

Connected tree-width

Reinhard Diestel and Malte Müller

May 21, 2016

Abstract

The *connected tree-width* of a graph is the minimum width of a tree-decomposition whose parts induce connected subgraphs. Long cycles are examples of graphs that have small tree-width but large connected tree-width. We show that a graph has small connected tree-width if and only if it has small tree-width and contains no long geodesic cycle.

We further prove a connected analogue of the duality theorem for tree-width: a finite graph has small connected tree-width if and only if it has no bramble whose connected covers are all large. Both these results are qualitative: the bounds are good but not tight.

We show that graphs of connected tree-width k are k -hyperbolic, which is tight, and that graphs of tree-width k whose geodesic cycles all have length at most ℓ are $\lfloor \frac{3}{2}\ell(k-1) \rfloor$ -hyperbolic. The existence of such a function $h(k, \ell)$ had been conjectured by Sullivan.

1 Introduction

Let us call a tree-decomposition $(T, (V_t)_{t \in T})$ of a graph G *connected* if its parts V_t are connected in G . For example, the standard minimum width tree-decomposition of a tree or a grid is connected. The *connected tree-width* $\text{ctw}(G)$ of G is the minimum width that a connected tree-decomposition of G can have.

Much of the practical use of tree-decompositions, connected or not, derives from the fact that G reflects the nested edge-separations of T obtained by deleting a single edge: these correspond to nested vertex-separations of G . If the tree-decomposition and G are connected, we also have a converse: then every (connected) subtree of T induces a connected subgraph of G .

We shall indicate below some contexts in which connected tree-width has been used for applications. The purpose of this paper, however, is to answer the natural first question one would ask: does connected tree-width differ from ordinary tree-width, and if so how? The answer will be unexpectedly

satisfying: we shall find an obstruction to when the two parameters are tied (in the sense that each is bounded by a function of the other), and be able to show that it is the only obstruction. Let us make this precise.

The ordinary tree-width $\text{tw}(G)$ of a graph G is clearly at most its connected tree-width, so large tree-width causes a graph to have large connected tree-width. But it is not the only possible cause.

A k -cycle C , for example, has tree-width 2 but connected tree-width $\lceil k/2 \rceil$. Indeed, its only connected subgraphs are its segments. In any tree-decomposition none of whose parts contains another, the intersection of two adjacent parts separates them in the whole graph [4, Lemma 12.3.1]. But the intersection of two such segments of C only separates them in C if they cover C – in which case one of them has at least $\lceil k/2 \rceil$ edges.

Long cycles as subgraphs do not necessarily raise the connected tree-width: consider wheels. But long geodesic cycles do: we shall be able to show that every graph containing a k -cycle geodesically also has connected tree-width at least $\lceil k/2 \rceil$ (Lemma 7.1). Here, a path or cycle in G is *geodesic* if it is an isometric subgraph of G . A subgraph $H \subseteq G$ is *isometric* in G if $d_H(x, y) = d_G(x, y)$ for all its vertices x and y , where $d(x, y)$ is the length of a shortest x - y path in H or G , respectively. Note that geodesic u - v paths are simply shortest u - v paths.

Our main theorem says that, conversely, large tree-width and long geodesic cycles are the only two obstructions to small connected tree-width:

Theorem 1.1. *The connected tree-width of a graph G is bounded above by a function of its tree-width and of the maximum length of its geodesic cycles.*

Specifically, if G is not a forest, $\text{tw}(G) < k \in \mathbb{N}$, and ℓ is the maximum length of a geodesic cycle in G , then $\text{ctw}(G) < f(k, \ell)$ for

$$f(k, \ell) = k + \binom{k}{2} (\ell(k - 2) - 1).$$

Theorem 1.1 is qualitatively best possible in that the two obstructions are independent: a large cycle (as a graph) contains a large geodesic cycle but has small tree-width, while a large grid has large tree-width but all its geodesic cycles are small. And both graphs have large connected tree-width.

Quantitatively, the bound has recently been improved by Hamann and Weißbauer [7].

Among the classical¹ obstructions to small ordinary tree-width there is one that gives a tight duality theorem: large-order brambles. A *bramble* is

¹See [5] for a more recent duality theorem for tree-width that is also tight.

a set of pairwise touching connected sets of vertices, where two vertex sets *touch* if they share a vertex or the graph has an edge between them. A set of vertices *covers* (or is a *cover* of) a bramble if it has a vertex in each of these sets. The *order* of a bramble is the smallest size of a cover.

Tree-width duality theorem (Seymour & Thomas [11, 4]). *A graph has tree-width at least $k \geq 0$ if and only if it contains a bramble of order $> k$.*

To adapt this duality to connected tree-width, let the *connected order* of a bramble be the least order of a connected cover, a cover inducing a connected subgraph. Given any tree-decomposition of a graph, it is well known and easy to show that every bramble is covered by some part of that tree-decomposition. Hence graphs of connected tree-width $< k$ cannot have brambles of connected order $> k$. Theorem 1.1 will enable us to prove a qualitative converse of this:

Theorem 1.2. *There is a function $g: \mathbb{N} \rightarrow \mathbb{N}$ such that every graph with no bramble of connected order $> k$ has connected tree-width $< g(k)$.*

We conjecture that, just as in the case of ordinary tree-width, the duality is also quantitatively tight:

Conjecture. *A graph has connected tree-width at least $k \geq 0$ if and only if it contains a bramble of connected order $> k$.*

We believe that the notion of connected tree-width is natural enough to merit further study: tree-decompositions, after all, are meant to exhibit the tree-like structure of a graph, and this aim is better served when its parts are connected (as in a minimum-width decomposition of a tree) than if not (as in a minimum-width decomposition of a long cycle). However, let us give a few pointers to how it relates to other contexts, both in graph theory and beyond.

