; RBR14]). On the other hand, our approach shows that UQC may be well-suited for problems that require finding a branch decomposition, since UQC is a tree reconstruction problem. It is possible that UQC or similar variants may be useful for other open XP membership problems that involve graph-to-tree representations (for instance module-width [Rao08; BV13] which is within a factor two of clique-width, or other similar problems stated in [Høg25]).

Although our results do not imply hardness for mim-width and variants between 3 and 1210, they indicate that there is little hope of recognizing graphs of small width efficiently. The notable cases of mim-width 1 and linear mim-width 1 remain open. These two graph classes have exploitable structural properties, for instance graphs of mim-width 1 are perfect [Vat12] and graphs of linear mim-width 1 have no so-called asteroidal edge triples [Høg25], and their membership in P is still possible. Recognizing *leaf powers*, one of the most natural subclasses of graphs of mim-width 1, was recently shown to be NP-hard [TLN25]. On the other hand, recognizing *linear leaf powers*, a subclass of graphs of linear mim-width 1, is in P [Ben+22].

2 Preliminaries

2.1 Notations and Definitions

All graphs considered are simple and undirected. The vertex set and edge set of a graph G are denoted by $V(G)$ and $E(G)$, respectively. An edge between two vertices u and v is denoted by uv (or equivalently vu). For a subset $X \subseteq V(G)$, let $G[X]$ denote the subgraph of G induced by X , and let $E_G[X]$ denote the edge set of $G[X]$. For a set of edges $F \subseteq E(G)$, we write $G - F$ for the graph obtained from G by deleting all edges in F .

A *matching* of G is a set of edges no two of which share a common endpoint. An *induced matching* of G is a matching M such that every edge of G is incident with at most one edge of M . A *cut* of G is a bipartition (A, B) of $V(G)$. We denote by $G[A, B]$ the bipartite subgraph of G with edge set $\{uv \in E(G) \mid u \in A, v \in B\}$.

Unless stated otherwise, all trees in this work are unrooted. A tree is *ternary* if each vertex has either 1 or 3 neighbors¹. A *caterpillar* is a tree in which the

¹Note, ternary trees are sometimes called binary, but here we prefer “ternary” to emphasize

set of internal vertices induces a path, called its *spine*. Given a total order $\leq$ on a set P of $n \geq 4$ points, we define the *ternary caterpillar realizing* $\leq$, with leaf set P as follows. If $p_1 < p_2 < \dots < p_{n-1} < p_n$, then the interval $[2, n-1]$ forms the spine of the caterpillar. Attach each leaf p_i to the integer i , and additionally attach p_1 to 2 and p_n to $n-1$. Note that a given ternary caterpillar with leaf set P realizes several total orders, since we may choose which end of the spine is the minimum/maximum and, at each end, which of the two incident leaves plays the “endpoint” role.

Width parameters: Given a cut (A, B) : The *mim-value* of (A, B) , denoted $\text{mim}_G(A, B)$, is the maximum size of an induced matching in $G[A, B]$.

The *sim-value* of (A, B) , denoted $\text{sim}_G(A, B)$, is the maximum size of a matching in $G[A, B]$ that is also an induced matching of G (hence, edges within A and within B are also considered).

For a subset $X \subseteq V(G)$, the *upper-induced matching number* of X is the maximum size of a matching in $G[X, V(G) \setminus X]$ that is an induced matching of the graph $G - E_G[V(G) \setminus X]$.

The *omim-value* of a cut (A, B) , denoted $\text{omim}_G(A, B)$, is the minimum of the upper-induced matching numbers of A and B .

The *Omim-value* of a cut (A, B) , denoted $\text{Omim}_G(A, B)$, is the maximum of the upper-induced matching numbers of A and B .

Note that it directly follows from the definitions of the parameters that $\text{mim}_G(A, B) \geq \text{Omim}_G(A, B) \geq \text{omim}_G(A, B) \geq \text{sim}_G(A, B)$.




Figure 1: A graph illustrating the differences of the width parameters on a specific cut (A, B)

A *branch decomposition* of a graph G is a ternary tree T whose leaves correspond exactly to the vertices of G . Given such a tree T , each edge e of T partitions the vertices of G into three sets, each of which is a union of leaves of T . that internal vertices have 3 neighbors, not 2.

T defines a cut (A_e, B_e) , where A_e and B_e are the vertex sets corresponding to the two connected components of $T - \{e\}$. The *mim-width* (resp. *sim-width*, *omim-width*, *Omim-width*) of T is the maximum, over all edges e of T , of $\text{mim}_G(A_e, B_e)$ (resp. $\text{sim}_G(A_e, B_e)$, $\text{omim}_G(A_e, B_e)$, $\text{Omim}_G(A_e, B_e)$). Finally, the *mim-width* (resp. *sim-width*, *omim-width*, *Omim-width*) of G is the minimum, over all branch decompositions T of G , of the *mim-width* (resp. *sim-width*, *omim-width*, *Omim-width*) of T .

Each of these width parameters has a *linear* variant, which require T to be a *caterpillar*. We denote such a variant by prefixing the parameter name with “linear”, e.g., $\text{linear-mim-width}(G)$ is the minimum *mim-width* of all branch decompositions of G that are caterpillars.

Lemma 1. *For any graph G , we have*

$$\text{mim-width}(G) \geq \text{Omim-width}(G) \geq \text{omim-width}(G) \geq \text{sim-width}(G).$$

This chain of inequalities also holds by replacing each parameter by their linear variant. Moreover, $\text{linear-mim-width}(G) \geq \text{mim-width}(G)$, and similarly for each variant of the width parameter.

2.2 Unrooted Quartet Consistency

Given a set of points P , a *quartet* consists of two unordered pairs of elements of P that do not intersect. For clarity, such a quartet $q = \{\{p_i, p_j\}, \{p_k, p_\ell\}\}$ is instead denoted $q = [p_i p_j | p_k p_\ell]$. Slightly abusing notation, we may write $p \in q$ when p is one of the four elements involved in q .

In a ternary tree T with leaves labeled $\{p_1, \dots, p_n\}$, consider any four leaves p_i, p_j, p_k, p_ℓ . If there exists an edge e of T such that p_i and p_j lie on one side of $T - \{e\}$, and p_k and p_ℓ lie on the other, we denote this relation by

$$p_i p_j \mid_T p_k p_\ell.$$

In this case, we say that T *satisfies* the quartet $q = [p_i p_j | p_k p_\ell]$. If Q is a set of quartets, then T satisfies Q if it satisfies each $q \in Q$.

Because T is ternary, given four leaves p_i, p_j, p_k, p_ℓ , exactly one of the three configurations holds:

$$p_i p_j \mid_T p_k p_\ell, \quad p_i p_k \mid_T p_j p_\ell, \quad \text{or} \quad p_i p_\ell \mid_T p_j p_k.$$

The NP-hard problem used in our reductions is the *Unrooted Quartet Consistency* (UQC) problem, introduced by Steel [Ste92].

Input: A set Q of quartets and a set of points $P = \{p_1, p_2, \dots, p_n\}$.

Question: Does there exist a ternary tree T with leaves labeled by the points in P that satisfies every quartet in Q ?

