Hyperbolicity Computation through Dominating Sets

David Coudert, André Nusser, Laurent Viennot

► To cite this version:

David Coudert, André Nusser, Laurent Viennot. Hyperbolicity Computation through Dominating Sets. ALENEX 2022 - SIAM Symposium on Algorithm Engineering and Experiments, Jan 2022, Alexandria, VA, United States. pp.78-90, 10.1137/1.9781611977042.7 . hal-03431155v2

HAL Id: hal-03431155

<https://hal.inria.fr/hal-03431155v2>

Submitted on 22 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Hyperbolicity Computation through Dominating Sets*

David Coudert¹, André Nusser², and Laurent Viennot³

¹Université Côte d’Azur, Inria, CNRS, I3S, France

²Max Planck Institute for Informatics and Graduate School of Computer Science, Saarland Informatics Campus, Saarbrücken, Germany

³Inria, Paris University, CNRS, Irif, France

Abstract

Hyperbolicity is a graph parameter related to how much a graph resembles a tree with respect to distances. Its computation is challenging as the main approaches consist in scanning all quadruples of the graph or using fast matrix multiplication as building block, both are not practical for large graphs. In this paper, we propose and evaluate an approach that uses a hierarchy of distance- k dominating sets to reduce the search space. This technique, compared to the previous best practical algorithms, enables us to compute the hyperbolicity of graphs with unprecedented size (up to a million nodes) and speeds up the computation of previously attainable graphs by up to 3 orders of magnitude while reducing the memory consumption by up to more than a factor of 23.

Keywords: Gromov hyperbolicity; graph algorithms; algorithm engineering.

1 Introduction

This paper aims at computing the hyperbolicity of graphs with millions of nodes. Hyperbolicity is a graph parameter used to classify complex networks [1, 5, 37]. It can be seen as a measure of how much a network is “democratic” [4, 11]. It has recently received growing attention as it appears to capture important properties of several large practical graphs such as Internet [42], the Web [41] and database relations [46]. Additionally, measuring hyperbolicity has applications in routing [9, 16, 38], and bioinformatics [14, 30]. See [1, 28] for recent surveys.

Hyperbolicity is usually defined through a 4-points condition that associate to each quadruple a δ -value defined through distances between the four nodes (see Section 2 for the formal definition). The hyperbolicity of a graph is the maximum of the δ -value over all quadruples in the graph. This definition trivially results in a $\Theta(n^4)$ algorithm for computing hyperbolicity, and the best known theoretical complexity is $O(n^{3.69})$, relying on an optimized (max,min)-matrix product [31]. However, the algorithms exhibiting the best performances in practice have time complexity in $\mathcal{O}(n^4)$ -time [12, 18, 21].

*This work has been supported by the French government, through the UCA^{JEDI} Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-15-IDEX-01, the ANR project Multimod with the reference number ANR-17-CE22-0016 and the ANR project Distancia with reference number ANR-17-CE40-0015.

The quest for computing the hyperbolicity of large graphs has lead to gain orders of magnitude in the size of the considered graphs. Indeed, the first attempts are based on brute force implementations of the trivial $\Theta(n^4)$ algorithm [14, 2]. More precisely, [14, 13] uses the revolving door algorithm to explore all quadruples, thus enabling to compute the hyperbolicity of graphs with few hundreds of nodes. Then, using massive parallelism (up to 1000 cores), [2] was able to consider graphs with up to 8000 nodes, but this approach is not scalable. The first noticeable progress is due to [18] which introduces pruning techniques to drastically reduce the number of quadruples to consider. With the addition of refined pruning techniques [12, 43], it enables to compute the hyperbolicity of graphs with up to 50000 nodes. Furthermore, the running time of this algorithm using a single core is orders of magnitude smaller than the running times reported in [2]. However, this algorithm reaches a memory bottleneck as it has space complexity in $\mathcal{O}(n^2)$. To go beyond this bottleneck, [21] engineered an algorithm, along with suitable data structures, that consumes significantly less memory while offering good performances in practice. It enables to compute the hyperbolicity of graphs with more than 100000 nodes. Nonetheless, the memory usage of this algorithm is still high, which limits its scalability. In this paper, we propose a new approach that uses a hierarchy of distance- k dominating sets to both reduce memory usage and further prune the search space, resulting in an algorithm enabling to compute for the first time the hyperbolicity of graphs with up to a million nodes.

1.1 Our approach.

In this paper, we propose to reduce the search space through the use of a hierarchy of distance- k dominating sets. The approach applies to graphs having small distance- k dominating sets for small k , which is often the case in practice. Recall that a distance- k dominating set is a set D such that any node of the graph is at distance at most k from a node in D . The main idea is to explore quadruples within the dominating set. Only when the δ -value of such a quadruple is large enough (depending on the highest δ -value found so far and k), we explore recursively the quadruples dominated by it. More precisely, each node is associated to its closest dominator, that is a node of D at distance at most k . Each quadruple of the graph is said to be dominated by the four associated dominators. For each dominator, we compute recursively a distance- k' dominating set of the nodes associated to it with $k' < k$, and repeat recursively this process with smaller and smaller values of k' . The search explores recursively quadruples of dominators in distance- k' dominating sets with smaller and smaller values of k' . At the deepest level, the search uses $k' = 0$ and all quadruples with potential high δ -value are explored. Although many quadruples are usually skipped, the quadruples with highest δ -value are guaranteed to be visited and the computation is exact. Besides having small distance- k dominating sets, the efficiency of the method also relies on the property that most of the quadruples have a relatively small δ -value which seems to be the case in practice when the hyperbolicity of the graph is large enough compared to k . More precisely, the graph should have hyperbolicity greater than $2k$ and a k -dominating set of few thousand nodes at most.

1.2 Main contributions.

Our main contribution is a new approach for scanning nodes of a graph through a hierarchy of dominating sets. In particular we show its effectiveness for computing hyperbolicity through a set of technical lemmas allowing to relate the delta value of a quadruple to that of four nodes

dominating them in a distance- k dominating set. We also provide an implementation [22] proving the efficiency of the method by computing the hyperbolicity of practical graphs with unprecedented size.

1.3 Other related works.

Graphs with hyperbolicity at most 1 can be recognized efficiently. Indeed, 0-hyperbolic graphs are block graphs (graphs whose biconnected components are cliques) [6, 36], which can be recognized in time $\mathcal{O}(n + m)$. Furthermore, using a 2-approximation algorithm, one can decide in time $\mathcal{O}(n^{2.69})$ if a graph has hyperbolicity at most 1 [31]. Finally, deciding if a graph is $\frac{1}{2}$ -hyperbolic is equivalent to decide if it contains an induced 4-cycle [19], which can be done in time $\mathcal{O}(n^{3.26})$ [39].

Preprocessing methods for reducing the size of the input graph have been proposed. In particular, the hyperbolicity of a graph is the maximum over the hyperbolicity of its biconnected components (see [18] for a proof). Taking the maximum hyperbolicity over the atoms of a decomposition of a graph by clique-minimal separators [45, 8] results in an additive +1 approximation of hyperbolicity and [17] shows how to modify the atoms to get the exact hyperbolicity of the graph. Moreover, [43] proved that the hyperbolicity of a graph is the maximum of the hyperbolicity of the graphs resulting from either a *modular* [32, 35] or a *split* [24, 23] decomposition of the graph and [20] proposed algorithms with time complexity in $\mathcal{O}(\text{mw}(G)^3 \cdot n + m)$ and $\mathcal{O}(\text{sw}(G)^3 \cdot n + m)$, where $\text{mw}(G)$ is the *modular width* and $\text{sw}(G)$ the *split width*. Observe that these last decompositions can be computed in time $\mathcal{O}(n + m)$ [15].

1.4 Organisation.

§ 2 introduces the formal definitions of hyperbolicity, δ -value and distance- k domination. § 3 provides the technical lemma relating the δ -value of a quadruple to that of four nodes dominating it. § 4 provides technical lemmas for extending [12] as well as our new algorithm. Some experiments are presented in § 5 before concluding with § 6.

2 Definitions and notations

We use the graph terminology of [10, 27]. All graphs considered in this paper are finite, connected, unweighted and simple. The graph $G = (V, E)$ has $n = |V|$ vertices and $m = |E|$ edges. The open neighborhood $N_G(S)$ of a set $S \subseteq V$ consists of all vertices in $V \setminus S$ with at least one neighbor in S .

Given two vertices u and v , a uv -path of length $\ell \geq 0$ is a sequence of vertices ($u = v_0 v_1 \dots v_\ell = v$), such that $\{v_i, v_{i+1}\}$ is an edge for every i . In particular, a graph G is *connected* if there exists a uv -path for all pairs $u, v \in V$, and in such a case the *distance* $d_G(u, v)$ is defined as the minimum length of a uv -path in G . When G is clear from the context, we write d (resp. N) instead of d_G (resp. N_G). The *eccentricity* $\text{ecc}(u)$ of a vertex u is the maximum distance between u and any other vertex $v \in V$, i.e., $\text{ecc}(u) = \max_{v \in V} d(u, v)$. The maximum eccentricity is the *diameter* $\text{diam}(G)$ and the minimum eccentricity is the *radius* $\text{rad}(G)$.