One of these is hyperbolic graphs. These are important especially when they are infinite and appear as Cayley graphs of hyperbolic groups. But hyperbolicity has also been exploited for the design of algorithms. We shall prove the following (see Section 6 for a more detailed statement and a definition of hyperbolicity):

Theorem 1.3. *Graphs of connected tree-width k are k -hyperbolic. This is best possible for every $k > 1$.*

Another width parameter, somewhat related to connected tree-width, was introduced recently by Dourisboure and Gavaille [6] and has since received some attention. They call it the *tree-length* of a graph G : the smallest value, minimized over all its tree-decompositions, of the maximum distance in G of any two vertices in a common part of this decomposition. Clearly, graphs of connected tree-width $< k$ also have tree-length $< k$, so by Theorem 1.1 the tree-length of G is bounded by the same function $f(k, \ell)$ as its connected tree-width. In fact, Reidl and Sullivan [12, 1] observed that the following better bound follows directly from one of our lemmas for the proof of Theorem 1.1:

Theorem 1.4. *If G has tree-width $< k$ and no geodesic cycle longer than ℓ , and G is not a forest, then the tree-length of G is at most $\ell(k - 2)$.*

Chepoi et al. [3] showed that graphs of tree-length at most k are $4k$ -hyperbolic. We shall give a simple direct proof showing that they are $\lfloor \frac{3}{2}k \rfloor$ -hyperbolic, which is best possible.

Theorem 1.4 thus implies that graphs of tree-width $< k$ that have no geodesic cycle longer than ℓ are $\lfloor \frac{3}{2}\ell(k - 2) \rfloor$ -hyperbolic. This confirms the conjecture of Sullivan [12] that the hyperbolicity of graphs can be bounded in terms of their tree-width and the maximum length of their geodesic cycles. Note that bounding just one of these parameters will not imply hyperbolicity: long cycles have small tree-width and grids have no geodesic cycles of length > 4 , but neither of these (classes of) graphs is hyperbolic.

Finally, connected tree-width has been used directly for the design of algorithms. Soon after this paper first appeared on the ArXiv [10], Jegou and Terrioux [8, 9] used connected tree-width in the context of constraint satisfaction problems in AI. They show how practical CSP algorithms can be improved if the constraint network considered has bounded connected tree-width. They also study the complexity of computing the connected tree-width of a graph.

We use the terminology of [4]. In particular, the *length* of a path is its number of edges, and the *distance* $d_G(u, v)$ between vertices u, v in a graph G is the minimum length of a u - v path in G . Given vertices u, v in a tree T , possibly just a path, we write uTv for the unique path in T between u and v . We also assume familiarity with the basic theory of tree-decompositions as described in [4, Ch.12.3], in particular with their axioms (T1)–(T3). The proof of our main result, Theorem 1.1, will be given in Sections 2–5. In Section 7 we use it to prove Theorem 1.2. Theorems 1.3–1.4 are proved in Section 6.

2 Making a tree-decomposition connected

Our approach to obtaining an upper bound for the connected tree-width of a graph G will be to start with an ordinary tree-decomposition of G of low width, and then to make its parts connected by adding paths joining its components in G . Since adding vertices to a part may invalidate axiom (T3) of a tree-decomposition, we may have to add the same vertices to other parts too. Our task will thus be to ensure that the number of vertices added to any given part in this way, either explicitly in order to decrease its number of components or implicitly when repairing (T3), remains bounded.

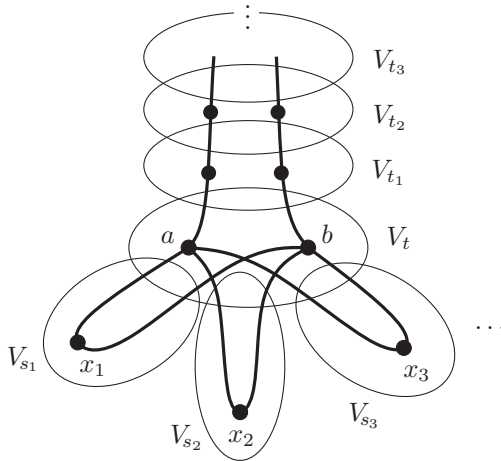


Figure 1: Making V_{t_1}, V_{t_2}, \dots connected can inflate V_t unboundedly

For example, consider the graph and tree-decomposition indicated in Figure 1. If for $i = 1, 2, \dots$ we use the path through x_i to make V_{t_i} connected, we will have to add all the x_i to the central part V_t while repairing (T3), because t lies between s_i and t_i in T .

Of course we made a bad choice here: we could use the path through x_1 for every V_{t_i} . This is the idea behind the following definition: if we already know a path joining two vertices a and b , we can re-use it whenever we have to connect a and b later.

Definition 2.1 (Navigational path systems, or *navs*). Let $G = (V, E)$ be a graph, $K \subseteq [V]^2$ a set of 2-element subsets $xy := \{x, y\}$ of its vertex set, and $\mathcal{N} = (P_{xy})_{xy \in K}$ a family of paths between these, where P_{xy} links x to y .

- (i) \mathcal{N} is a K -nav if for every path $P \in \mathcal{N}$ and for every two vertices $a, b \in P$ we have $ab \in K$ and $P_{ab} = aPb$.

- (ii) A *nav* is a K -nav for $K = [V]^2$.
- (iii) When $\mathcal{D} := (T, (V_t)_{t \in T})$ is a tree-decomposition of G , then a K -nav satisfying $[V_t]^2 \subseteq K$ for every $t \in T$ is a \mathcal{D} -nav.
- (iv) \mathcal{N} is *geodesic* if every $P_{xy} \in \mathcal{N}$ is a shortest x - y path in G .
- (v) The *length* $\ell(\mathcal{N})$ of \mathcal{N} is the maximum length of a path in \mathcal{N} .