The hardness of UQC is proved by a reduction from BETWEENNESS, a classical problem that was shown to be NP-complete by Opatrny [Opa79]. In that problem, the input is a set B of ordered triples whose elements are from a

set of points U . One must decide whether there is a total order $<$ on U such that, for each $(a, b, c) \in B$, one of $a < b < c$ or $c < b < a$ holds. Steel's reduction proves a slightly stronger result than NP-hardness, which we will use.

Theorem 3 ([Ste92]). *There is a polynomial-time reduction from BETWEENNESS to UQC. Moreover, if an instance of BETWEENNESS is a YES-instance, then there exists a caterpillar that satisfies the corresponding UQC instance.*

By analyzing the chain of reductions used to prove Theorem 3, we can also obtain a slightly stronger result under ETH.

Theorem 4. *Assuming the ETH, the UQC problem cannot be solved in time $2^{o(n+m)}$, where n is the number of points of an instance and m is the number of quartets.*

Proof. The result can be deduced by simply revisiting the chain of reductions resulting in the NP-hardness of UQC. The first reduction is from 3-SET SPLITTING to BETWEENNESS. In 3-SET SPLITTING, we receive sets S of size 3 over a universe U and must color U with two colors so that no set of S is monochromatic. This cannot be solved in time $2^{o(|U|+|S|)}$ under the ETH (a proof appears in [Ant+24, Proposition 5.1]).

Opatrný's reduction from 3-SET SPLITTING to BETWEENNESS [Opa79, Lemma 2] produces from (S, U) an instance (B, U_B) with $|B| = 3m$ and $|U_B| = 2n + 1$ (each set in S becomes three betweenness triples, each element of U has two corresponding elements in U_B , plus an extra point). It follows that BETWEENNESS cannot be solved in time $2^{o(|B|+|U_B|)}$.

Finally, Steel's reduction [Ste92] from BETWEENNESS to UQC transforms (B, U_B) to an instance (Q, P) with $|Q| = 6|B|$ (each betweenness triple becomes 6 quartets) and $|P| = |U_B| + 2 + 4|B|$ (the point set consists of the elements of U_B and two extra points, and each betweenness constraint adds four points). Since $|P| + |Q|$ is linear in $|B| + |U_B|$, UQC cannot be solved in time $2^{o(|P|+|Q|)} = 2^{o(n+m)}$. $\square$

3 Sim-width, omim-width and Omim-width

We start with a single reduction showing that it is NP-hard to decide whether a given graph has sim-width, omim-width, or Omim-width equal to 1. The same holds for their linear variants. Given a UQC instance (Q, P) , we define the following graph G :

- **Vertices:** For each element $p \in P$, we introduce a vertex p in G . For each quartet $q = [p_i p_j \mid p_k p_\ell] \in Q$, we introduce four additional vertices in G :

$$\{u_i^q, u_j^q, u_k^q, u_\ell^q\}.$$

Let U denote the set of all such u -vertices, so that $|U| = 4|Q|$.

- **Edges:** For each quartet $q = [p_i p_j \mid p_k p_\ell] \in Q$, add the edges

$$p_i u_i^q, u_i^q u_j^q, u_j^q p_j, \text{ and } p_k u_k^q, u_k^q u_\ell^q, u_\ell^q p_\ell.$$

In addition, connect every pair of vertices $u, u' \in U$ whenever they correspond to different quartets. That is, for $u_i^q, u_j^r \in U$, we add the edge $u_i^q u_j^r$ whenever $q \neq r$.

Observe that P forms an independent set, and that each $u_i^q \in U$ has only two non-neighbors in U .




Figure 2: Subgraph of G corresponding to a quartet $q = [p_i p_j \mid p_k p_\ell]$.

Note that $|V(G)| = |P| + 4|Q|$. The fact that $|V(G)|$ is linear in $|P|$ and $|Q|$ will be used to obtain the ETH-based hardness part of our result. Note that we are not able to obtain the stronger lower bound under ETH of $2^{o(|E(G)|)}$, because the number of edges in this reduction is quadratic.

The combination of Lemma 1, Theorem 4, and the following proposition will imply our first hardness result, Theorem 1.

Proposition 1. *If there is a caterpillar satisfying the UQC instance (Q, P) , then $\text{linear-omim-width}(G) = 1$. Otherwise, if $(Q, P) \notin \text{UQC}$, then any branch decomposition T of G satisfies $\text{sim-width}_G(T) = 2$, and in particular $\text{sim-width}(G) = 2$.*

We split the proof in two lemmas.

Lemma 2. *If $(Q, P) \notin \text{UQC}$, then any branch decomposition T of G satisfies $\text{sim-width}_G(T) = 2$, and in particular $\text{sim-width}(G) = 2$.*

Proof. We start by proving that $\text{sim-width}(G) \leq 2$.

This follows directly from the fact that the maximum induced matching in the full graph G has size at most 2. Suppose, for contradiction, that there exists an induced matching of size at least 3. Since the set P is an independent set in G , at least one endpoint of every edge in the matching must belong to U . By construction, each vertex in U is adjacent to all other vertices of U except two, which themselves are adjacent. Hence, among any three such vertices of U , there must exist at least one edge between them, contradicting the assumption that they form an induced matching.

We now prove that if $(Q, P) \notin \text{UQC}$, then $\text{sim-width}(G) \geq 2$.

Let $(Q, P) \notin \text{UQC}$, and let T be a branch decomposition of G . There must exist a quartet $q = [p_i p_j \mid p_k p_\ell] \in Q$ not satisfied by T , otherwise (Q, P) would be in UQC . Assume without loss of generality that $p_i p_k \mid_T p_j p_\ell$, and let e be an edge of T inducing the cut (A_e, B_e) such that $p_i, p_k \in A_e$ and $p_j, p_\ell \in B_e$. At least one of the edges along the path $p_i - u_i^q - u_j^q - p_j$ must cross this cut, and similarly, one of the edges along the path $p_k - u_k^q - u_\ell^q - p_\ell$ must also cross the cut. Because there is no edge with one endpoint in the first path and the other in the second path, these two edges form an induced matching in G , and therefore $\text{sim-width}_G(T) \geq 2$ for every branch decomposition T . Consequently, $\text{sim-width}(G) \geq 2$. $\square$

Lemma 3. *If there is a caterpillar satisfying the UQC instance (Q, P) , then $\text{linear-Omim-width}(G) = 1$.*

Proof. Let T be such a caterpillar, and let $\leq$ be a total order on P that T realizes. For each $p_i \in P$, define

$$C_i = \{p_i\} \cup \{u_i^q \mid q \in Q, p_i \in q\},$$

that is, C_i consists of all vertices in $P \cup U$ whose index is i . We extend $\leq$ to a total order $\leq'$ on $P \cup U$ as follows.

For all $p_i, p_j \in P$ with $p_i < p_j$, and for all $x \in C_i, y \in C_j$, we set $x <' y$. Within each C_i , the relative position of p_i and any u_i^q depends on the quartet q :

If $q = [p_i p_j \mid p_k p_\ell]$, then, since T satisfies q , we have either $p_i, p_j \leq p_k, p_\ell$ or $p_k, p_\ell \leq p_i, p_j$. In the former case $(p_i, p_j \leq p_k, p_\ell)$ we set $p_i <' u_i^q$, in the latter case $(p_k, p_\ell \leq p_i, p_j)$ we set $u_i^q <' p_i$. Note that this implies that we have either

$$p_i <' u_i^q <' u_k^q, u_\ell^q, \text{ or } u_k^q, u_\ell^q <' u_i^q <' p_i.$$

We choose an arbitrary order within C_i among the u -vertices to the left of p_i and among those to the right of p_i . We then consider the ternary caterpillar T' realizing $\leq'$. We now prove that $\text{Omim-width}_G(T') = 1$, that is, $\text{Omim}_G(A_e, B_e) = 1$ for every edge e of T' .