Domination. Given an integral $k > 0$, we say that a node u *k-dominates* a node v when $d(u, v) \leq k$. We define a *k-dominating set* of G as a set $D \subseteq V$ of nodes such that any node $v \in V$ is *k-dominated* by some node $u \in D$. Given a *k-dominating set* D , we associate to any node $v \in V$ such a *k-dominating node* $D(v)$ in D . $D(v)$ is called the *associated dominator* of v . We denote

$D^{-1}(u)$ the set of vertices that are k -dominated by $u \in D$, i.e., $D^{-1}(u) = \{v \in V : D(v) = u\}$. For each node $u \in D$, we define its *domination radius* as $k_u = \max_{v \in D^{-1}(u)} d(u, v)$. When a node v is k -dominated by several nodes of D , the choice of its associated dominator $D(v)$ can be arbitrary. However, we will typically choose $D(v)$ as a closest node to v in D as a heuristic to obtain smaller values of k_u with $u \in D$.

Hyperbolicity. This notion has been introduced to measure how the shortest-path metric space (V, d_G) of a connected graph $G = (V, E)$ deviates from a tree metric when its vertices are mapped into the vertices of an edge-weighted tree. This additive stretch of the distances, denoted δ , is called the *hyperbolicity* of the graph and a graph is said to be δ -hyperbolic if it satisfies the 4-point condition below.

Definition 1 (4-points Condition, [34]). *Let G be a connected graph. For every quadruple u, v, x, y of vertices of G , we define $\delta(u, v, x, y)$ as half of the difference between the two largest sums among $S_1 = d(u, v) + d(x, y)$, $S_2 = d(u, x) + d(v, y)$, and $S_3 = d(u, y) + d(v, x)$.*

The hyperbolicity of G , denoted by $\delta(G)$, is equal to $\max_{u, v, x, y \in V(G)} \delta(u, v, x, y)$. Moreover, given a value $\bar{\delta}$, we say that G is $\bar{\delta}$ -hyperbolic whenever $\delta(G) \leq \bar{\delta}$.

Note that if G is a tree, or a clique, one can easily show $S_1 = S_2$ for all quadruples, implying $\delta(G) = 0$. If G is a cycle of order $n = 4p + \varepsilon$, with $p \geq 1$ and $0 \leq \varepsilon < 4$, then $\delta(G) = p - 1/2$ when $\varepsilon = 1$, and $\delta(G) = p$ otherwise. If G is a $n \times m$ grids, with $2 \leq n \leq m$, then we have $\delta(G) = n - 1$. Other definitions of hyperbolicity have been proposed [7, 25, 34] and differ only by a small constant factor.

3 Using dominating set for approximating hyperbolicity

In this section, we show how to obtain an additive $+4k$ approximation of hyperbolicity from a k -dominating set.

Lemma 1. *Given a k -dominating set D of G , a quadruple $u', v', x', y' \in V$ with respective associated dominators $u, v, x, y \in D$, we then have $\delta(u, v, x, y) - K_4 \leq \delta(u', v', x', y') \leq \delta(u, v, x, y) + K_4$ where $K_4 = k_u + k_v + k_x + k_y$.*

Proof. We first prove the second inequality. Assume without loss of generality that the first sum for quadruple u', v', x', y' is $S'_1 = d(u', v') + d(x', y')$, the second sum is $S'_2 = d(u', x') + d(v', y')$ and the third sum is $S'_3 = d(u', y') + d(v', x')$ where $S'_1 \geq S'_2 \geq S'_3$. We let S_1, S_2, S_3 denote the corresponding sums with u, v, x, y : $S_1 = d(u, v) + d(x, y)$, $S_2 = d(u, x) + d(v, y)$ and $S_3 = d(u, y) + d(v, x)$.

When $\delta(u', v', x', y') \leq K_4$, we obviously have $\delta(u', v', x', y') \leq \delta(u, v, x, y) + K_4$ as $\delta(u, v, x, y) \geq 0$.

Now consider the case $\delta(u', v', x', y') \geq K_4$. We can then show $S_1 \geq S_2$ and $S_1 \geq S_3$ as follows. Note that by triangle inequality, we have $d(u', v') \leq d(u', u) + d(u, v) + d(v, v') \leq d(u, v) + k_u + k_v$. Using similar inequalities for other pairs, we can deduce $S_1 - S_2 \geq S'_1 - K_4 - (S'_2 + K_4) \geq S'_1 - S'_2 - 2K_4$. By hyperbolicity definition, we have $S'_1 - S'_2 = 2\delta(u', v', x', y')$, and we get $S_1 - S_2 \geq 0$ as we assume $\delta(u', v', x', y') \geq K_4$. Similarly, we have $S_1 - S_3 \geq S'_1 - S'_3 - 2K_4 \geq S'_1 - S'_2 - 2K_4$ since we have $S'_3 \leq S'_2$. We thus also have $S_1 - S_3 \geq 0$.

We finally prove $\delta(u', v', x', y') \leq \delta(u, v, x, y) + K_4$. Having $S_1 \geq S_2$ and $S_1 \geq S_3$, the greatest sums for quadruple u, v, x, y are either S_1 and S_2 or S_1 and S_3 , implying $2\delta(u, v, x, y) = \min\{S_1 -$

$S_2, S_1 - S_3\}$. As both $S_1 - S_2$ and $S_1 - S_3$ are at least $S'_1 - S'_2 - 2K_4$ from above inequalities, we obtain $2\delta(u, v, x, y) \geq S'_1 - S'_2 - 2K_4 = 2\delta(u', v', x', y') - 2K_4$ as desired.

We now prove the first inequality. We assume without loss of generality that the first sum for quadruple u, v, x, y is $S_1 = d(u, v) + d(x, y)$, the second sum is $S_2 = d(u, x) + d(v, y)$ and the third sum is $S_3 = d(u, y) + d(v, x)$. When $\delta(u, v, x, y) \leq K_4$, we obviously have $\delta(u, v, x, y) - K_4 \leq \delta(u', v', x', y')$ as $\delta(u', v', x', y') \geq 0$. Otherwise, we can prove similarly as before that we have $S'_1 - S'_2 \geq S_1 - S_2 - 2K_4 \geq 0$ and $S'_1 - S'_3 \geq S_1 - S_3 - 2K_4 \geq S_1 - S_2 - 2K_4 \geq 0$. This ensures that S'_1 is the greatest sum for u', v', x', y' in that case too. We thus have $2\delta(u', v', x', y') = \min\{S'_1 - S'_2, S'_1 - S'_3\} \geq \min\{S_1 - S_2 - 2K_4, S_1 - S_3 - 2K_4\} = S_1 - S_2 - 2K_4 = 2\delta(u, v, x, y) - 2K_4$. $\square$

From Lemma 1, we can get an additive $+4k$ approximation of the hyperbolicity of the graph.

Corollary 1. *Given a k -dominating set D of G , let $\delta_L = \max_{u,v,x,y \in D} \delta(u, v, x, y)$. We have $\delta_L \leq \delta(G) \leq \delta_L + 4k$.*

4 Extending Borassi et al. to k -domination

In this section, we build upon the algorithm proposed in [12] to design a new algorithm for hyperbolicity that exploits k -domination.

4.1 Skippable, acceptable and valuable.

In this section, we show how to extend the pruning techniques proposed in [12, 18] to the use of k -dominating sets. As in [12, 18], we define for any quadruple u, v, x, y :

$$\tau(u, v, x, y) := \frac{1}{2} (d(u, v) + d(x, y) - \max\{d(x, u) + d(y, v), d(x, v) + d(y, u)\}).$$

As noted in [12, 18], we have $\delta(G) = \max_{u,v,x,y \in V} \tau(u, v, x, y)$. Indeed, if $d(u, v) + d(x, y)$ is the largest sum as defined in Definition 1, we have $\tau(u, v, x, y) = \delta(u, v, x, y)$, and otherwise $\tau(u, v, x, y) \leq 0$.

Given a k -dominating set D of a connected graph G and a quadruple $u, v, x, y \in D$, we now establish upper bounds on the values $\delta(u', v', x', y')$ and $\tau(u', v', x', y')$ for every quadruple u', v', x', y' with respective associated dominators u, v, x, y .