We shall use navs to make tree-decompositions connected, as follows:

Lemma 2.2. *Let G be a graph, $\mathcal{D} = (T, (V_t)_{t \in T})$ a tree-decomposition of G of width $< k$, and $\mathcal{N} = (P_{uv})_{uv \in K}$ a \mathcal{D} -nav of G . For all $t \in T$ let*

$$W_t := V_t \cup \bigcup_{xy \in [V_t]^2} V(P_{xy})$$

Then $(T, (W_t)_{t \in T})$ is a connected tree-decomposition of G of width less than

$$k + \binom{k}{2} (\ell(\mathcal{N}) - 1).$$

Proof. The family $(T, (W_t)_{t \in T})$ satisfies the axioms (T1) and (T2) for tree-decompositions, because \mathcal{D} does. To prove (T3) let t_1, t_2 and t_3 be distinct nodes of T with $t_2 \in t_1 T t_3$, and let $z \in W_{t_1} \cap W_{t_3}$ be given. We have to show that $z \in W_{t_2}$.

If $z \in V_{t_2}$ this is the case, so we assume that $z \notin V_{t_2}$. Then V_{t_2} , which separates V_{t_1} from V_{t_3} in G , also separates z from one of these, say from V_{t_3} . Hence as $z \in W_{t_3}$ there are $x, y \in V_{t_3}$ such that $z \in P_{xy}$, and z lies on a subpath $uP_{xy}v$ with $u, v \in V_{t_2}$. Since \mathcal{N} is a \mathcal{D} -nav, this subpath coincides with $P_{uv} \subseteq W_{t_2}$, so $z \in W_{t_2}$ as desired.

The sets W_t are connected in G by construction. Since $|V_t| \leq k$ by assumption, and each path in \mathcal{N} has at most $\ell(\mathcal{N}) - 1$ inner vertices, they also satisfy $|W_t| \leq k + \binom{k}{2} (\ell(\mathcal{N}) - 1)$. \square

Lemma 2.2 reduces our task of making a given tree-decomposition \mathcal{D} connected without increasing its width too much to that of finding a \mathcal{D} -nav of bounded length. This will be achieved by a geodesic nav of G .

While the existence of some nav in a connected graph is trivial – for example, we can route all its paths through some fixed spanning tree – the existence of a geodesic nav is more surprising. Spanning trees, for example, are never geodesic unless the entire graph is a tree.

Lemma 2.3. *Every connected graph has a geodesic nav.*

Proof. Let $G = (V, E)$ be a connected graph. Fix a linear ordering on V , and consider the following ordering on the subsets of V . If v is the first vertex of V that lies in the symmetric difference of two sets $U_1, U_2 \subseteq V$, let $U_1 < U_2$ if $v \in U_2$. Then $U' \subsetneq U$ implies $U' < U$, and if we replace in U a subset $U' \subseteq U$ with a smaller set $U'' < U'$, the resulting set $(U \setminus U') \cup U''$ is smaller than U .

For every two vertices $x, y \in V$ let P_{xy} be a shortest x - y path in G whose vertex set is smallest in our ordering among all such paths. We claim that $\mathcal{N} := (P_{xy})_{xy \in [V]^2}$ is a nav in G ; if so, it will be geodesic by definition.

Suppose not. Then there are $x, y \in V$ and $a, b \in P_{xy}$ such that P_{ab} differs from $aP_{xy}b$. As both these paths are geodesic, it is easy to see that their vertex sets differ too; then $V(P_{ab}) < V(aP_{xy}b)$ by the choice of P_{ab} . By our earlier observation, replacing $aP_{xy}b$ with P_{ab} in P_{xy} yields a shortest x - y path with a vertex set smaller than $V(P_{xy})$, contrary to the choice of P_{xy} . \square

A geodesic nav $\mathcal{N} = (P_{xy})_{xy \in [V]^2}$ of G induces, for any tree-decomposition $\mathcal{D} = (T, (V_t)_{t \in T})$ of G , a geodesic \mathcal{D} -nav $\mathcal{N}_{\mathcal{D}} := (P_{uv})_{uv \in K_{\mathcal{D}}}$ of G simply by restricting $[V]^2$ to its subset

$$K_{\mathcal{D}} := \bigcup_{t \in T} \bigcup_{xy \in [V_t]^2} [V(P_{xy})]^2;$$

note that $\mathcal{N}_{\mathcal{D}}$ is indeed a nav. Its length $\ell(\mathcal{N}_{\mathcal{D}})$ is bounded by the maximum distance of two vertices in a common part of \mathcal{D} .

We have thus reduced our task further: it only remains to find a tree-decomposition of width $\text{tw}(G)$ such that the distance between any two vertices in a common part is bounded by a function of $\text{tw}(G)$ and of the maximum length of a geodesic cycle in G .

3 Atomic tree-decompositions

We shall work with tree-decompositions that cannot be refined in a certain sense. These were introduced by Thomas [13] and have also been used in [2]. We therefore only sketch the proofs of the few easy lemmas we need. Explicit proofs can be found in [10].

Definition 3.1 (Atomic tree-decomposition). The *fatness* of a tree-decomposition of a graph of order n is the n -tuple (a_0, \dots, a_n) in which a_i denotes the number of parts of order $n - i$. A tree-decomposition of lexicographically minimum fatness is *atomic*.

Clearly, atomic tree-decompositions of a graph G have width $\text{tw}(G)$. The point of the notion is to minimize not just the largest part but the smaller ones too, in turn as their size decreases.

Let $\mathcal{D} = (T, (V_t)_{t \in T})$ be an atomic tree-decomposition of a graph G .