Because T' is a caterpillar, it suffices to consider the cuts induced by the edges of its spine. Let e be an edge of the spine, and let (A_e, B_e) be the cut induced by e . Since T' realizes the order $\leq'$, one side of the cut corresponds to an initial segment of $\leq'$. Thus, we have either for all $a \in A_e$ and $b \in B_e$, $a <' b$, or for all $a \in A_e$ and $b \in B_e$, $b <' a$.

To prove that $\text{Omim}_G(A_e, B_e) = 1$, we need to prove that the upper induced matching numbers of A_e and B_e are both equal to one. This means that there is no matching of size 2 in $G[A_e, B_e]$ that is an induced matching in the graph $G - E_G[A_e]$ or in the graph $G - E_G[B_e]$. Let $a_1 b_1, a_2 b_2 \in E(G)$ be two edges with disjoint endpoints such that $a_1, a_2 \in A_e$ and $b_1, b_2 \in B_e$. To prove that it is not an induced matching in $G - E_G[A_e]$ or $G - E_G[B_e]$, it suffices to show that one of the two following conditions holds:

1. One of the diagonals $a_1 b_2$ or $a_2 b_1$ is an edge, or

2. Both a_1a_2 and b_1b_2 are edges.

Each vertex $p_i \in P$ is connected only to vertices of the form u_i^q . Consider two edges $p_iu_i^q$ and $p_ju_j^{q'}$ with $i \neq j$. By the construction of the order $\leq'$, we have either $p_i, u_i^q <' p_j, u_j^{q'}$ or $p_j, u_j^{q'} <' p_i, u_i^q$. Therefore, it is impossible for the cut (A_e, B_e) to contain both edges $p_iu_i^q$ and $p_ju_j^{q'}$, as the cut can separate at most one p_i vertex from one of its neighbors in G . Consequently, at most one vertex among $\{a_1, a_2, b_1, b_2\}$ can belong to P .

Suppose that exactly one vertex among $\{a_1, a_2, b_1, b_2\}$ belongs to P . Without loss of generality, assume that $a_1 = p_i$. Then b_1 must be of the form u_i^q . Suppose by contradiction that condition 1 is false and that a_2b_1 is not an edge. Noting that $b_1 \in U$, the only vertices in U that are not adjacent to u_i^q are those corresponding to the other side of the same quartet. That is, if $q = [p_ip_j \mid p_kp_\ell]$, then these vertices are u_k^q and u_ℓ^q . By the definition of the order $\leq'$, we have either $p_i <' u_i^q <' u_k^q, u_\ell^q$, or $u_k^q, u_\ell^q <' u_i^q <' p_i$. Since the edge e lies on the path between $a_1 = p_i$ and $b_1 = u_i^q$ in T' , the vertices u_k^q and u_ℓ^q cannot belong to A_e . This yields a contradiction.

Now suppose that all vertices $\{a_1, a_2, b_1, b_2\}$ belong to U . Let q be the quartet associated with a_1 and q' the quartet associated with b_1 . As before, assume for contradiction that condition 1 is false, and that neither a_1b_2 nor a_2b_1 is an edge. The only vertices in U not adjacent to a_1 correspond to the other side of the same quartet, so q must also be the quartet of b_2 . Similarly, q' must be the quartet of a_2 .

If $q = q'$, then we have $\{a_1, a_2, b_1, b_2\} = \{u_i^q, u_j^q, u_k^q, u_\ell^q\}$ for some $q = [p_ip_j \mid p_kp_\ell] \in Q$. The only edges among these vertices are $u_i^qu_j^q$ and $u_k^qu_\ell^q$, which are both separated by e . Consequently, the quartet $[u_i^qu_j^q \mid u_k^qu_\ell^q]$ is not satisfied by T' , which, by the definition of $\leq'$, implies that q is not satisfied by T , and we have a contradiction.

On the other hand, if $q \neq q'$, then both a_1a_2 and b_1b_2 are edges, and hence condition 2 holds. $\square$

Proposition 1 follows from Lemma 2 and Lemma 3. We are now ready to prove Theorem 1.

Theorem 1. *The problem of deciding whether a graph has sim-width, omim-width, or Omim-width equal to 1 is NP-complete. Moreover, assuming the ETH, it cannot be solved in time $2^{o(n)}$, where n is the number of vertices of the graph. The same holds for the linear variant of these three parameters.*

Proof. Let σ be one of the parameters sim-width, omim-width, Omim-width, linear-sim-width, linear-omim-width, or linear-Omim-width. We want to show that deciding whether $\sigma(G) = 1$ is NP-complete. Membership in NP is easy since a branch decomposition is an easily verifiable certificate.

For NP-hardness, for technical accuracy we need to use Theorem 3 and reduce BETWEENNESS to deciding $\sigma(G) = 1$. Take an instance I of BETWEENNESS and let (Q, P) be the UQC instance obtained from Theorem 3, and then obtain

G as constructed above from that UQC instance. If I is a YES-instance of BETWEENNESS, then $(Q, P) \in \text{UQC}$ and by Theorem 3 there is a ternary caterpillar that satisfies Q . Then, Proposition 1 implies that $\text{linear-omim-width}(G) = 1$. This implies that $\sigma(G) = 1$, since $\text{linear-omim-width}(G)$ is an upper bound on all the possible σ parameters listed above (by Lemma 1).

Conversely, if I is a NO-instance of BETWEENNESS, then no tree satisfies the UQC instance (Q, P) . By Proposition 1, $\text{sim-width}(G) = 2$, which implies $\sigma(G) \geq 2$ since $\text{sim-width}(G)$ is a lower bound on all the possible σ parameters listed above (by Lemma 1). Therefore, I is a YES-instance of BETWEENNESS if and only if $\sigma(G) = 1$.

Moreover $|V(G)| = |P| + 4|Q|$, therefore an algorithm recognising graphs with σ parameter equal to one in time $2^{o(|V(G)|)}$ could be used to solve UQC in time $2^{o(|P|+|Q|)}$, contradicting Theorem 4 assuming ETH. $\square$

Remarks on the graph class of hard instances. Let us observe that, denoting by C_k the chordless cycle on k vertices, our construction produces a graph G in which the U vertices induce the complement of a set of disjoint C_4 's, and that all P vertices have only neighbors in U and are simplicial i.e., their neighbors form a clique (since two distinct neighbors of a p_i vertex have the form $u_i^q, u_i^{q'}$ with $q \neq q'$). This implies that G is C_k -free for every $k \geq 5$: the p_i 's cannot be in a C_k because they are simplicial, and can show that the U vertices cannot have a C_k since they are the complement of disjoint C_4 's. We note that for chordal graphs, i.e., graphs that are C_k -free for every $k \geq 4$, whether $\text{linear-sim-width}(G) = 1$ can be decided in polynomial time [Zie+18]. The set of forbidden induced cycles from our construction is therefore, in some sense, tight for the case linear-sim-width equal to 1.