Lemma 2 ([18]). *For every quadruple u, v, x, y of vertices of a connected graph G , we have $\delta(u, v, x, y) \leq \min_{a,b \in \{u,v,x,y\}} d(a, b)$. Furthermore, if $S_1 = d(u, v) + d(x, y)$ is the largest of the sums defined in Definition 1 (which can be assumed w.l.o.g.), we have $\delta(u, v, x, y) \leq \frac{1}{2} \min\{d(u, v), d(x, y)\}$.*

Lemma 3. *Given a k -dominating set D of a connected graph G and a quadruple $u, v, x, y \in D$, then for each quadruple u', v', x', y' with respective associated dominators u, v, x, y , we have*

$$\delta(u', v', x', y') \leq \min_{a,b \in \{u,v,x,y\}} k_a + k_b + d(a, b) \leq 2k + \min_{a,b \in \{u,v,x,y\}} d(a, b)$$

Furthermore, assuming that $S_1 = d(u, v) + d(x, y)$ is the largest of the sums defined in Definition 1 (which can be assumed w.l.o.g.), for each quadruple u', v', x', y' with respective associated dominators u, v, x, y we have:

$$\begin{aligned}\delta(u', v', x', y') &\leq k_u + k_v + k_x + k_y + \frac{1}{2} \min \{d(u, v), d(x, y)\} \leq 4k + \frac{1}{2} \min \{d(u, v), d(x, y)\} \\ \text{and} \\ \tau(u', v', x', y') &\leq \frac{1}{2} \min \{k_u + k_v + d(u, v), k_x + k_y + d(x, y)\} \leq k + \frac{1}{2} \min \{d(u, v), d(x, y)\}.\end{aligned}$$

Proof. By Lemma 2, we have that $\delta(u', v', x', y') \leq \min_{a', b' \in \{u', v', x', y'\}} d(a', b')$. Since for any pair $\{a', b'\} \subset \{u', v', x', y'\}$ such that $D(a') = a$ and $D(b') = b$ with $\{a, b\} \subset \{u, v, x, y\}$, we have $d(a', b') \leq d(a', a) + d(a, b) + d(b, b') \leq k_a + k_b + d(a, b)$, by triangle inequality. We thus get $\delta(u', v', x', y') \leq \min_{a, b \in \{u, v, x, y\}} k_a + k_b + d(a, b)$.

Next, assuming that $S_1 = d(u, v) + d(x, y)$ is the largest of the sums defined in Definition 1, and that we have $\delta(u', v', x', y') > 0$, we get by Lemmas 1 and 2 that $\delta(u', v', x', y') \leq k_u + k_v + k_x + k_y + \delta(u, v, x, y) \leq k_u + k_v + k_x + k_y + \frac{1}{2} \min \{d(u, v), d(x, y)\}$.

Define the sums $S'_1 = d(u', v') + d(x', y')$, $S'_2 = d(u', x') + d(v', y')$ and $S'_3 = d(u', y') + d(v', x')$. We have $2\tau(u', v', x', y') = S'_1 - \max\{S'_2, S'_3\} \leq S'_1 - \frac{1}{2}(S'_2 + S'_3)$ and $S'_2 + S'_3 = d(u', x') + d(v', y') + d(u', y') + d(v', x') = (d(u', x') + d(u', y')) + (d(v', x') + d(v', y'))$. By triangle inequality, we get $S'_2 + S'_3 \geq 2d(x', y')$ and so $2\tau(u', v', x', y') \leq d(u', v') \leq k_u + k_v + d(u, v)$. Similarly, we have by triangle inequality that $S'_2 + S'_3 = (d(u', x') + d(v', x')) + (d(u', y') + d(v', y')) \geq 2d(u', v')$, and so $2\tau(u', v', x', y') \leq d(x', y') \leq k_x + k_y + d(x, y)$. $\square$

Interestingly, the upper bound on $\tau(u', v', x', y')$ established in Lemma 3 is significantly stronger than the upper bound on $\delta(u', v', x', y')$. This will help reducing the size of the search space, and in particular to decide when to stop the exploration of quadruples and to return the result.

By [12, Lemma 8], we know that if $x, y, u \in V$ are such that $\text{ecc}(u) + d(x, y) - d(x, u) - d(y, u) \leq 4\delta_L + 2$, where δ_L is a lower bound on the hyperbolicity of G , then for each $v \in V$, we have $\tau(u, v, x, y) \leq \delta_L$. Hence, for each such triple (x, y, u) , we can skip the exploration of the n quadruples (u, v, x, y) for $v \in V$. We now adapt this result to the case $x, y, u \in D$.

Lemma 4. *Given a k -dominating set D of a graph G , a lower bound δ_L on the hyperbolicity of G and the vertices $x, y, u \in D$ satisfying $2\text{ecc}(u) + d(x, y) - d(x, u) - d(y, u) + K_8 \leq 4\delta_L + 1$ where $K_8 = 4k_u + 2k_x + 2k_y \leq 8k$, then for each quadruple u', v', x', y' with respective associated dominators u, v, x, y , we have $\tau(u', v', x', y') \leq \delta_L$.*

Proof. Given $x, y, u \in D$ satisfying $2\text{ecc}(u) + d(x, y) - d(x, u) - d(y, u) + K_8 \leq 4\delta_L + 1$, we assume for the sake of contradiction that there exists $x' \in D^{-1}(x)$, $y' \in D^{-1}(y)$, $u' \in D^{-1}(u)$ and $v' \in V$ such that $\delta_L < \tau(u', v', x', y')$.

Then, $2\delta_L + 1 \leq 2\tau(u', v', x', y') = d(u', v') + d(x', y') - \max\{d(u', x') + d(v', y'), d(u', y') + d(v', x')\} \leq d(u', v') + d(x', y') - \frac{1}{2}(d(u', x') + d(v', y') + d(u', y') + d(v', x')) \leq d(x', y') + \text{ecc}(u') - \frac{1}{2}(d(u', x') + d(u', y') + d(x', y')) \leq \text{ecc}(u') - \frac{1}{2}(d(u', x') + d(u', y') - d(x', y'))$.

Now, observe that we have $\text{ecc}(u') \leq \text{ecc}(u) + k_u$ by triangle inequality. Similarly, we have $d(u', x') + d(u', y') - d(x', y') \geq d(u, x) - k_u - k_x + d(u, y) - k_u - k_y - d(x, y) - k_x - k_y$. Hence, we get $2\delta_L + 1 \leq \text{ecc}(u) + k_u - \frac{1}{2}(d(u, x) + d(u, y) - d(x, y) - 2(k_u + k_x + k_y))$, and so $4\delta_L + 2 \leq 2\text{ecc}(u) + d(x, y) - d(u, x) - d(u, y) + K_8$ which contradicts the hypothesis. $\square$

Similarly, by [12, Lemma 9], we know that if $x, y, u \in V$ are such that $\text{ecc}(u) + d(x, y) - 3\delta_L - \frac{3}{2} \leq \max\{d(u, x), d(u, y)\}$, then for each $v \in V$, we have $\tau(u, v, x, y) \leq \delta_L$. Hence, for each such triple we can skip the exploration of n quadruples. We now adapt this result to the case $x, y, u \in D$.

Lemma 5. *Given a k -dominating set D of a graph G , a lower bound δ_L on the hyperbolicity of G and vertices $x, y, u \in D$ satisfying $\text{ecc}(u) + d(x, y) - 3\delta_L - 1 + K_4 \leq \max\{d(u, x) - k_x, d(u, y) - k_y\}$*

where $K_4 = 2k_u + k_x + k_y \leq 4k$, then for each quadruple u', v', x', y' with respective associated dominators u, v, x, y we have $\tau(u', v', x', y') \leq \delta_L$.

Proof. Given $x, y, u \in D$ satisfying $\text{ecc}(u) + d(x, y) - 3\delta_L - 1 + K_4 \leq \max\{d(u, x) - k_x, d(u, y) - k_y\}$, we assume for the sake of contradiction that there exists $x' \in D^{-1}(x)$, $y' \in D^{-1}(y)$, $u' \in D^{-1}(u)$ and $v' \in V$ such that $\delta_L < \tau(u', v', x', y')$ or equivalently $2\delta_L + 1 \leq 2\tau(u', v', x', y')$.

By Lemma 2, we have $d(u', y') > \delta_L$, that is $d(u', y') \geq \delta_L + \frac{1}{2}$. Consequently, we get $2\delta_L + 1 \leq 2\tau(u', v', x', y') = d(u', v') + d(x', y') - \max\{d(u', x') + d(v', y'), d(u', y') + d(v', x')\} \leq d(u', v') + d(x', y') - d(u', x') - d(u', y') \leq \text{ecc}(u') + d(x', y') - d(u', x') - \delta_L - \frac{1}{2}$.