Lemma 3.2. *No part of \mathcal{D} contains another part.*

Proof. If there are $t \neq t'$ such that $V_{t'} \subseteq V_t$, then by (T3) such t, t' can be chosen adjacent. Contracting the edge tt' in T and assigning V_t as a part to the new contracted node yields a tree-decomposition of smaller fatness than \mathcal{D} , a contradiction. \square

Given an edge $e = t_1t_2$ of T and $i \in \{1, 2\}$, let T_i be the component of $T - e$ containing t_i , and put $U_i := \bigcup_{t \in T_i} V_t$ and $G_i := G[U_i]$. Then $\{U_1, U_2\}$ is a separation of G , with $X := U_1 \cap U_2 = V_{t_1} \cap V_{t_2}$ [4, Lemma 12.3.1].

Lemma 3.3. *$G_1 - X$ and $G_2 - X$ have components C_1 and C_2 , respectively, in which every vertex from X has a neighbour.*

Proof. If the neighbourhood of every component of $G_1 - X$ (say) is a proper subset of X , we can obtain a tree-decomposition of smaller fatness than \mathcal{D} by replacing its portion that decomposes G_1 with the decompositions that \mathcal{D} induces on the graphs $G_1[V(C) \cup N(C)]$, where C ranges over the components of $G_1 - X$. This decomposition has smaller fatness, because we lose the part $V_{t_1} \supseteq X$ and gain only parts that are proper subsets of parts we lose.

More formally, we replace T_1 in T with copies $T_{1,C}$ of T_1 , one for every C , making the nodes $t_{1,C}$ corresponding to t_1 adjacent to t_2 . For every C and every $t \in T_1$, we associate $V_{t_C} := V_t \cap (V(C) \cap N(C))$ with the node t_C of $T_{1,C}$ corresponding to t . It can be checked that the decomposition thus obtained has smaller fatness than \mathcal{D} . \square

Lemma 3.3 implies at once:

Corollary 3.4. *Every two vertices $u, v \in X$ are linked by an X -path in G_1 and another in G_2 , which thus form a cycle.* \square

Our final lemma ensures that any two non-adjacent vertices in a common part of \mathcal{D} lie in some such X . We shall then be able to use Corollary 3.4 to find a cycle through them, which will be our aim when we prove Theorem 1.1.

Lemma 3.5. *For all $t \in T$ and distinct $u, v \in V_t$, either $uv \in E(G)$ or t has a neighbour s in T such that $u, v \in V_s$.*

Proof. If not, we split V_t into $V_{t'} := V_t \setminus \{u\}$ and $V_{t''} := V_t \setminus \{v\}$, and replace t in T with new adjacent nodes t', t'' . This yields another tree-decomposition (of smaller fatness) since, by assumption, every V_s with $st \in E(T)$ fails to contain either u or v , so we can join s to t' or t'' , respectively, in the expanded decomposition tree. \square

4 Use of the cycle space

As shown at the end of Section 2, our remaining task for a proof of Theorem 1.1 was to show that, in a suitable tree-decomposition of width $\text{tw}(G) < k$, say, the distance in G between any two vertices in a common part is bounded by a function of k and of the maximum length ℓ of a geodesic cycle in G .

The tree-decomposition that we shall choose will be any atomic one, say $\mathcal{D} = (T, (V_t)_{t \in T})$. By Lemma 3.5 we can then find, for any non-adjacent vertices u, v in a common part of \mathcal{D} , adjacent nodes $t_1, t_2 \in T$ such that $\{u, v\} \subseteq X := V_{t_1} \cap V_{t_2}$. Corollary 3.4 then provides X -paths $P_1 \subseteq G_1$ and $P_2 \subseteq G_2$ linking u to v (the G_i being defined as in Section 3), so that $C_{uv} := P_1 \cup P_2$ will be a cycle.

If these cycles C_{uv} were geodesic, the distances $d_G(u, v)$ would be bounded by $\ell/2$ and achieve our goal. Unfortunately they need not be. However, we shall use an algebraic argument to deduce from their existence that we can link the vertices u, v using only geodesic cycles meeting X , and only a bounded number of these.

Given $X \subseteq V(G)$ and a set \mathcal{C} of cycles in G , let us write

$$\mathcal{C}[X] := \bigcup \{C \in \mathcal{C} \mid V(C) \cap X \neq \emptyset\}.$$

The set \mathcal{C} will later be the set of geodesic cycles in G . The *diameter* of X in G is the maximum distance in G between any two vertices in X .

Lemma 4.1. *If $G[X] \cup \mathcal{C}[X]$ is connected, $1 \leq \ell \in \mathbb{N}$, and every cycle in \mathcal{C} has length at most ℓ , then X has diameter at most $\ell(|X| - 1)$ in G .*

Proof. If $\mathcal{C}[X] = \emptyset$ then $G[X]$ is connected, and the assertion holds. If not, then $\ell \geq 3$, which we shall now assume.

Let $x_1, x_2 \in X$ be given. By assumption, $G[X] \cup \mathcal{C}[X]$ contains an x_1 - x_2 walk W consisting of segments that are each either an edge of $G[X]$ or a path on a cycle in $\mathcal{C}[X]$. By definition of $\mathcal{C}[X]$, these latter segments, too, can be chosen so as to meet X . Cutting loops out of W as necessary we may assume that no two of these segments meet X in the same vertex x , unless they are consecutive segments meeting at x .

Let us combine each pair of consecutive segments $u \dots x$ and $x \dots v$ with $x \in X$ and $u \notin X$ to a new segment $u \dots v$. This can be done simultaneously with all relevant segments since, by our conditions on x and u , a given segment can be combined with at most one of its two adjacent segments. Choosing $u \dots x$ and $x \dots v$ as either an edge of $G[X]$ or the shorter half of a cycle in $\mathcal{C}[X]$, we can ensure that the combined segment has length at most ℓ .