Additionally, the complement $\overline{G}$ is also C_k -free for $k \geq 5$. To see this, suppose that $\overline{G}$ has an induced C_k and notice that it must contain some p_i vertex (as the U vertices induce disjoint C_4 's). Since P is a clique in $\overline{G}$, p_i must have some neighbor u_j^q on C_k , with $i \neq j$. The other neighbor of u_j^q on C_k can only be u_i^q since a p_j vertex or a u_ℓ^q vertex with $i \neq \ell$ would form a triangle with p_i . But then, the other neighbor of u_i^q is either some u_ℓ^q or some p_j vertex, which are both neighbors of p_i , and so no C_k is possible for $k \geq 5$. Thus the hard instances do not contain induced cycles of length 5 or more nor their complements. In particular, they are perfect graphs and the NP-hardness also holds on this class of graphs.

Using a similar reduction for mim-width 1 (or linear mim-width 1)? It is interesting to ponder the possibility of using a similar reduction for $\text{mim-width} = 1$. It appears difficult to do so and here we provide the intuitive reasons for this. The high-level idea in our reduction is to construct a graph G that includes vertices representing the elements of the UQC instance and, for each quartet q , a gadget H_q . A reduction for $\text{mim-width} = 1$ would have the following intended properties.

1. In any branch decomposition of G , if the four vertices of a quartet q do not

satisfy q , then some edge of the decomposition induces a cut within H_q containing an induced matching of size 2. Thus, H_q is meant to “force” a particular shape for the subtree containing these four vertices in any branch decomposition of mim-width 1.

2. If the UQC instance is satisfiable, then there exists a branch decomposition of G whose mim-width is exactly 1.

The difficulty arises when we try to satisfy these conditions simultaneously. Consider two different quartets q and q' with corresponding gadgets H_q and $H_{q'}$. In the branch decomposition guaranteed by property 2, the subtrees corresponding to these two gadgets may overlap (it seems hard to avoid this). As a consequence, an edge of H_q used in the induced matching when q is unsatisfied (property 1) may lie in the same cut as an edge of $H_{q'}$ used when q' is unsatisfied.

To ensure mim-width 1, we must prevent these two edges to form an induced matching. This forces us to add *blocking edges* between certain endpoints of edges coming from H_q and $H_{q'}$ in order to destroy the unwanted induced matchings. In our reduction, the blocking edges correspond to the edges between the u -vertices corresponding to different quartets.

However, these blocking edges create a new problem: they may themselves form an induced matching of size 2 across some cut, contradicting property 2 and preventing the existence of a mim-width 1 decomposition even when the UQC instance is satisfiable. In our reduction, this correspond to the last case in the proof of Lemma 3: the matching of the type $u_x^q u_y^{q'}, u_w^{q'} u_z^q$ for different q, q' , which is only destroyed with edges not in the cut.

In short, enforcing quartet constraints independently via the gadgets H_q inevitably causes interactions (either among the gadgets or among the blocking edges) that create new induced matchings. We found no way around this, which is why our approach does not seem extendable to mim-width 1 (or linear mim-width 1).

4 Mim-width and linear mim-width

In this section, we describe two closely related reductions. The first is used to prove the hardness of recognizing graphs with mim-width 2, and the second establishes the hardness of recognizing graphs with linear-mim-width 2.

We modify the previous reduction by adding four new vertices per quartet. Additionally, one extra vertex is introduced in the non-linear version.

Given a UQC instance (Q, P) , we define the following graph G :

- **Vertices:** For each element $p \in P$, we introduce a vertex p . For each quartet $q = [p_i p_j \mid p_k p_\ell] \in Q$, we introduce eight additional vertices:

$$\{u_i^q, u_j^q, u_k^q, u_\ell^q, \gamma_i^q, \gamma_j^q, \gamma_k^q, \gamma_\ell^q\}$$

Let U (resp. Γ) denote the set of all such u -vertices (resp. γ -vertices).

- **Edges:** For each quartet $q = [p_i p_j \mid p_k p_\ell] \in Q$, add the edges

$$\{p_i, u_i^q\}, \{u_i^q, u_j^q\}, \{u_j^q, p_j\}, \{p_k, u_k^q\}, \{u_k^q, u_\ell^q\}, \{u_\ell^q, p_\ell\},$$

together with

$$\{p_i, \gamma_i^q\}, \{p_j, \gamma_j^q\}, \{p_k, \gamma_k^q\}, \{p_\ell, \gamma_\ell^q\}.$$

Moreover, add all possible edges so that Γ forms a clique, add an edge between every pair of vertices $u, u' \in U$ whenever they correspond to different quartets (i.e., whenever their superscript differs), and between every pair of vertices $u \in U$ and $\gamma \in \Gamma$, whenever they correspond to different quartets.

The graph H is obtained from G by adding an extra vertex ω , sharing an edge with every vertex in U .




Figure 3: Gadget corresponding to a single quartet $q = [p_i p_j \mid p_k p_\ell]$, with the extra vertex ω of H represented.

Note, the constructed graphs G and H have mim-width at least 2, since Vatshelle argued that graphs of mim-width 1 [Vat12, Corollary 3.7.4] have no induced cycle of length 5 or more, whereas these graphs have such cycles of the form $p_i - u_i^q - u_j^q - p_j - \gamma_j^q - \gamma_i^q - p_i$.

Proposition 2. *If there is a caterpillar satisfying the UQC instance (Q, P) , then $\text{linear-mim-width}(G) = 2$, and $\text{mim-width}(H) = 2$. Otherwise, if $(Q, P) \notin \text{UQC}$ then $\text{linear-mim-width}(G) \geq 3$, and $\text{mim-width}(H) \geq 3$.*

The proof have the same structure as in the previous section, but require a more involved case analysis. We split it in several lemmas.

Lemma 4. *If $(Q, P) \notin \text{UQC}$ then $\text{linear-mim-width}(G) \geq 3$.*

Lemma 5. *If $(Q, P) \notin \text{UQC}$ then $\text{mim-width}(H) \geq 3$.*

The first step in the proofs of these two lemmas is similar. We begin by showing a result on any branch decomposition of G . Any branch decomposition of H can be transformed into one of G by deleting the ω -leaf. The result will therefore also apply to branch decompositions of H .

Let $(Q, P) \notin \text{UQC}$, and let T be any branch decomposition of G . Assume by contradiction that $\text{mim}_G(T) \leq 2$. There must exist a quartet $q = [p_i p_j \mid p_k p_\ell] \in Q$ that is not satisfied by T ; otherwise, (Q, P) would belong to UQC .

Without loss of generality, assume that $p_i p_k \mid_T p_j p_\ell$ is true. Let π be the set of all edges e of T inducing a cut (A_e, B_e) such that $p_i, p_k \in A_e$ and $p_j, p_\ell \in B_e$. Note that π forms a path. Let A (resp. B) denote $\bigcap_{e \in \pi} A_e$ (resp. $\bigcap_{e \in \pi} B_e$); that is, A (resp. B) consists of all vertices on the side of p_i, p_k (resp. p_j, p_ℓ) of π .