Since $\text{ecc}(u') \leq \text{ecc}(u) + k_u$, $d(x', y') \leq d(x, y) + k_x + k_y$ and $d(u', x') \geq d(u, x) - k_u - k_x$, we obtain $2\delta_L + 1 \leq \text{ecc}(u) + k_u + d(x, y) + k_x + k_y - d(u, x) + k_u + k_x - \delta_L - \frac{1}{2}$, that is $d(u, x) - k_x \leq \text{ecc}(u) + d(x, y) - 3\delta_L - \frac{3}{2} + K_4$. By exchanging the role of x' and y' , we similarly get $d(u, y) - k_y \leq \text{ecc}(u) + d(x, y) - 3\delta_L - \frac{3}{2} + K_4$. These two inequalities contradict the hypothesis. $\square$

We use Lemmas 3, 4 and 5 to identify *acceptable* and *skippable* vertices. More precisely,

Definition 2. Given a k -dominating set D of a graph G , a lower bound δ_L on the hyperbolicity of G and a pair $x, y \in D$ of vertices, we say that a vertex $u \in D$ is *acceptable* if it satisfies all the following conditions. Otherwise, it is *skippable*.

1. $d(u, x) + k_u + k_x > \delta_L$ and $d(u, y) + k_u + k_y > \delta_L$ (Lemma 3);
2. $2\text{ecc}(u) + d(x, y) - d(u, x) - d(u, y) + 4k_u + 2k_x + 2k_y > 4\delta_L + 2$ (Lemma 4);
3. $\text{ecc}(u) + d(x, y) - 3\delta_L - \frac{3}{2} + 2k_u + k_x + k_y > \max\{d(u, x) - k_x, d(u, y) - k_y\}$ (Lemma 5).

Thanks to Definition 2, we can avoid exploring all quadruples for which either u or v is skippable, thus saving a significant amount of computations. To further reduce the search space, [12] introduce the notion of c -valuable vertices (or simply valuable vertices), where $c \in V$ is an arbitrarily chosen vertex, for instance a vertex with small centrality or eccentricity.

Lemma 6. Given a k -dominating set D of a graph G , a pair $x, y \in D$ of vertices and any fixed node $c \in V$, define $f_c(z) = \frac{1}{2}(d(x, y) - d(x, z) - d(y, z)) + d(z, c) + 2k_z + k_x + k_y$ for any $z \in D$. Then, for each pair $u, v \in D$ and each quadruple u', v', x', y' with respective associated dominators u, v, x, y , we have $2\tau(u', v', x', y') \leq f_c(u) + f_c(v)$.

Proof. We have $2\tau(u', v', x', y') = d(u', v') + d(x', y') - \max\{d(u', x') + d(v', y'), d(u', y') + d(v', x')\} \leq d(u', v') + d(x', y') - \frac{1}{2}(d(u', x') + d(v', y') + d(u', y') + d(v', x'))$. Furthermore, we have $d(u', v') \leq d(u', c) + d(v', c)$, $d(u', c) \leq d(u, c) + k_u$ and $d(v', c) \leq d(v, c) + k_v$. Hence, we get $2\tau(u', v', x', y') \leq d(x, y) + d(u, c) + d(v, c) - \frac{1}{2}(d(u, x) + d(u, y) + d(v, x) + d(v, y)) + 2k_u + 2k_v + 2k_x + 2k_y = f_c(u) + f_c(v)$. $\square$

As a consequence of Lemma 6, if $2\tau(u, v, x, y) + 2(k_u + k_v + k_x + k_y) > 2\delta_L$, then either $f_c(u) > \delta_L$ or $f_c(v) > \delta_L$.

Definition 3. Given a k -dominating set D of a graph G , a lower bound δ_L on the hyperbolicity of G and a pair $x, y \in D$ of vertices, we say that an acceptable vertex $u \in D$ is c -valuable (or simply valuable) if $f_c(u) > \delta_L$.

Summarizing, if $\tau(u, v, x, y) > \delta_L$, then u and v are both acceptable and at least one of them is furthermore valuable.

Observe that, for a given pair $(x, y) \in D^2$, deciding if a vertex $u \in D$ is skippable, acceptable or valuable requires a constant number of operations, assuming that distances and eccentricities are known. We will discuss this aspect in § 4.3.

4.2 Algorithm.

We can now present an exact algorithm (Algorithm 1) that exploits the notion of k -domination to prune the search space and significantly reduce the memory usage compared to the algorithms proposed in [12, 18, 21]. Algorithm 1 takes as input a connected graph G and a sequence $k_i, k_{i-1}, \dots, k_0$ of domination distances, with $i \geq 0$ and $k_i = k > k_{i-1} > \dots > k_0 = 0$.

The general principle of the algorithm is to first scan the quadruples of D_i where D_i is a k_i -dominating set. Then, thanks to Lemma 1, we know that if a quadruple $u, v, x, y \in D_i$ is such that $\delta(u, v, x, y) + 4k > \delta_L$, where δ_L is the current lower bound on $\delta(G)$, it may dominate a quadruple u', v', x', y' such that $\delta(u', v', x', y') > \delta_L$. In this case, instead of exploring directly all the quadruples of V that are dominated by u, v, x, y , we scan the quadruples made of the nodes of a k_{i-1} -dominating set D_{i-1} that are dominated by u, v, x, y . We proceed similarly for smaller and smaller domination radii k_j with $0 < j < i$ as long as the condition of Lemma 1 is satisfied for k_j . When a quadruple does not satisfies the conditions of Lemma 1, we prune the exploration of all dominated quadruples. We also use Lemma 3 and Definitions 2 and 3 to prune more quadruples, generalizing in some extend the approach of [12].

Let us now present the algorithm in details.

The first step of the algorithm is the construction of a hierarchical dominating set (line 1 of Algorithm 1). This is done as follows. We compute a k_i -dominating set D_i for G , using a greedy algorithm (see § 4.3). We associate to each vertex v the closest vertex $D_i(v)$ that dominates it and let $V_{i,u} = D_i^{-1}(u)$ denote the set of vertices associated to u for each $u \in D_i$. Note that $\{V_{i,u} : u \in D_i\}$ forms a partition of V . We let $k_{i,u} = \max_{v \in V_{i,u}} d(u, v)$ denote the domination radius of $u \in D_i$. Obviously, we have $k_{i,u} \leq k_i$. We proceed similarly for each $j < i$ as follows. Given a subset $V' \subseteq V$ we define a k -dominating set of V' as a set $D \subseteq V'$ such that each vertex $v \in V'$ is dominated by some vertex $u \in D$ such that $d(u, v) \leq k$. Given a partition of V into sets $V_{j+1,u}$ for $u \in D_{j+1}$, we compute a k_j -dominating set $D_{j,u}$ of $V_{j+1,u}$ for each $u \in D_{j+1}$. Each vertex $v' \in V_{j+1,u}$ is associated to its closest dominator $D_{j,u}(v')$ in $D_{j,u}$. We then let $V_{j,u'} = D_{j,u}^{-1}(u')$ denote the set of vertices associated to u' for each $u' \in D_{j,u}$ and let $k_{j,u'} = \max_{v' \in V_{j,u'}} d(u', v')$ denote its domination radius. We then obtain a new partition of V as $\{V_{j,u'} : u' \in D_j\}$ where $D_j = \cup_{u \in D_{j+1}} D_{j,u}$. As $k_0 = 0$, this process ends with $D_{0,u} = V_{1,u} = \{u\}$ for each $u \in D_1$ and $D_0 = V$.

Next, Algorithm 1 computes the distances between all pairs of vertices of the k -dominating set D_i and stores these distances in matrix d (line 2). The construction of d can be done in time $\mathcal{O}(|D_i|(n + m))$ using BFS and uses spaces $\mathcal{O}(|D_i|^2)$. We assume here that $|D_i|$ is small enough to ensure that matrix d can fit into memory.

Algorithm 1 uses the notion of mates, introduced in [12]. Roughly, v is a mate of u if $d(u, v) \geq d(x, y)$ and the pair (u, v) has been considered before the pair (x, y) that is currently considered. Hence, when all quadruples involving a pair (x, y) have been considered, x becomes a mate of y and y becomes a mate of x . This notion is used line 12 of the algorithm to reduce the number of pairs (u, v) to consider. The maintenance of the mates requires overall time and space in $\mathcal{O}(|D_i|^2)$.

The main part of the algorithm is lines 5 to 19 which uses the best value found so far (δ_L) and the bounds of Lemma 3 to stop computations as soon as possible. More precisely, the algorithm considers all pairs (x, y) of vertices in D_i sorted by non-increasing distance. This ordering is a consequence of Lemma 3. Indeed, if a pair (x, y) is such that $d(x, y) + 2k_i \leq 2\delta_L$, then we also have $d(x', y') + 2k_i \leq 2\delta_L$ for all pairs (x', y') such that $d(x', y') \leq d(x, y)$. Hence, we can stop computation and return the result as soon as the condition is satisfied (line 6). Furthermore, Lemma 3 also ensures that for a pair (x, y) such that $d(x, y) + k_{i,x} + k_{i,y} \leq 2\delta_L$, we have $\tau(u, v, x, y) \leq \delta_L$ for all pairs (u, v) , and so we can skip computations involving this pair (x, y) (line 8).