The resulting segments all meet X and have distinct first vertices in X . They each have length at most ℓ . The first segment, starting at x_1 , can in fact be chosen of length at most $\ell/2$. So can the last segment, the one ending at x_2 , unless it is a combined segment, in which case x_2 is not its first vertex in X . Altogether we now have either $|X| - 1$ segments of length at most ℓ , or $|X| - 2$ inner segments of length at most ℓ and two outer segments of length at most $\ell/2$. So W has length at most $\ell(|X| - 1)$, as claimed. \square

In our next lemmas we use $+$ to denote addition in the binary cycle space of G , that is, for the symmetric difference of edge sets. To avoid clutter, we shall not always distinguish notationally between a subgraph and its edge set, or between an edge e and the set $\{e\}$.

Lemma 4.2. *Let $P = u \dots v$ be a path in G , where $uv =: e \notin E(G)$. Let D be an element of the cycle space of G . Then $P + D$ contains a u - v path.*

Proof. Since $(P + e) + D$ lies in the cycle space of $G + e$, it is also an edge-disjoint union of cycles [4, Prop. 1.9.1]. If C is the cycle containing e , then $C - e$ is the desired u - v path in $P + D$. \square

Lemma 4.3. *Let $G = G_1 \cup G_2$, and let $u, v \in X = V(G_1 \cap G_2)$ be non-adjacent in G . Assume that each G_i contains an X -path $P_i = u \dots v$, and that \mathcal{C} is a set of cycles generating $C = P_1 \cup P_2$. Then $\mathcal{C}[X]$ contains a u - v path.*

Proof. As \mathcal{C} generates C , we can find subsets $\mathcal{C}_1, \mathcal{C}_X, \mathcal{C}_2 \subseteq \mathcal{C}$ such that

$$C = \sum \mathcal{C}_1 + \sum \mathcal{C}_X + \sum \mathcal{C}_2,$$

where every cycle in \mathcal{C}_i has its vertices in $G_i - X$ and every cycle in \mathcal{C}_X meets X . By Lemma 4.2 there is a u - v path P in $P_1 + \sum \mathcal{C}_1$. This latter set of edges is disjoint from $P_2 + \sum \mathcal{C}_2$, so

$$P \subseteq P_1 + \sum \mathcal{C}_1 + P_2 + \sum \mathcal{C}_2 = C + \sum \mathcal{C}_1 + \sum \mathcal{C}_2 = \sum \mathcal{C}_X \subseteq \mathcal{C}[X].$$

\square

Lemma 4.4. *Let \mathcal{D} be any atomic tree-decomposition of G , and $1 \leq \ell \in \mathbb{N}$. Let \mathcal{C} be a set of cycles in G , all of length at most ℓ , that generate its cycle space. Then any two non-adjacent vertices in a common part W of \mathcal{D} have distance at most $\ell \cdot (|W| - 2)$ in G .*

Proof. Let $\mathcal{D} = (T, (V_t)_{t \in T})$. Let W be any part of \mathcal{D} , and let $u, v \in W$ be non-adjacent. We show that $d_G(u, v) \leq \ell \cdot (|W| - 2)$.

By Lemma 3.5 we can find adjacent nodes $t_1, t_2 \in T$ such that $\{u, v\} \subseteq X := V_{t_1} \cap V_{t_2}$. For every two non-adjacent vertices u', v' in X , Corollary 3.4 ensures that there are X -paths $P_1 \subseteq G_1$ and $P_2 \subseteq G_2$ linking u' to v' (the G_i being defined as in Section 3). By Lemma 4.3 there is a $u'-v'$ path in $\mathcal{C}[X]$. Thus, $G[X] \cup \mathcal{C}[X]$ is connected.

By Lemma 4.1, the vertices u, v have distance at most $\ell \cdot (|X| - 1)$ in G . By Lemma 3.2 we have $|X| < |W|$, which completes the proof. \square

5 Proof of Theorem 1.1

We are now ready to prove our main result, which we restate.

Theorem 1.1. *The connected tree-width of a graph G is bounded above by a function of its tree-width and of the maximum length of its geodesic cycles.*

Specifically, if G is not a forest, $\text{tw}(G) < k \in \mathbb{N}$, and ℓ is the maximum length of a geodesic cycle in G , then $\text{ctw}(G) < f(k, \ell)$ for

$$f(k, \ell) = k + \binom{k}{2} (\ell(k - 2) - 1).$$

Proof. The connected tree-width of forests coincides with their ordinary tree-width, so we assume that G contains a cycle. Let ℓ denote the maximum length of a geodesic cycle in G , and let $k \in \mathbb{N}$ be such that $\text{tw}(G) < k$. Assume without loss of generality that G is connected.

By Lemma 2.3, G has a geodesic nav \mathcal{N} . Let $\mathcal{D} = (T, (V_t)_{t \in T})$ be any atomic tree-decomposition of G . Then \mathcal{N} induces a geodesic \mathcal{D} -nav $\mathcal{N}_{\mathcal{D}} := (P_{uv})_{uv \in K_{\mathcal{D}}}$ of G , where

$$K_{\mathcal{D}} := \bigcup_{t \in T} \bigcup_{xy \in [V_t]^2} [V(P_{xy})]^2;$$

For each $t \in T$ let

$$W_t := V_t \cup \bigcup_{xy \in [V_t]^2} V(P_{xy}).$$

By Lemma 2.2, $(T, (W_t)_{t \in T})$ is a connected tree-decomposition of G of width less than

$$k + \binom{k}{2} (\ell(\mathcal{N}_{\mathcal{D}}) - 1).$$

It remains to show that $\ell(\mathcal{N}_{\mathcal{D}}) \leq \ell(k - 2)$, i.e., that any two vertices x, y in a common part of \mathcal{D} have distance at most $\ell(k - 2)$ in G . This is clear if x, y are adjacent, since G contains a cycle and hence $k, \ell \geq 3$. If x, y are non-adjacent, it follows from Lemma 4.4, since the geodesic cycles of G generate its cycle space [4, Ex. 1.32] and \mathcal{D} realizes the tree-width of G . \square