Claim 1. *Either $\gamma_i^q, \gamma_j^q, \gamma_k^q, \gamma_\ell^q$ are all contained in A , or they are all contained in B .*

Proof. Suppose, for the sake of contradiction, that there exists an edge e of π such that one of the γ -vertices lies in A_e and another lies in B_e . We will show that we can always find an induced matching of size 3 in $G[A_e, B_e]$. We proceed by a case analysis.

First, consider the case where both γ_i^q and γ_k^q lie in A_e . Choose $x \in \{j, \ell\}$ such that γ_x^q is a γ -vertex in B_e . Along the path $p_i - u_i^q - u_j^q - p_j$ of G , at least one edge crosses the cut (A_e, B_e) ; likewise, at least one edge of the path $p_k - u_k^q - u_\ell^q - p_\ell$ also crosses the cut. These two crossing edges, together with the edge $\gamma_i^q \gamma_x^q$, form an induced matching in $G[A_e, B_e]$. Indeed, the only other edges that could appear among these six vertices in G are $p_i \gamma_i^q$, which lies entirely within A_e , and $p_x \gamma_x^q$, which lies entirely within B_e .

Therefore, we may assume that at least one of γ_i^q or γ_k^q is in B_e . Without loss of generality, we may assume that $\gamma_i^q \in B_e$. By applying the symmetric argument on the pair $\gamma_j^q, \gamma_\ell^q$, we may assume that one of the vertices γ_j^q or γ_ℓ^q lies in A_e .

Suppose now that $u_j^q \in A_e$. Then one of the edges on the path $p_k - u_k^q - u_\ell^q - p_\ell$, together with the edges $p_i \gamma_i^q$ and $u_j^q p_j$, forms an induced matching in $G[A_e, B_e]$. Hence we may assume $u_\ell^q \in B_e$.

Next, suppose that $u_i^q \in A_e$. Then one of the edges on the path $p_k - u_k^q - u_\ell^q - p_\ell$, together with the edges $p_i \gamma_i^q$ and $u_\ell^q p_\ell$, forms an induced matching in $G[A_e, B_e]$. Hence we may assume $u_i^q \in B_e$.

If we had $\gamma_j^q \in A_e$, then one of the edges on the path $p_k - u_k^q - u_\ell^q - p_\ell$, together with the edges $p_i \gamma_i^q$ and $\gamma_j^q p_j$, forms an induced matching in $G[A_e, B_e]$. Therefore $\gamma_j^q \in B_e$.

It follows that $\gamma_\ell^q \in A_e$, and by symmetry again, we have $\gamma_k^q, u_k^q, u_\ell^q \in A_e$ as well. In this situation, the three edges $u_\ell^q p_\ell$, $p_i u_i^q$, and $\gamma_k^q \gamma_j^q$ form an induced matching in $G[A_e, B_e]$.


This establishes that either $\gamma_i^q, \gamma_j^q, \gamma_k^q, \gamma_\ell^q$ are all contained in A , or they are all contained in B . $\square$

We will assume without loss of generality that all the gammas are in B . We are now ready to finish the proof of Lemma 4.

Proof of Lemma 4. We now assume for contradiction that T is a caterpillar satisfying $\text{linear-mim-width}(T) \leq 2$, implying in particular that $\text{mim-width}(T) \leq 2$. Note that Claim 1 is applicable to T since it holds for any branch decomposition. Now, one of p_j or p_ℓ is adjacent to an endpoint of π in T (recall that π is the set of edges separating A and B). Without loss of generality, suppose it is p_j , and let z be the neighbor of p_j on the spine of the caterpillar. Let e be the edge incident to z that is not in π , other than zp_j (so, e is the next edge on the spine beyond π in the direction past p_j , going towards p_ℓ). This induces a cut (A_e, B_e) with p_ℓ and all the γ -vertices in B_e , and $p_i, p_j, p_k \in A_e$. Consequently, the edges $p_i\gamma_i^q, p_j\gamma_j^q$, and $p_k\gamma_k^q$ form an induced matching in $G[A_e, B_e]$. This completes the proof. $\square$

The proof of Lemma 5 is slightly more complicated, because T might not be a caterpillar. Therefore, we need a case analysis depending on which vertices of B lie in the branch containing p_j and which lie in the branch containing p_ℓ . The additional vertex ω is required to force an induced matching of size 3 in all configurations.

Proof of Lemma 5. We assume for contradiction that T is a branch decomposition of H satisfying $\text{mim-width}(T) \leq 2$. Let v be the endpoint of π that lies on the path between p_j and p_ℓ in T (thus, one incident “side” of v contains p_j , one contains p_ℓ , and the remaining side contains both p_i and p_k , see figure). Let e_{mid} be the edge from v toward p_i and p_k , e_{top} the edge toward p_j , and e_{bot} the edge toward p_ℓ . We denote by $(A_{\text{mid}}, B_{\text{mid}})$, $(A_{\text{bot}}, B_{\text{bot}})$, and $(A_{\text{top}}, B_{\text{top}})$ the cuts induced by $e_{\text{mid}}, e_{\text{bot}}, e_{\text{top}}$, respectively. We have $B_{\text{mid}} = B$, and we partition B as $B = B_{\text{bot}} \cup B_{\text{top}}$, where B_{bot} (resp. B_{top}) consists of the vertices of B in the branch toward p_ℓ (resp. toward p_j).



The remainder of the proof accumulates a series of facts under our contradiction assumptions, ultimately showing that every possible placement of the u -vertices, the γ -vertices, and ω leads to a 3-induced matching across one of the edges $e_{\text{mid}}, e_{\text{bot}}$, or e_{top} .

Fact 1. All the γ -vertices $\gamma_i^q, \gamma_j^q, \gamma_k^q, \gamma_\ell^q$ of the quartet q lie in B .

Proof. Claim 1 still holds, because T can be turned into a branch decomposition of G by deleting the ω -leaf, which does not increase the mim-width of the decomposition. $\square$

Fact 2. *All u -vertices $u_i^q, u_j^q, u_k^q, u_\ell^q$ of the quartet q lie in B .*

Proof. Suppose not, and assume without loss of generality that $u_i^q \notin B$ or $u_j^q \notin B$ (the cases $u_k^q \notin B$ or $u_\ell^q \notin B$ are symmetric). Then one of the edges $u_i^q u_j^q$ or $u_j^q p_j$ crosses the cut (A_{mid}, B_{mid}) (with $u_i^q \in A_{mid}$ in the first case). One of the edges of the path $p_k - u_k^q - u_\ell^q - p_\ell$ must also cross the cut. The vertex γ_i^q is in $B = B_{mid}$ by Fact 1. These two edges, together with the edge $p_i \gamma_i^q$, form an induced matching of size 3. $\square$

Fact 3. *The vertex ω is in B .*

Proof. Assume for contradiction that $\omega \notin B$. Using Fact 1 and Fact 2, all the u and γ vertices are in $B = B_{mid}$. In the cut (A_{mid}, B_{mid}) , the edges ωu_j^q , $p_i \gamma_i^q$, and $p_k \gamma_k^q$ form an induced matching of size 3. $\square$