Then, the algorithm determines, through a call to method `computeAccVal`(x, y, D_i), the acceptable and valuable vertices of D_i for pair (x, y) according Definitions 2 and 3 (line 10). It simply consists in a linear scan of D_i , as explained in § 4.3. The algorithm then considers all quadruples (u, v, x, y) made of the pair (x, y) , a valuable vertex u and an acceptable vertex v that is also a mate of u . The algorithm first checks whether this quadruple enables to improve the lower bound δ_L (line 14). Then, it checks using Lemma 1 whether this quadruple may dominates another quadruple that could help improving the bound (lines 16-17). If this is the case, it calls method `explore` (Algorithm 2) that we describe next.

Algorithm 1 Compute hyperbolicity using k -domination.

Require: $G = (V, E)$, a connected graph

Require: $k_i = k > k_{i-1} > \dots > k_0 = 0$, a sequence of domination distances

```

1:  $\mathcal{H} = \{(D_j, (D_{j-1,u})_{u \in D_j}, k_j), 0 \leq j \leq i\} \leftarrow \text{hierarchicalDominatingSet}(G, (k_i, k_{i-1}, \dots, 0))$ 
2: Let  $d$  be the distance matrix between the vertices in  $D_i$ 
3:  $\delta_L \leftarrow 0$ 
4:  $\text{mates}[v] \leftarrow \emptyset$  for each  $v \in D_i$ 
5: for all  $(x, y) \in D_i^2$  sorted by non increasing distances do
6:   if  $d(x, y) + 2k_i \leq 2\delta_L$  then
7:     return  $\delta_L$ 
8:   if  $d(x, y) + k_{i,x} + k_{i,y} \leq 2\delta_L$  then
9:     continue
10:   $(\text{acceptable}, \text{valuable}) \leftarrow \text{computeAccVal}(x, y, D_i)$ 
11:  for  $u \in \text{valuable}$  do
12:    for  $v \in \text{mates}[u]$  do
13:      if  $v \in \text{acceptable}$  then
14:         $\delta_L \leftarrow \max\{\delta_L, \tau(u, v, x, y)\}$ 
15:         $\delta_{UB} \leftarrow \tau(u, v, x, y) + \sum_{z \in \{u, v, x, y\}} k_{i,z}$ 
16:        if  $\delta_{UB} > \delta_L$  then
17:           $\delta_L \leftarrow \max\{\delta_L, \text{explore}(u, v, x, y, \mathcal{H}, i - 1)\}$ 
18:      add  $y$  to  $\text{mates}[x]$ 
19:      add  $x$  to  $\text{mates}[y]$ 
20: return  $\delta_L$ 
```

Given a quadruple (u, v, x, y) of vertices in D_{j+1} , method `explore` (Algorithm 2) considers all pairs (x', y') such that $x' \in D_{j,x}$ and $y' \in D_{j,y}$. Recall that $D_{j,x}$ is the set of vertices of a distance k_{j-1} dominating set of the set of vertices $V_{j,x}$ that are dominated by x in a distance k_j dominating set. Thanks to Lemma 3, the algorithm do not consider the pairs (x', y') that can not

Algorithm 2 Method $\text{explore}(u, v, x, y, \mathcal{H}, j)$.

Require: $G = (V, E)$, a connected graph

Require: j , the current index for domination

Require: $u, v, x, y \in D_{j+1}$, a quadruple of G

Require: $\mathcal{H} = \{(D_{j'}, (D_{j'-1, u})_{u \in D_{j'}}, k_{j'})\}$, $0 \leq j' \leq i\}$, hierarchical dominating set of G

```
1: if  $j = -1$  then
2:   return  $\tau(u, v, x, y)$ 
3: for  $x' \in D_{j, x}$  do
4:   for  $y' \in D_{j, y}$  do
5:     if  $d(x', y') + k_{j-1, x'} + k_{j-1, y'} \leq 2\delta_L$  then
6:       continue
7:      $(\text{acceptable}(u), \text{valuable}(u)) \leftarrow \text{computeAccVal}(x', y', D_{j, u})$ 
8:      $(\text{acceptable}(v), \text{valuable}(v)) \leftarrow \text{computeAccVal}(x', y', D_{j, v})$ 
9:     for  $(a, b) \in \{(u, v), (v, u)\}$  do
10:      for  $u' \in \text{valuable}(a)$  do
11:        for  $v' \in \text{acceptable}(b)$  do
12:           $\delta_L \leftarrow \max\{\delta_L, \tau(u', v', x', y')\}$ 
13:           $k_4 \leftarrow \sum_{z' \in \{u', v', x', y'\}} k_{j-1, z'}$ 
14:          if  $\tau(u', v', x', y') + k_4 > \delta_L$  then
15:             $\delta_L \leftarrow \max\{\delta_L, \text{explore}(u', v', x', y', \mathcal{H}, j-1)\}$ 
16: return  $\delta_L$ 
```

help improving the lower bound (line 5).

Next, Algorithm 2 determines the subsets of vertices of $D_{j, u}$ and $D_{j, v}$ that are acceptable or valuable for (x', y') according Definitions 2 and 3 (lines 7 and 8). Then, it considers the quadruples (u', v', x', y') such that either u' is a valuable vertex of $D_{j, u}$ and v' an acceptable vertex of $D_{j, v}$, or u' is an acceptable vertex of $D_{j, u}$ and v' a valuable vertex of $D_{j, v}$. For each quadruple (u', v', x', y') , it checks using Lemma 1 whether this quadruple may dominate another quadruple that could help improving the bound (line 14). If this is the case, it performs a recursive call to **explore** to consider the dominated quadruples in a distance k_{j-1} dominating set (line 15).

Observation 1. When the domination distance is zero ($k = 0$), the test on line 16 of Algorithm 1 will always be false, and so the algorithm is almost the one proposed in [12]. Indeed, the only difference is that the algorithm of [12] considers the list of *far-apart pairs* instead of the list of all pairs, which helps further pruning the search space. Roughly, u is far from v if $d(u, v) \geq d(u, w)$ for all $w \in N(v)$, and the pair (u, v) is far-apart if u is far from v and v is far from u . It has been proved in [43] that any graph contains two far-apart pairs (u, v) and (x, y) satisfying $\delta(u, v, x, y) = \delta(G)$. An interesting open question is therefore how to use the notion of far-apart pairs in Algorithms 1 and 2 to further prune the search space.

Observation 2. Since $\max_{a', b' \in V_{i, u}} d(a', b') \leq 2k$ for all $u \in D_i$, we know by Lemma 3 that, when $\delta_L \geq 2k$, Algorithm 1 has to consider only quadruples of D_i composed of four distinct nodes. However, when $\delta_L < 2k$, it is necessary to extend the list of pairs (x, y) to consider line 5 of Algorithm 1 by the list of pairs $\{(x, x) : x \in D_i\}$ in order to also consider quadruples of D_i with two or more identical nodes.

4.3 Implementation notes.

We now discuss how to efficiently implement Algorithms 1 and 2.

Let us start with the method used to create the dominating sets. We use a greedy approach where we order the nodes decreasingly according to their degrees and then, in this order, add a node to the dominating set if and only if it is not dominated yet for the considered radius. During this process, we ensure that a dominated vertex is attached to the closest vertex that dominates it, and we record for each vertex of the dominating set its domination radius (i.e., largest distance to a vertex it dominates). While we tried different strategies to construct dominating sets (e.g., order the vertices by increasing/decreasing degree/eccentricity, ensure that the vertices dominated by a vertex induce a connected subgraph, etc.), we could not observe large performance changes on the hyperbolicity computation. In particular, we do not believe that this part of the algorithm is worthwhile optimizing further.

Let us now consider the data used in Algorithm 1, and in particular for the classification of vertices as skippable, acceptable or valuable. This step of the algorithm requires the knowledge of the eccentricities of the vertices (Definition 2), of the distances from a central vertex (Lemma 6), of the effective domination distance of each considered vertex (given by the hierarchical dominating set), and of the distances between the considered vertices. Firstly, it is possible to efficiently compute the eccentricities of the vertices using the algorithms proposed in [44, 29, 40] that perform a smart management of the upper and lower bounds on the eccentricity of each vertex in order to avoid computing distances from each vertex. In our implementation, we have chosen to use the algorithm proposed in [29]. Secondly, as noticed in [12], a good choice for the central vertex is a vertex with small eccentricity or farness (sum of the distances from a vertex to all other vertices). Since we have already computed the eccentricities of the vertices, we choose a vertex with minimum eccentricity as central vertex, and we store distances from that central vertex (computed using BFS). Now, observe that Algorithm 1 uses the distance matrix between all the vertices in D_i (line 2). This matrix of distances is obtained solving a BFS from each vertex in D_i . Consequently, each required data for the classification of the vertices (method `computeAccVal`, line 10 of Algorithm 1) can be accessed in constant time and so this classification requires a constant number of operations per vertex.