6 Tree-length and hyperbolicity

The *tree-length* of a graph G is the smallest value, minimized over all tree-decompositions of G , of the maximum diameter in G of a part of this decomposition [6]. Reidl and Sullivan, see [12, 1], noticed that our lemmas from Sections 3 and 4 can be applied directly to give a better bound on the tree-length of G than the bound of $f(\ell, k)$ implied by Theorem 1.1:

Theorem 1.4. *If G has tree-width $< k$ and no geodesic cycle longer than ℓ , and G is not a forest, then the tree-length of G is at most $\ell(k - 2)$.*

Proof. This follows from Lemma 4.4, as the geodesic cycles of G generate its cycle space and every atomic tree-decomposition realizes its tree-width. \square

Reidl and Sullivan point out that the graph obtained from the $n \times n$ grid by subdividing every edge $m - 1$ times has tree-width n , no geodesic cycle longer than $\ell = 4m$, and its tree-length is at most nm .² The bound of $4m(n - 1)$ offered by Theorem 1.4 is thus off by no more than a factor of 4.

Given a real number $d \geq 0$, a graph is called *d-hyperbolic* if whenever vertices x, y, z are linked by shortest paths $P_{xy} = x \dots y$ and $P_{yz} = y \dots z$ and $P_{xz} = x \dots z$, each of these paths lies within distance d of the other two. Trees, for example, are 0-hyperbolic, since $P_{xz} \subseteq P_{xy} \cup P_{yz}$.

Chepoi et al [3] showed that graphs of tree-length d are $4d$ -hyperbolic.³ It is not hard to make this bound sharp:

²In fact, it has connected tree-width nm : the standard width n tree-decompositions of the $n \times n$ grid induce connected tree-decompositions of its subdivision of width nm , and the well-known brambles of order $n + 1$ in the grid, see e.g. [4, Ex. 12.21], extend to brambles of connected order $nm + 1$. See Section 7 for connected bramble order.

³Their result, Proposition 13, says ' d -hyperbolic'. But they use a different notion of hyperbolicity that translates to ours at a loss of a factor of 4. See their Proposition 1.

Lemma 6.1. *Graphs of tree-length d are $\lfloor \frac{3}{2}d \rfloor$ -hyperbolic. This value is sharp for every $d \geq 0$.*

Proof. Let G be a graph with a tree-decomposition $(T, (V_t)_{t \in T})$ into parts V_t each of diameter at most d in G . Consider vertices $x, y, z \in G$ and shortest paths P_{xy}, P_{yz}, P_{xz} between them as in the definition of hyperbolicity. We show that $P = P_{xz}$ lies within distance $\lfloor \frac{3}{2}d \rfloor$ of $P' = P_{xy} \cup P_{yz}$.

Let $t_0 \dots t_n$ be the path in T between a node t_0 such that $x \in V_{t_0}$ and a node t_n with $z \in V_{t_n}$. Since P is connected and meets V_{t_0} as well as V_{t_n} , it meets every set $S_i := V_{t_i} \cap V_{t_{i-1}}$ with $i = 1, \dots, n$ [4, Lemma 12.3.1], in a vertex $s_i \in S_i$ say. Similarly, P' has a vertex s'_i in every S_i . Put $s_0 := x =: s'_0$ and $s_{n+1} := z =: s'_{n+1}$, and let $S := \{s_0, \dots, s_{n+1}\}$ and $S' := \{s'_0, \dots, s'_{n+1}\}$.

For each $i = 0, \dots, n$ we have $s_i, s_{i+1} \in V_{t_i}$ and hence $d(s_i, s_{i+1}) \leq d$. Since P is a shortest path, its segments $s_i P s_{i+1}$ thus have length at most d , so every vertex of P lies within distance $\lfloor d/2 \rfloor$ of S . Similarly, $d(s_i, s'_i) \leq d$ for every i , so S lies within distance d of $S' \subseteq V(P')$. Hence P lies within distance $\lfloor \frac{3}{2}d \rfloor$ of P' , as claimed.

To show that the bound of $\lfloor \frac{3}{2}d \rfloor$ is sharp, consider a path-decomposition of adhesion 2 into many copies of K_4 and with disjoint adhesion sets. Subdividing every edge $d - 1$ times yields a graph G of tree-length at most d : the tree-decomposition witnessing this is a path-decomposition whose parts are the 4-vertex-sets from the K_4 s together with pendent parts each consisting of a subdivided edge.

It is easy to find in G three vertices x, y, z with shortest paths P_{xy}, P_{yz} and P_{xz} between them such that each subdivided K_4 (other than the first and the last) meets P_{xz} in exactly one subdivided edge P and $P_{xy} \cup P_{yz}$ in exactly one subdivided edge Q disjoint from P (Fig. 2). For each subdivided K_4 except the leftmost and the rightmost one, the vertex closest to the midpoint of $P \subseteq P_{xz}$ has distance exactly $\lfloor \frac{3}{2}d \rfloor$ not only from Q but from all of $P_{xy} \cup P_{yz}$. \square




Figure 2: Geodesic paths between x, y, z showing that $\lfloor \frac{3}{2}d \rfloor$ -hyperbolicity is best possible for tree-length d

Lemma 6.1 and Theorem 1.4 imply:

Corollary 6.2. *If G has tree-width $< k$, has no geodesic cycle longer than ℓ , and G is not a forest, then G is $\lfloor \frac{3}{2}\ell(k-2) \rfloor$ -hyperbolic. \square*

If we strengthen the premise in Lemma 6.1 so as to assume that not just the tree-length but even the connected tree-width is bounded, we also get a stronger bound, which is again sharp:

Theorem 1.3. *The following statements hold for every integer $k > 1$:*

- (i) *Every graph of connected tree-width k is k -hyperbolic.*
- (ii) *There exists a graph of connected tree-width k that is not $(k-1)$ -hyperbolic.*

The graphs of connected tree-width 1 are forests, which are 0-hyperbolic.