Fact 4. *Each side B_{top} and B_{bot} contains at least one u -vertex.*

Proof. Note that by Fact 2, each u -vertex belongs to either B_{top} or B_{bot} . If all the four u -vertices lie on the same side, say B_{top} , then across the cut (A_{top}, B_{top}) the edges $p_i u_i^q$, $p_k u_k^q$, and $p_\ell u_\ell^q$ form an induced matching of size 3. Likewise, not all four u -vertices can lie in B_{bot} . $\square$

Fact 5. *At least one of $\gamma_k^q, \gamma_\ell^q$ lies in B_{bot} . Symmetrically, at least one of γ_i^q, γ_j^q lies in B_{top} .*

Proof. Using Fact 1, each γ -vertex is either in B_{bot} or B_{top} . Assume γ_k^q and γ_ℓ^q are both in B_{top} . Then one of the edges of the path $p_i - u_i^q - u_j^q - p_j$ must cross the cut (A_{top}, B_{top}) . Together with the edges $p_k \gamma_k^q$ and $p_\ell \gamma_\ell^q$, this yields an induced matching of size 3, a contradiction. $\square$

Using Fact 3, the vertex ω is either in B_{bot} or in B_{top} . In the rest of the proof, we will assume without loss of generality that $\omega \in B_{bot}$ (to see why this is without loss of generality, notice that if this is not the case, then by symmetry we may swap the subscripts of p_i and p_k , and the subscripts of p_j and p_ℓ without affecting the previous arguments).

Fact 6. *The vertices u_i^q and u_j^q are both in B_{top} .*

Proof. Consider u_i^q first. Using Fact 2, u_i^q is either in B_{bot} or in B_{top} . Assume by contradiction that it is in B_{bot} . Across the cut (A_{bot}, B_{bot}) , the following 3 edges form an induced matching:

- The edge $p_i u_i^q$.
- One of the edge $u_x^q \omega$ with $x \in \{j, k, \ell\}$. At least one u vertices is in B_{top} by Fact 4.
- An edge joining a γ -vertex in B_{top} to either γ_k^q or γ_ℓ^q in B_{bot} . Such an edge exists using Fact 5.

We may thus assume that $u_i^q \in B_{top}$.

Next assume for contradiction that $u_j^q \in B_{bot}$. Then across the cut (A_{bot}, B_{bot}) we may form the induced matching containing the edges $p_j u_j^q$, $u_i^q \omega$, and as before an edge joining a γ -vertex in B_{top} to either γ_k^q or γ_ℓ^q in B_{bot} . $\square$

Fact 7. *At least one of γ_j^q, γ_k^q lies in B_{top} .*

Proof. Using Fact 1, each γ -vertex is either in B_{bot} or B_{top} . Assume for contradiction that γ_j^q and γ_k^q are both in B_{bot} . Across the cut (A_{bot}, B_{bot}) , the following 3 edges form an induced matching: $p_j \gamma_j^q$, $p_k \gamma_k^q$, and $u_i^q \omega$ (we know that $u_i^q \in B_{top}$ using Fact 6). $\square$

Fact 8. *The vertex u_k^q is in B_{bot} .*

Proof. Using Fact 2, u_k^q is either in B_{bot} or in B_{top} . Assume by contradiction that it is in B_{top} . Across the cut (A_{top}, B_{top}) , the following 3 edges form an induced matching:

- $p_i u_i^q$. We have $u_i^q \in B_{top}$ by Fact 6.
- $u_\ell^q u_k^q$. We are assuming that $u_k^q \in B_{top}$. Moreover, we know that $u_i^q, u_j^q \in B_{top}$ (Fact 6), and that the u -vertices are not all in B_{top} (Fact 4); therefore u_ℓ^q is in B_{bot} .
- An edge joining a γ -vertex in B_{bot} to either γ_j^q or γ_k^q in B_{top} . Such an edge exists using Fact 7.

$\square$

Fact 9. *The vertex γ_k^q is the only gamma vertex in B_{bot} .*

Proof. Assume by contradiction that there is a γ -vertex other than γ_k^q in B_{bot} . Across the cut (A_{bot}, B_{bot}) , the following 3 edges form an induced matching:

- $p_k u_k^q$. We know that u_k^q is in B_{bot} by Fact 8.
- $u_i^q \omega$. We know that u_i^q is in B_{top} by Fact 6.
- An edge joining either γ_i^q or γ_j^q in B_{top} (Fact 5) to a γ -vertex in B_{bot} that is not γ_k^q .

$\square$

We can finally conclude the proof. By combining Fact 1 and Fact 9, we know that γ_i^q and γ_ℓ^q are both in B_{top} . Across the cut (A_{top}, B_{top}) , the following 3 edges form an induced matching: $p_i \gamma_i^q$, $p_\ell \gamma_\ell^q$, and ωu_j^q (since u_j^q is in B_{top} by Fact 6). This completes the proof: in all cases we obtain an induced matching of size 3.

$\square$

Lemma 6. *If there is a caterpillar satisfying the UQC instance (Q, P) , then $\text{linear-mim-width}(G) = 2$.*

Proof. As we already observed, $\text{mim-width}(G) \geq 2$ because it contains induced cycles of length 6, implying $\text{linear-mim-width}(G) \geq 2$. We focus on the upper bound.

Let T be a caterpillar satisfying (Q, P) , and let $\leq$ be a total order on P that T realizes. For each $p_i \in P$, define

$$C_i = \{p_i\} \cup \{u_i^q \mid q \in Q, p_i \in q\} \cup \{\gamma_i^q \mid q \in Q, p_i \in q\},$$

that is, C_i consists of all vertices in $P \cup U \cup \Gamma$ whose index is i . We extend $\leq$ to a total order $\leq'$ on $P \cup U \cup \Gamma$ as follows: for all $p_i, p_j \in P$ with $p_i < p_j$, and for all $x \in C_i, y \in C_j$, we set $x <' y$; within each C_i , the order is arbitrary. Let T' be the ternary caterpillar realizing $\leq'$. We prove that $\text{mim-width}_G(T') \leq 2$, i.e., $\text{mim}_G(A_e, B_e) \leq 2$ for every edge e of T' .

Because T' is a caterpillar, it suffices to consider the cuts induced by the edges of its spine. Let e be an edge of the spine, and let (A_e, B_e) be the cut induced by e . Since T' realizes $\leq'$, one side of the cut is an initial segment of $\leq'$. Thus, either for all $a \in A_e$ and $b \in B_e$ we have $a <' b$, or for all $a \in A_e$ and $b \in B_e$ we have $b <' a$.

Let $a_1b_1, a_2b_2, a_3b_3 \in E(G)$ be three edges with pairwise disjoint endpoints such that $a_1, a_2, a_3 \in A_e$ and $b_1, b_2, b_3 \in B_e$. Suppose, for contradiction, that these three edges form an induced matching in $G[A_e, B_e]$.

First, note that at most one of the six vertices lies in P . Indeed, a vertex $p_i \in P$ is adjacent only to vertices of C_i . If p_i were among the six vertices, then e would separate p_i from another element of C_i ; by the construction of $\leq'$, this can happen for at most one index i .