Let us continue with the data used in Algorithm 2. Apart from the eccentricities of the vertices, the distances from the central vertex and the effective domination distances that have already been computed for Algorithm 1, we need the distances between the vertices dominated by the quadruple u, v, x, y . Similarly, method `computeAccVal` linearly scans all vertices u' in a dominating set and requires distances between u' and a pair x', y' of nodes to determine whether u' is acceptable and/or valuable. In order to get all these distances, we use a hub labeling (also called 2-hop labeling) of the graph [33]. Roughly, this data structure assigns to each vertex u of the graph a *label*, that is list of vertices with shortest path distance from u to each of these vertices. The construction of the lists ensures that for each pair of vertices $u, v \in V$, the intersection of the lists of u and v contains a vertex that is on a shortest path between u and v . Then, to get the distance between u and v , it suffices to search for the vertex minimizing the distance in the intersection of the lists. In practice, the labels are shorts (i.e., of poly-logarithmic length in the size of the graph) and so distance queries can be answered efficiently for graphs with millions of nodes. Furthermore, this data structure uses little space. Practical algorithms for building such data structure have time complexity in $\mathcal{O}(nm)$, and we choose to use the algorithm proposed in [3] and [26] that combines several methods to significantly speedup the construction.

Now, since the distance between say $u' \in D_j^{-1}(u)$ and $v' \in D_j^{-1}(v)$ is used several times in Algorithm 2, we first store these distances in a matrix, thus enabling to access these distances in constant time when needed, and so to classify a vertex in constant time. More precisely, at the beginning of Algorithm 2, we build for each pair (u, v) , (x, u) , (y, u) , (x, v) , (y, v) and (x, y) a rectangular matrix storing the required distances (e.g., all distances between $u' \in D_j^{-1}(u)$ and $v' \in D_j^{-1}(v)$). We thus query the hub labeling data structure only once per distance. In order to further reduce the number of queries to the hub labeling data structure, we use a cache of distance matrices. Indeed, two consecutive calls to Algorithm 2 for domination index j might involve a same pair (x, y) , and possibly a same triple u, x, y . This cache is implemented as a bounded size doubly linked list of matrices along with a mapping from pairs of vertices to elements of this list. This way, if the matrix needed for a pair of vertices is in the cache, we move it to the end of the doubly linked list and return it. Otherwise, we build a new matrix, insert it at the end of the list and record it in the mapping. When the maximum size of the cache is reached, we remove the first element of the list and the corresponding entry from the mapping, and then create a new matrix and insert it at the end of the list. Since we use only 7 matrices at the same time (the 6 matrices listed above and the matrix for the pair (v, u)), this method is safe if the size of the cache is at least 7. In practice the matrices are small and so we can store millions of them. However, if we observe that a matrix is too large to be stored (e.g., if its side is more than 50 000), we avoid the construction of the matrices and directly query the hub labeling data structure.

Finally, in order to further reduce the overall running time of the algorithm, we perform two pass of lines 4-19 of Algorithm 1. During the first pass, we omit lines 16-17, and so the calls to **explore**, and we obtain a good lower bound δ_L such that $\delta_L \leq \delta(G) \leq \delta_L + 4k$. Then, during the second pass of lines 4-19 of Algorithm 1, we call **explore** on the quadruples that can help improving the lower bound δ_L . Since we start the second pass with a better lower bound δ_L , we prune many calls to **explore**. In addition, we also record for each pair (x, y) the best found value $\tau(u, v, x, y) + k_{i,u} + k_{i,v} + k_{i,x} + k_{i,y}$ over all considered pairs (u, v) of valuable and acceptable vertices during the first pass. This value is in fact an upper bound on the value $\tau(u', v', x', y')$ of any quadruple u', v', x', y' having respective associated dominators u, v, x, y such that $d(u, v) \geq d(x, y)$. Hence, during the second pass, we can avoid considering the pairs (x, y) for which this value is less than the current lower bound δ_L , and so that cannot lead to any call to **explore**.

4.4 Correctness of the algorithm.

The validity of the algorithms mainly follows from Lemmas 1 and 3, and Definitions 2 and 3. More precisely, the validity of the pruning at line 16 in Algorithm 1 comes from Lemma 1. The general proof of correctness is similar to that of [12] using Definitions 2 and 3, and relying on Lemmas 3, 4 and 5. Observe that a quadruple u, v, x, y may be considered several times. Indeed, it dominates itself, and so can be considered with a smaller radii during the recursive calls to method **explore**. However, a quadruple is considered at most $i + 1$ times (i.e., the number of radii) and so the worst case time complexity is $\mathcal{O}(n^4)$ when i is a small constant. The space usage is dominated by the hub labeling and the distance matrix for nodes in D_i . Although this could be $\mathcal{O}(n^2)$ in the worst case, it is close to linear when $|D_i| = \mathcal{O}(\sqrt{n})$ and the average hubset size is small.

Table 1: Selected graph parameters of all the graphs that we use in our experiments. Note that for all graphs, we extracted the largest biconnected component and restrict to this subgraph in our experiments.

Graph	#nodes	#edges	eccentricity			degree		
			radius	mean	diameter	min	mean	max
NotreDame	134 958	833 732	18	20.99	36	2	12.36	10 721
web-Stanford	181 906	1 676 077	46	48.81	92	2	18.43	35 488
web-BerkStan	489 296	5 939 242	46	49.19	92	2	24.27	80 733
t.CAL	1 267 004	1 671 989	1 149	1 656.89	2 298	2	2.64	7
t.FLA	691 175	941 893	890	1 378.52	1 780	2	2.73	8
roadNet-PA	863 105	1 313 732	401	689.81	793	2	3.04	9
buddha	543 652	1 631 574	244	360.44	487	3	6.0	17
froz	749 520	2 895 228	812	1 130.38	1 451	2	7.73	8
grid300-10	90 211	162 152	300	450.50	600	2	3.6	4
grid500-10	250 041	449 831	500	750.49	1 000	2	3.6	4
z-alue7065	34 040	54 835	213	319.43	426	2	3.22	4

5 Experimental Evaluation

We now turn to experimentally evaluating the algorithm presented in § 4. To this end, we first explain the setup of the experiments to then compare our approach to the previous state of the art. As our algorithm depends on input parameters, we subsequently evaluate how the choice of parameters affects the running time.

5.1 Setup.

For our experimental evaluation, we used web graphs (NotreDame, web-BerkStan, web-Stanford), road networks (t.CAL, t.FLA, roadNet-PA), a 3D triangular mesh (buddha), a grid-like graphs from VLSI applications (z-alue7065) and a graph from a computer game (froz). The data is available from snap.stanford.edu, www.dis.uniroma1.it/challenge9, graphics.stanford.edu, steinlib.zib.de, and movingai.com. We also use synthetic graphs (grid300-10, grid500-10) which are square grids with respective sides 301 and 501 where 10% of the edges have been randomly deleted. Each graph is taken as an undirected unweighted graph and we consider only its largest biconnected component. See Table 1 for the characteristics of these graphs. The chosen graphs have a large number of nodes compared to the graph sizes that were feasible for previous implementations that compute the graph hyperbolicity. Furthermore, our approach relies on the fact that pruning of quadruples is actually possible. Thus, the graphs that we consider do not exhibit a very low hyperbolicity — the lowest is 8 (NotreDame).

We implemented our algorithms in C++ and our code can be found at [22]. For all except two experiments, we used a computer equipped with Intel Xeon Gold 6240 CPUs operating at 2.6GHz and 192GB RAM. The other two experiments (the very long experiments in Table 2) were run on a computer equipped with an Intel Core i9-10900K CPU operating at 3.7GHz and 64GB RAM. Both machines run Linux. Furthermore, all computations were conducted using a single thread.

For the hub labeling, we use the algorithm proposed in [3], except for road networks for which

Table 2: Experiments on different graphs using algorithm [21] with time limit of 60h (which are 216 000 seconds) and memory limit of 192GB and the algorithm proposed in this paper (including the computation of the dominating sets and the hub labeling). A skull (A) in the time/memory column means the time/memory limit was reached before terminating.

Graph	Algorithm [21]			This paper				
	time (s)	memory	hyperb.	time (s)	memory	hyperb.	Parameters dom dist ratio	
NotreDame	4 514	53.02 GB	8.0	249	3.27 GB	8.0	2	2
web-Stanford	8 249	23.28 GB	23.0	15	3.22 GB	23.0	8	2
web-BerkStan	65 134	76.93 GB	23.0	59	3.33 GB	23.0	4	4
t.CAL	29 358	A	[379.0, 1025.5]	119 055	22.00 GB	506.5	50	1.5
t.FLA	A	143.3 GB	[81.0, 818.0]	1 199 907	18.04 GB	229.5	25	1.5
roadNet-PA	A	148.2 GB	[109.0, 370.5]	1 357 512	23.32 GB	170.5	20	1.5
buddha	A	88.35 GB	[93.0, 211.5]	134 421	52.84 GB	112.0	8	1.5
froz	A	106.4 GB	[387.5, 599.0]	16 011	11.74 GB	401.5	27	1.5
grid300-10	10	1.08 GB	280.0	23	5.28 GB	280.0	10	1.5
grid500-10	95	2.99 GB	463.0	98	6.14 GB	463.0	10	2
z-alue7065	33	431.18 MB	138.0	1 927	3.48 GB	138.0	2	2

we use [26] as we observed that it produces shorter labels in this case. Reported running times have been measured using `/usr/bin/time -v` and include all steps of the program, from reading data from file to returning the result. Memory limits have been set to 192GB using `ulimit`.