Proof. Repeat verbatim the first two paragraphs of the proof of Lemma 6.1, except that the tree-decomposition $(T, (V_t)_{t \in T})$ is now connected and of width k , and we wish to show that P lies within distance k of P' . Let $v \in P$ be any vertex, say in its segment $P_i := s_i P s_{i+1}$.

Let T_i be a spanning tree of $G[V_{t_i}]$. Since P is a geodesic path, P_i is at most as long as the path $Q_i := s_i T_i s_{i+1}$. Let R_i be the unique $s'_i - Q_i$ path in T_i , and let $q_i \in Q_i$ be its last vertex.

Since $s_i P_i s_{i+1} Q_i s_i$ is a closed walk that has $\|P_i\| + \|Q_i\| \leq 2\|Q_i\|$ edges and contains both v and q_i , we have $d(q_i, v) \leq \|Q_i\|$. As R_i is an $s'_i - q_i$ path with at most $\|T_i\| - \|Q_i\| \leq k - \|Q_i\|$ edges, we thus have $d(s'_i, v) \leq k$.




Figure 3: The vertex $v \in P_{xz}$ has distance k from $P_{xy} \cup P_{yz}$

Figure 3 shows a graph G of connected tree-width $k = 5$ that is not $(k-1)$ -hyperbolic. Indeed, G has a path-decomposition into parts of order at most $k+1$ that form roughly vertical slices, each spanning a path in G . Two (non-adjacent) such parts are shown in the figure. To see that G is not $(k-1)$ -hyperbolic, consider the shortest paths P_{xy} , P_{yz} , P_{xz} between x , y and z whose union is the perimeter cycle, and the middle vertex v of P_{xz} . The example can be adapted to other values of $k \geq 3$ by extending the

underlying ladder and altering the number of small vertices on its rungs, and to $k = 2$ by keeping the rungs unsubdivided and contracting the edge below y . \square

In the notation of the proof of Theorem 1.3 (i), replacing the segments $P_i \subseteq P$ with $Q_i \subseteq T_i$ shows that, given a connected tree-decomposition $(T, (V_t)_{t \in T})$ of width $< k$ of G , the distance in G between two vertices $v \in V_t$ and $v' \in V_{t'}$ is at most $k \cdot (d_T(t, t') + 1)$. It thus seems that graphs of bounded connected tree-width are quasi-isometric to their decomposition trees (and, in particular, hyperbolic). However, this is not the case, since the converse inequality may fail: there need not be a constant c_k such that $d_T(t, t') \leq c_k \cdot d_G(v, v')$, even if we choose the tree-decomposition well.

For example, let G be obtained from a long path P by joining a new vertex v to every vertex of P . All its tree-decompositions of low width, connected or not, have a long path in their decomposition tree T : otherwise T would have a node t of large degree, and V_t would be a small separator of G leaving many components; but G has no such separator. So while the distance of two vertices $v, v' \in P$ is always at most 2 in G , the distance in T of nodes t, t' with $v \in V_t$ and $v' \in V_{t'}$ can be arbitrarily large.

7 Duality

We have postponed until now a proof that not only do long cycles themselves have large connected tree-width (as we saw in the Introduction), but more generally that any graph with a long geodesic cycle does. Indeed, this is not entirely straightforward to prove. It is, however, an immediate consequence of connected tree-width duality, our topic in this last section.

Duality theorems for width parameters give us witnesses for large width. A general framework for this, which includes all the classical duality theorems and several others, has recently been developed in [5]. For tree-width, the classical duality is between tree-decompositions and brambles:

Tree-width duality theorem (Seymour & Thomas [11, 4]). *A graph has tree-width at least $k \geq 0$ if and only if it contains a bramble of order $> k$.*

Rephrased as a min-max theorem: 1 + the minimum width of a tree-decomposition equals the maximum order of a bramble.

To adapt this duality to connected tree-width, let the *connected order* of a bramble be the least order of a *connected cover*, a cover spanning a connected subgraph. (See the Introduction, or [4, Ch. 12], for definitions.)

Conjecture. *Let $k \geq 0$ be an integer. A graph has connected tree-width at least k if and only if it contains a bramble of connected order $> k$.*

The proof of the backward implication, or of \geq in the min-max version, is the same as for ordinary tree-width. Given any bramble \mathcal{B} and any tree-decomposition \mathcal{D} , one shows that some part of \mathcal{D} covers \mathcal{B} . Hence no bramble has greater order than the largest size of a part of \mathcal{D} . And if \mathcal{D} is connected, no bramble has greater connected order.

Let us apply this to show that long geodesic cycles raise the connected tree-width of a graph by giving rise to a bramble of large connected order.

Lemma 7.1. *If a graph G contains a geodesic cycle of length $k \in \mathbb{N}$, then G has a bramble of connected order at least $\lceil k/2 \rceil + 1$, and $\text{ctw}(G) \geq \lceil k/2 \rceil$.*

Proof. Let $C \subseteq G$ be a geodesic cycle of length k , and let \mathcal{B} be the set of all connected subsets of $V(C)$ of size $\lceil k/2 \rceil$. This is clearly a bramble. We shall prove that every connected cover X of \mathcal{B} has at least $\lceil k/2 \rceil + 1$ vertices. This will also imply that $\text{ctw}(G) \geq \lceil k/2 \rceil$, as noted above.