We say that a vertex v of G is of *type* q for a quartet q if v is one of the eight vertices in $U \cup \Gamma$ corresponding to q (i.e., the type is the superscript q). Suppose a_1b_1 is an edge with a_1 of type q and b_1 of type $q' \neq q$. In our construction, the only vertices not adjacent to a_1 are those in P or of type q ; hence b_2 and b_3 must each lie in P or be of type q . Symmetrically, the only vertices not adjacent to b_1 are those in P or of type q' , so a_2 and a_3 must each lie in P or be of type q' . If, without loss of generality, a_3 is the only possible vertex possibly in P , then a_2 is of type q' and b_3 is of type q , which forces $a_2b_3 \in E(G)$, contradicting that the three edges form an induced matching.

Therefore, at most one of the six vertices in the induced matching is in P and the remaining vertices in $U \cup \Gamma$ all have the same type $q = [p_i p_j \mid p_k p_\ell]$. The three matching edges must then be of the following kinds: $\gamma-\gamma$, $u-u$, $p-u$, or $p-\gamma$.

We cannot have two $\gamma-\gamma$ edges, since Γ is a clique and there would be an extra edge between the γ -endpoints, violating the induced property.

Assume there are two $u-u$ edges. The only possibilities are $u_i^q u_j^q$ and $u_k^q u_\ell^q$. Without loss of generality, let $u_i^q, u_k^q \in A_e$ and $u_j^q, u_\ell^q \in B_e$. By the definition of $\leq'$, this implies $u_i^q, u_k^q <' u_j^q, u_\ell^q$, hence $p_i, p_k < p_j, p_\ell$ in $\leq$, which contradicts that T satisfies the quartet q .

Since at most one of the six vertices lies in P , it follows that *exactly one* matching edge is of type $u-u$ (say a_1b_1), *exactly one* is of type $\gamma-\gamma$ (say a_2b_2), and the remaining one (say a_3b_3) is either $p-u$ or $p-\gamma$.

Without loss of generality, let $a_3 = p_i \in P$. If $b_3 = \gamma_i^q$, then b_3 is adjacent to a_2 (both are γ -vertices), so the matching is not induced. Otherwise, if $b_3 = u_i^q$, then $a_1 b_1$ must be $u_k^q u_\ell^q$. Since e separates p_i from u_i^q , the cut lies inside C_i ; together with $u_k^q u_\ell^q$ crossing the cut, this forces p_i to lie between p_k and p_ℓ in the order $\leq$, contradicting that T satisfies the quartet q . This contradiction shows that no induced matching of size 3 can appear in $G[A_e, B_e]$, and hence $\text{mim}_G(A_e, B_e) \leq 2$ for every spine edge e of T' . Therefore, $\text{mim-width}_G(T') \leq 2$. $\square$

Lemma 7. *If there is a caterpillar satisfying the UQC instance (Q, P) , then $\text{mim-width}(H) = 2$.*




Figure 4: Expansion of p_i in T into the subtree corresponding to C_i in T' .

Proof. Again, $\text{mim-width}(H) \geq 2$ follows from the existence of induced cycles of length 6 and we show that $\text{mim-width}(H) \leq 2$.

Let T be a caterpillar satisfying (Q, P) . For each $p_i \in P$, define

$$C_i = \{p_i\} \cup \{u_i^q \mid q \in Q, p_i \in q\} \cup \{\gamma_i^q \mid q \in Q, p_i \in q\},$$

that is, C_i consists of all vertices in $P \cup U \cup \Gamma$ whose index is i .

We define a new tree T' from T by adding the leaf ω anywhere on the caterpillar. For each i , replace the leaf p_i by a rooted subtree whose leaves are exactly the vertices of C_i . The leaf p_i is attached directly to the root of this subtree, and all the other leaves lie on the other branch (see Figure 4).

We prove that $\text{mim-width}_G(T') \leq 2$, i.e., $\text{mim}_G(A_e, B_e) \leq 2$ for every edge e of T' . Fix any edge e of T' .

Let $a_1 b_1, a_2 b_2, a_3 b_3 \in E(G)$ be three edges with pairwise disjoint endpoints such that $a_1, a_2, a_3 \in A_e$ and $b_1, b_2, b_3 \in B_e$. Suppose, for contradiction, that these three edges form an induced matching in $G[A_e, B_e]$.

First note that none of the six endpoints lies in P . Indeed, each $p_i \in P$ is adjacent exactly to the vertices of C_i , and by the placement of p_i in T' , any cut (A_e, B_e) with $p_i \in A_e$ that separates p_i from some vertex of C_i must be of one of two forms. Either $A_e = \{p_i\}$, then there is no matching of size 3; or $B_e \subseteq C_i$ because e is in the subtree containing $C_i \setminus \{p_i\}$, then, since p_i is adjacent to every vertex of C_i , the matching would not be induced.

As in the previous proof, we say that a vertex v of G is of *type* q for a quartet q if v is one of the eight vertices in $U \cup \Gamma$ corresponding to q . Suppose a_1b_1 is an edge with a_1 of type q and b_1 of type $q' \neq q$. In our construction of G , the only vertices nonadjacent to a_1 are those in P , those of type q , and possibly ω . Since no endpoint is in P , it follows that b_2 and b_3 must be of type q or be ω ; symmetrically, a_2 and a_3 must be of type q' or be ω . If, without loss of generality, a_3 is the only vertex that might be ω , then a_2 is of type q' and b_3 is of type q , which forces $a_2b_3 \in E(G)$, contradicting that the three edges form an induced matching.

Hence at most one of the six endpoints is ω , and all remaining endpoints in $U \cup \Gamma$ belong to a common type (say q). Thus the three matching edges must be among the following kinds: $\gamma\text{--}\gamma$, $u\text{--}u$, or $\omega\text{--}u$.

We cannot have two $\gamma\text{--}\gamma$ edges, since Γ is a clique, which would create an extra edge between γ -endpoints and violate the induced property.

Assume there are two $u\text{--}u$ edges. The only possibilities are $u_i^q u_j^q$ and $u_k^q u_\ell^q$. Without loss of generality, let $u_i^q, u_k^q \in A_e$ and $u_j^q, u_\ell^q \in B_e$. This implies that in T' , the $u_i^q - u_j^q$ path intersects the $u_k^q - u_\ell^q$ path (in particular, they intersect at the ends of e). By our construction of T' from T , this in turn implies that in T' , the $p_i - p_j$ path intersects the $p_k - p_\ell$ path, and thus T' does not satisfy q . Since our transformation from T' to T does not alter its quartets but only adds new ones, this contradicts that T satisfies q .

Moreover, we cannot have simultaneously a $u\text{--}u$ edge and an $\omega\text{--}u$ edge, because ω is adjacent to all u -vertices, creating a forbidden cross-edge and destroying the induced property. Since there is at most one $\omega\text{--}u$ edge, any third edge would have to be either $u\text{--}u$ or $\gamma\text{--}\gamma$, both of which have been ruled out. Therefore no induced matching of size 3 can exist, a contradiction.