5.2 Parameters.

In Algorithm 2, we assume that a hierarchical dominating set is given as input. Especially, we require the domination distances to be given as parameters. Note that the number of possible sequences of domination distances is huge and it is impractical as parameter of an algorithm. We thus reduce the potential sequences by just requiring two parameters: the largest domination distance and the ratio by which it is reduced in each round. More formally, given a maximal domination distance k and a ratio r , the sequence of domination distances $k_i, \dots, k_0$ of the dominating sets then is defined as

$$k_i = k, \quad k_{i-1} = \left\lfloor \frac{k_i}{r} \right\rfloor, \quad \dots, \quad k_0 = \left\lfloor \frac{k_1}{r} \right\rfloor = 0.$$

Note that the value of i is implicitly given by the number of steps until this process reaches zero.

5.3 Comparison with Previous Work.

To evaluate the improvement of our algorithm over previous work, we compare to the implementation of [21], which was shown to outperform — considering the trade-off between running time and memory consumption — the algorithm of [12]. See [21, Figures 1 and 2] for a comparison of the running time and memory consumption of [12] and [21]. In particular, the memory consumption of [12] is prohibitive for all the graph sizes considered in this work except for the graph **z-alue7065**, which we mainly use to conduct experiments with different parameter choices.

To compare our new approach with [21], we let both algorithms run on the same graphs. The parameters of our approach are manually chosen and we give more details about this choice in § 5.4.

We show the results of our experiments in Table 2. The most notable points of these results are:

- For the first three graphs in the experiments (**NotreDame**, **web-Stanford**, **web-BerkStan**), we reduce the memory consumption compared to [21] by factors of 16.2, 7.2, and 23.1. The running time is reduced by factors of 18.1, 549.9, and 1104, i.e., up to 3 orders of magnitude.
- Computing the hyperbolicity of the second and third set of graphs (**t.CAL**, **t.FLA-w**, **roadNet-PA**, **buddha**, **froz**) was not feasible using previous algorithms but it can be computed using our approach. Especially, we are able to compute the hyperbolicity of a real-world graph with more than a million nodes for the first time. We want to highlight that the memory consumption of our algorithm is below 25 GB for all graphs except **buddha**. Even in the cases where the algorithm of [21] hits the time limit of 60 hours and not the memory limit, increasing the time limit most probably does not make these graphs attainable, as the lower and upper bounds are still very far from matching.
- The class of graphs where our algorithm fails to outperform previous work is grid-like graphs (**grid300-10**, **grid500-10**, **z-blue7065**). This is probably due to the notion of far-apart pairs introduced in [12] and also used in [21], which we do not use in our approach. Computing the hyperbolicity iterating over far-apart pairs is very fast on grid-like graphs as they contain only very few far-apart pairs — a perfect grid actually just contains two far-apart pairs, which also form the quadruple that implies its hyperbolicity.

Figure 1: The running times for the **z-blue7065** graph for different parameters.

5.4 Parameter Choice.

The main difficulty for applying our algorithm to new graphs is to determine the parameters that lead to the best performance. Note that these parameters only influence the running time and memory consumption of our algorithm but not the output. As described in § 5.2, we use the two parameters *maximal domination distance* and *ratio*, which in turn determine the dominating set distances that are used. To understand the effect on the running time, we ran our algorithm with different parameters on the graph **z-blue7065**. More precisely, we did an experiment for all

combinations of the maximal domination distance being from the set $\{2, 3, \dots, 8\}$ and the ratio being from the set $\{1.1, 1.5, 2, 2.5, 3\}$. The running times of these experiments are shown in Figure 1. The results suggest that too large or too small values for the ratio are detrimental to a fast running time. Furthermore, on this graph a small maximal domination distance is better. As can be seen by the parameter choices in Table 2, the moderate value for the ratio is also a good choice for other graphs, while the maximal domination distance can vary greatly depending on the graph. We leave it to future research to investigate how specific graph features influence the best parameter choice.

Figure 2: The number of considered quadruples and the number of quadruples that we call **explore** on, for each level of the hierarchy of dominating sets.

To understand how different sequences of domination distances help to reduce the running time, we conduct an experiment that examines the number of recursive calls on different levels of

the hierarchy of dominating sets. To this end, we plot the number of considered quadruples and the number of quadruples that we call **explore** on (i.e., that we recurse on) for each level of the hierarchy of dominating sets, see Figure 2. Let us first consider the figures 2a and 2b, which are both runs on the **buddha** graph: one with ratio 1.01 and the other with ratio 1.5. Note that the run with the larger ratio skips some domination distances, however, the shape and magnitude of the number of recursive calls does not significantly change. Thus, the pruning of recursive calls works equally well, even when omitting domination distances. This explains how our technique can significantly reduce the computation time. When considering Figure 2c, we can see that the number of recursive calls does not have to be unimodal, but can also exhibit different distributions. Furthermore, as can be seen in Figure 2d, the maximal number of recursive calls can also occur at quite small domination distances, even though we need to start with a large one to have a fast overall running time.

6 Conclusion

In this work, we presented a new practical algorithm to compute the graph hyperbolicity of very large graphs (compared to the graph sizes that were feasible for previous algorithms). The main idea of our algorithm is to construct a hierarchy of dominating sets and use these to prune a large number of quadruples that need to be considered as candidates for the hyperbolicity of the graph. Our approach is especially suited for graphs with small dominating sets and a hyperbolicity that is not very small. For the first time we can compute the hyperbolicity of non-trivial graphs with more than a million nodes and we obtain speed-ups of three orders of magnitude while reducing the memory consumption by more than an order of magnitude.

A problem that we leave open is how to derive the parameters that lead to the lowest running time and memory consumption for our algorithm in an automated fashion. Another open question is how to use the notion of far-apart pairs to further reduce the number quadruples to consider.

References

- [1] Muad Abu-Ata and Feodor F. Dragan. Metric tree-like structures in real-life networks: an empirical study. *Networks*, 67(1):49–69, 2016. doi:10.1002/net.21631.
- [2] Aaron B. Adcock, Blair D. Sullivan, Oscar R. Hernandez, and Michael W. Mahoney. Evaluating openmp tasking at scale for the computation of graph hyperbolicity. In *9th International Workshop on OpenMP - IWOMP*, volume 8122 of *Lecture Notes in Computer Science*, pages 71–83, Canberra, Australia, 2013. Springer. doi:10.1007/978-3-642-40698-0_6.
- [3] Takuya Akiba, Yoichi Iwata, and Yuichi Yoshida. Fast exact shortest-path distance queries on large networks by pruned landmark labeling. In *ACM SIGMOD International Conference on Management of Data - SIGMOD*, pages 349–360, 2013. doi:10.1145/2463676.2465315.
- [4] Réka Albert, Bhaskar DasGupta, and Nasim Mobasher. Topological implications of negative curvature for biological and social networks. *Physical Review E*, 89(3):032811, 2014. doi:10.1103/PhysRevE.89.032811.
- [5] Hend Alrasheed and Feodor F. Dragan. Core-periphery models for graphs based on their δ -hyperbolicity: An example using biological networks. In *6th Workshop on Complex Networks -*