Choose $x, y \in X \cap V(C)$ at maximum distance $d_C(x, y) = d_G(x, y)$. If this distance is $k/2 \in \mathbb{N}$, then X contains an x - y path with $\lceil k/2 \rceil + 1$ vertices.

If $d_C(x, y) < k/2$, then X has another vertex z in the longer x - y segment of C . Let C_{xy}, C_{yz}, C_{zx} be the shortest paths on C between their indices. By the choice of x, y, z , none of these contains another, so $C_{xy} \cup C_{yz} \cup C_{zx} = C$. Hence $d(x, y) + d(y, z) + d(z, x) \geq k$, in C and hence also in G .

Consider a spanning tree T of $G[X]$. The lengths of the paths xTy , yTz and zTx sum to at least $d(x, y) + d(y, z) + d(z, x) \geq k$. But no edge of T lies on more than two of these paths, so T has at least $\lceil k/2 \rceil$ edges. Thus, $|X| = |T| \geq \lceil k/2 \rceil + 1$ as desired. \square

We already saw in the Introduction that k -cycles themselves have connected tree-width at least $\lceil k/2 \rceil$. Lemma 7.1 says that graphs containing such cycles geodesically inherit this. This may lead one to suspect that, whenever $H \subseteq G$ is embedded in G isometrically, $\text{ctw}(H) \leq \text{ctw}(G)$. The graph G of Figure 4, however, shows that this is not so. Indeed, G itself has a connected tree-decomposition of width 3, with v and its neighbours as a central part of order 4. Its isometric subgraph $H = G - v$, however, has a bramble of connected order 5 (and hence $\text{ctw}(H) \geq 4$): the set of paths of order 4 on its outer 9-cycle.

Still, it would be interesting to explore further the relationship between both the connected tree-width and the largest connected order of brambles of a graph and the corresponding parameters of its isometric subgraphs.




Figure 4: $G - v$ is isometric in G but has larger connected tree-width

While deleting edges can increase the connected tree-width of a graph (for example, delete the inner rungs of a long ladder), it is easy to see that contracting edges cannot increase the connected tree-width.

Let us finish by proving a qualitative connected version of the hard implication of the tree-width duality theorem, to establish at least a weakening of our connected tree-width duality conjecture:

Theorem 1.2. *There is a function $g: \mathbb{N} \rightarrow \mathbb{N}$ such that every graph of connected tree-width at least $g(k)$ has a bramble of connected order $> k$.*

Proof. Let $g(k) = 2$ for $k \leq 2$, and $g(k) = f(k, 2k - 2)$ for $k > 2$, where f is the function from Theorem 1.1.

Let G be a graph with no bramble of connected order $> k \in \mathbb{N}$. If G is a forest then $\text{ctw}(G) < 2 \leq g(k)$, as desired. Assume now that G is not a forest. Then $k > 2$, because cycles contain brambles of order 3.

Since every bramble in G has (connected) order $\leq k$, we have $\text{tw}(G) < k$ by the classical tree-width duality theorem, and every geodesic cycle in G has length at most $2k - 2$ by Lemma 7.1. By Theorem 1.1 this implies $\text{ctw}(G) < f(k, 2k - 2) = g(k)$, as desired. \square

8 Acknowledgement

We thank Daniel Weißauer for pointing out an error in an earlier version of Lemma 3.3. Hamann and Weißauer [7] have recently disproved our conjecture in Section 7 by constructing a graph whose connected tree-width is larger than the largest connected order of any of its brambles. They have proposed a modified conjecture, which is still open.

References

- [1] A.B. Adcock, B.D. Sullivan, and M.W. Mahoney. Tree decompositions and social graphs. arXiv:1411.1546, 2014.
- [2] P. Bellenbaum and R. Diestel. Two short proofs concerning tree-decompositions. *Comb., Probab. Comput.*, 11:1–7, 2002.
- [3] Victor Chepoi, Feodor Dragan, Bertrand Estellon, Michel Habib, and Yann Vaxès. Diameters, centers, and approximating trees of delta-hyperbolic geodesic spaces and graphs. In *Proceedings of the twenty-fourth annual symposium on computational geometry*, SCG '08, pages 59–68, New York, NY, USA, 2008. ACM.
- [4] R. Diestel. *Graph Theory*. Springer, 4th edition, 2010.
- [5] R. Diestel and S. Oum. Unifying duality theorems for width parameters, I. Weak and strong duality. arXiv:1406.3797, 2014.
- [6] Y. Dourisboure and C. Gavaille. Tree-decompositions with bags of small diameter. *Discrete Math.*, 307:2008–2029, 2007.
- [7] M. Hamann and D. Weißbauer. Bounding connected tree-width. arXiv:1503.01592, 2015.
- [8] P. Jegou and C. Terrioux. Connected tree-width: a new parameter for graph decomposition.
http://www.cs.uic.edu/pub/Isaim2014/WebPreferences/ISAIM2014_Jegou_Terrioux
2014.
- [9] P. Jegou and C. Terrioux. Tree-decompositions with connected clusters for solving constraint networks. In B. O’Sullivan, editor, *CP 2014*, volume 8656, pages 403–407. Springer Lecture Notes in Computer Science, 2014.
- [10] Malte Müller. Connected tree-width. arXiv:1211.7353, 2012.
- [11] P. Seymour and R. Thomas. Graph searching and a min-max theorem for tree-width. *J. Combin. Theory (Series B)*, 58(1):22–33, 1993.
- [12] Blair Sullivan. Personal communication, Dagstuhl 2013.
- [13] R. Thomas. A Menger-like property of tree-width; the finite case. *J. Combin. Theory (Series B)*, 48:67–76, 1990.