This proves that $\text{mim}_G(A_e, B_e) \leq 2$ for every edge e of T' , and hence $\text{mim-width}_G(T') \leq 2$. $\square$

Theorem 2. *The problem of deciding whether a graph has mim-width or linear mim-width at most 2 is NP-complete. Moreover, assuming the ETH, it cannot be solved in time $2^{o(n)}$, where n is the number of vertices of the graph.*

Proof. We proceed as in Theorem 1. Take an instance I of BETWEENNESS and let (Q, P) be the UQC instance obtained from Theorem 3. Then let G and H be obtained from (Q, P) as described above. If I is a YES-instance, then $(Q, P) \in \text{UQC}$ and by Theorem 3 some ternary caterpillar satisfies Q . Then by Proposition 2 we have $\text{linear-mim-width}(G) = 2$ and $\text{mim-width}(H) = 2$. If I is a NO-instance, then $(Q, P) \notin \text{UQC}$ and by Proposition 2 we have $\text{linear-mim-width}(G) \geq 3$ and $\text{mim-width}(H) \geq 3$. Thus I is a YES-instance of BETWEENNESS if and only if $\text{linear-mim-width}(G) \leq 2$, and if and only if $\text{mim-width}(H) \leq 2$, which proves the NP-hardness of our two problems (and membership in NP is easy).

As for the ETH lower bound, we observe that $|V(G)| = |P| + 8|Q|$ and $|V(H)| = |P| + 8|Q| + 1$. Since $|V(G)|$ and $|V(H)|$ are linear in $|P| + |Q|$, Theorem 4 implies that no time $2^{o(n)}$ algorithm can recognize mim-width 2 and linear-mim-width 2 graphs under the ETH. $\square$

5 Open problems

Many questions remain open. The most important one is whether graphs of mim-width 1 and graphs of linear mim-width 1 can be recognized in polynomial time.

Our result also does not imply that deciding whether $\text{mim-width}(G) \leq c$ is NP-hard for every fixed constant $c > 2$. One possible direction is to address the following question: Is there a graph operation that transforms any graph G into a graph G' satisfying $\text{mim-width}(G') = \text{mim-width}(G) + 1$? Alternatively, is there a graph operation that multiplies the mim-width of any graph by a constant factor? The same question remains open for other width parameters as well.

Another question is whether UQC can be used to prove hardness of other graph-to-tree representations?

Finally, can one obtain inapproximability results, for instance via a reduction that yields an arbitrarily large approximation gap?

References

- [Ant+24] Dhanyamol Antony, Yixin Cao, Sagartanu Pal, and RB Sandeep. “Switching classes: Characterization and computation”. In: *arXiv preprint arXiv:2403.04263* (2024).
- [BBD25] Benjamin Bergougnoux, Édouard Bonnet, and Julien Duron. “Mim-Width Is paraNP-Complete”. In: *52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025)*. Ed. by Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis. Vol. 334. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025, 25:1–25:17.
- [BDJ23] Benjamin Bergougnoux, Jan Dreier, and Lars Jaffke. “A logic-based algorithmic meta-theorem for mim-width”. In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2023, pp. 3282–3304.
- [Ben+22] Bergougnoux Benjamin, Svein Høgemo, Jan Arne Telle, and Martin Vatshelle. “Recognition of Linear and Star Variants of Leaf Powers is in P”. In: *Graph-Theoretic Concepts in Computer Science*. Ed. by Michael A. Bekos and Michael Kaufmann. Cham: Springer International Publishing, 2022, pp. 70–83.
- [BKR23] Benjamin Bergougnoux, Tuukka Korhonen, and Igor Razgon. “New width parameters for independent set: One-sided-mim-width and neighbor-depth”. In: *International Workshop on Graph-Theoretic Concepts in Computer Science*. Springer. 2023, pp. 72–85.

- [BPT22] Benjamin Bergougnoux, Charis Papadopoulos, and Jan Arne Telle. “Node multiway cut and subset feedback vertex set on graphs of bounded mim-width”. In: *Algorithmica* 84.5 (2022), pp. 1385–1417.
- [BS01] David Bryant and Mike Steel. “Constructing optimal trees from quartets”. In: *Journal of Algorithms* 38.1 (2001), pp. 237–259.
- [BV13] Rémy Belmonte and Martin Vatshelle. “Graph classes with structured neighborhoods and algorithmic applications”. In: *Theoretical Computer Science* 511 (2013), pp. 54–65.
- [Høg25] Svein Høgemo. “Mapping Graphs to Trees: Partition Trees, Leaf Powers and Decompositions”. In: *PhD thesis. The University of Bergen* (2025).
- [IP01] Russell Impagliazzo and Ramamohan Paturi. “On the complexity of k-SAT”. In: *Journal of Computer and System Sciences* 62.2 (2001), pp. 367–375.
- [Jaf+19] Lars Jaffke, O-joung Kwon, Torstein JF Strømme, and Jan Arne Telle. “Mim-width III. Graph powers and generalized distance domination problems”. In: *Theoretical Computer Science* 796 (2019), pp. 216–236.
- [JKT20a] Lars Jaffke, O-joung Kwon, and Jan Arne Telle. “Mim-width I. Induced path problems”. In: *Discrete Applied Mathematics* 278 (2020), pp. 153–168.
- [JKT20b] Lars Jaffke, O-joung Kwon, and Jan Arne Telle. “Mim-width II. The feedback vertex set problem”. In: *Algorithmica* 82.1 (2020), pp. 118–145.
- [Kan+17] Dong Yeap Kang, O-joung Kwon, Torstein JF Strømme, and Jan Arne Telle. “A width parameter useful for chordal and co-comparability graphs”. In: *Theoretical Computer Science* 704 (2017), pp. 1–17.
- [MY23] Andrea Munaro and Shizhou Yang. “On algorithmic applications of sim-width and mim-width of (H1, H2)-free graphs”. In: *Theoretical Computer Science* 955 (2023), p. 113825.
- [Opa79] J. Opatrny. “Total Ordering Problem”. In: *SIAM Journal on Computing* 8.1 (1979), pp. 111–114.
- [OST24] Yota Otachi, Akira Suzuki, and Yuma Tamura. “Finding Induced Subgraphs from Graphs with Small Mim-Width”. In: *19th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2024)*. Ed. by Hans L. Bodlaender. Vol. 294. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, 38:1–38:16.
- [Rao08] Michaël Rao. “Clique-width of graphs defined by one-vertex extensions”. In: *Discrete Mathematics* 308.24 (2008), pp. 6157–6165.

- [RBR14] Rezwana Reaz, Md Shamsuzzoha Bayzid, and M Sohel Rahman. “Accurate phylogenetic tree reconstruction from quartets: A heuristic approach”. In: *PloS one* 9.8 (2014), e104008.
- [Ste92] Michael Steel. “The complexity of reconstructing trees from qualitative characters and subtrees”. In: *Journal of Classification* 9.1 (Jan. 1992), pp. 91–116.
- [SV16] Sigve Hortemo Sæther and Martin Vatshelle. “Hardness of computing width parameters based on branch decompositions over the vertex set”. In: *Theoretical Computer Science* 615 (2016), pp. 120–125.
- [TLN25] Max Dupré la Tour, Manuel Lafond, and Ndiame Ndiaye. *Recognizing Leaf Powers and Pairwise Compatibility Graphs is NP-Complete*. 2025.
- [Vat12] Martin Vatshelle. “New width parameters of graphs”. In: *PhD Thesis. The University of Bergen* (2012).
- [Zie+18] Emile Ziedan, Deepak Rajendraprasad, Rogers Mathew, Martin Charles Golumbic, and Jérémie Dusart. “The Induced Separation Dimension of a Graph”. In: *Algorithmica* 80.10 (Oct. 2018), pp. 2834–2848.