- CompleNet*, volume 597 of *Studies in Computational Intelligence*, pages 65–77. Springer, 2015. doi:10.1007/978-3-319-16112-9_7.
- [6] Hans-Jürgen Bandelt and Henry Martyn Mulder. Distance-hereditary graphs. *Journal of Combinatorial Theory, Series B*, 41(2):182–208, 1986. doi:10.1016/0095-8956(86)90043-2.
 - [7] Sergio Bermudo, José M. Rodríguez, José M. Sigarreta, and Jean-Marie Vialaire. Gromov hyperbolic graphs. *Discrete Mathematics*, 313(15):1575–1585, 2013. doi:10.1016/j.disc.2013.04.009.
 - [8] Anne Berry, Romain Pogorelcnik, and Geneviève Simonet. An introduction to clique minimal separator decomposition. *Algorithms*, 3(2):197–215, 2010. doi:10.3390/a3020197.
 - [9] Marián Boguñá, Fragkiskos Papadopoulos, and Dmitri V. Krioukov. Sustaining the Internet with hyperbolic mapping. *Nature Communications*, 1(62):1–18, October 2010. doi:10.1038/ncomms1063.
 - [10] John Adrian Bondy and Uppaluri Siva Ramachandra Murty. *Graph theory with applications*, volume 290. Macmillan London, 1976.
 - [11] Michele Borassi, Alessandro Chessa, and Guido Caldarelli. Hyperbolicity measures democracy in real-world networks. *Physical Review E*, 92(3):032812, 2015. doi:10.1103/PhysRevE.92.032812.
 - [12] Michele Borassi, David Coudert, Pierluigi Crescenzi, and Andrea Marino. On computing the hyperbolicity of real-world graphs. In *European Symposium on Algorithms - ESA*, volume 9294 of *Lecture Notes in Computer Science*, pages 215–226. Springer, September 2015. doi:10.1007/978-3-662-48350-3_19.
 - [13] John Chakerian and Susan Holmes. distory: Distance between phylogenetic histories, 2010. URL: <http://cran.r-project.org/web/packages/distory/>.
 - [14] John Chakerian and Susan Holmes. Computational tools for evaluating phylogenetic and hierarchical clustering trees. *Journal of Computational and Graphical Statistics*, 21(3):581–599, 2012. doi:10.1080/10618600.2012.640901.
 - [15] Pierre Charbit, Fabien de Montgolfier, and Mathieu Raffinot. Linear time split decomposition revisited. *SIAM Journal on Discrete Mathematics*, 26(2):499–514, 2012. doi:10.1137/10080052X.
 - [16] Victor Chepoi, Feodor F. Dragan, Bertrand Estellon, Michel Habib, Yann Vaxès, and Yang Xiang. Additive spanners and distance and routing labeling schemes for hyperbolic graphs. *Algorithmica*, 62(3-4):713–732, 2012. doi:10.1007/s00453-010-9478-x.
 - [17] Nathann Cohen, David Coudert, Guillaume Ducoffe, and Aurélien Lancin. Applying clique-decomposition for computing gromov hyperbolicity. *Theoretical Computer Science*, 690:114–139, 2017. doi:10.1016/j.tcs.2017.06.001.
 - [18] Nathann Cohen, David Coudert, and Aurélien Lancin. On computing the gromov hyperbolicity. *ACM Journal of Experimental Algorithmics*, 20:1.6:1–1.6:18, 2015. doi:10.1145/2780652.

- [19] David Coudert and Guillaume Ducoffe. Recognition of C_4 -free and 1/2-hyperbolic graphs. *SIAM Journal on Discrete Mathematics*, 28(3):1601–1617, September 2014. doi:10.1137/140954787.
- [20] David Coudert, Guillaume Ducoffe, and Alexandru Popa. Fully polynomial FPT algorithms for some classes of bounded clique-width graphs. *ACM Transactions on Algorithms*, 15(3):1–57, June 2019. doi:10.1145/3310228.
- [21] David Coudert, André Nusser, and Laurent Viennot. Enumeration of far-apart pairs by decreasing distance for faster hyperbolicity computation. Research report, Inria ; I3S, Université Côte d’Azur, April 2021. URL: <https://hal.inria.fr/hal-03201405>.
- [22] David Coudert, André Nusser, and Laurent Viennot. Hyperbolicity (version 2.0). <https://gitlab.inria.fr/dcoudert/hyperbolicity/>, 2021.
- [23] William H. Cunningham. Decomposition of directed graphs. *SIAM Journal on Algebraic Discrete Methods*, 3(2):214–228, 1982. doi:10.1137/0603021.
- [24] William H. Cunningham and Jack Edmonds. A combinatorial decomposition theory. *Canadian Journal of Mathematics*, 32(3):734–765, 1980. doi:10.4153/CJM-1980-057-7.
- [25] Pierre de La Harpe and Etienne Ghys. *Sur les groupes hyperboliques d’après Mikhael Gromov*, volume 83. Progress in Mathematics, 1990. doi:10.1007/978-1-4684-9167-8.
- [26] Daniel Delling, Andrew V. Goldberg, Thomas Pajor, and Renato F. Werneck. Robust distance queries on massive networks. In Andreas S. Schulz and Dorothea Wagner, editors, *European Symposium on Algorithms - ESA*, volume 8737 of *Lecture Notes in Computer Science*, pages 321–333. Springer, 2014. doi:10.1007/978-3-662-44777-2_27.
- [27] Reinhard Diestel. *Graph Theory, 5th edition*, volume 173 of *Graduate Texts in Mathematics*. Springer, Heidelberg, 2017.
- [28] Feodor F. Dragan. Tree-like structures in graphs: A metric point of view. In *39th International Workshop on Graph-Theoretic Concepts in Computer Science - WG*, volume 8165 of *Lecture Notes in Computer Science*, pages 1–4. Springer, 2013. doi:10.1007/978-3-642-45043-3_1.
- [29] Feodor F. Dragan, Michel Habib, and Laurent Viennot. Revisiting radius, diameter, and all eccentricity computation in graphs through certificates. *CoRR*, abs/1803.04660, 2018. arXiv:1803.04660.
- [30] Andreas Dress, Katharina Huber, Jacobus Koolen, Vincent Moulton, and Andreas Spillner. *Basic Phylogenetic Combinatorics*. Cambridge University Press, Cambridge, UK, December 2011. doi:10.1017/CB09781139019767.
- [31] Hervé Fournier, Anas Ismail, and Antoine Vigneron. Computing the gromov hyperbolicity of a discrete metric space. *Information Processing Letters*, 115(6):576–579, 2015. doi:10.1016/j.ipl.2015.02.002.
- [32] Tibor Gallai. Transitiv orientierbar graphen. *Acta Mathematica Hungarica*, 18(1):25–66, 1967. doi:10.1007/BF02020961.

- [33] Cyril Gavaille, David Peleg, Stéphane Pérennes, and Ran Raz. Distance labeling in graphs. *Journal of Algorithms*, 53(1):85–112, 2004. doi:10.1016/j.jalgor.2004.05.002.
- [34] Micha Gromov. Hyperbolic groups. In S.M. Gersten, editor, *Essays in Group Theory*, volume 8 of *Mathematical Sciences Research Institute Publications*, pages 75–263. Springer, New York, 1987. doi:10.1007/978-1-4613-9586-7_3.
- [35] Michel Habib and Christophe Paul. A survey of the algorithmic aspects of modular decomposition. *Computer Science Review*, 4(1):41–59, 2010. doi:10.1016/j.cosrev.2010.01.001.
- [36] Edward Howorka. On metric properties of certain clique graphs. *Journal of Combinatorial Theory, Series B*, 27(1):67–74, 1979. doi:10.1016/0095-8956(79)90069-8.
- [37] W. Sean Kennedy, Iraj Saniee, and Onuttom Narayan. On the hyperbolicity of large-scale networks and its estimation. In *International Conference on Big Data*, pages 3344–3351. IEEE, 2016. doi:10.1109/BigData.2016.7840994.
- [38] Robert Krauthgamer and James R. Lee. Algorithms on negatively curved spaces. In *IEEE Symposium on Foundations of Computer Science - FOCS*, pages 119–132. IEEE, 2006. doi:10.1109/FOCS.2006.9.
- [39] François Le Gall. Faster algorithms for rectangular matrix multiplication. In *IEEE Symposium on Foundations of Computer Science - FOCS*, pages 514–523, New Brunswick, NJ, USA, 2012. IEEE. doi:10.1109/focs.2012.80.
- [40] Wentao Li, Miao Qiao, Lu Qin, Ying Zhang, Lijun Chang, and Xuemin Lin. Exacting eccentricity for small-world networks. In *IEEE International Conference on Data Engineering - ICDE*, pages 785–796, April 2018. doi:10.1109/ICDE.2018.00076.
- [41] Tamara Munzner and Paul Burchard. Visualizing the structure of the world wide web in 3d hyperbolic space. In David R. Nadeau and John L. Moreland, editors, *Symposium on Virtual Reality Modeling Language - VRML*, pages 33–38. ACM, 1995. doi:10.1145/217306.217311.
- [42] Yuval Shavitt and Tomer Tankel. On the curvature of the internet and its usage for overlay construction and distance estimation. In *Annual Joint Conference of the IEEE Computer and Communications Societies - INFOCOM*, 2004. doi:10.1109/INFOCOM.2004.1354510.
- [43] Mauricio Abel Soto Gómez. *Quelques propriétés topologiques des graphes et applications à internet et aux réseaux*. PhD thesis, Univ. Paris Diderot (Paris 7), 2011.
- [44] Frank W. Takes and Walter A. Kosters. Computing the eccentricity distribution of large graphs. *Algorithms*, 6(1):100–118, 2013. doi:10.3390/a6010100.
- [45] Robert Endre Tarjan. Decomposition by clique separators. *Discrete Mathematics*, 55(2):221–232, 1985. doi:10.1016/0012-365X(85)90051-2.
- [46] Jörg A. Walter and Helge J. Ritter. On interactive visualization of high-dimensional data using the hyperbolic plane. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD*, pages 123–132. ACM, 2002. doi:10.1145/775047.775065.