Fast Approximation and Exact Computation of Negative Curvature Parameters of Graphs

Jérémie Chalopin
CNRS, Aix-Marseille Université, Université de Toulon, LIS, Marseille, France
jeremie.chalopin@lis-lab.fr

Victor Chepoi
Aix-Marseille Université, CNRS, Université de Toulon, LIS, Marseille, France
victor.chepoi@lis-lab.fr

Feodor F. Dragan
Computer Science Department, Kent State University, Kent, USA
dragan@cs.kent.edu

Guillaume Ducoffe
National Institute for Research and Development in Informatics and Research Institute of the University of Bucharest, Bucureşti, Romînia
guillaume.ducoffe@ici.ro

Abdulhakeem Mohammed
Computer Science Department, Kent State University, Kent, USA
amohamm4@kent.edu

Yann Vaxès
Aix-Marseille Université, CNRS, Université de Toulon, LIS, Marseille, France
yann.vaxes@lis-lab.fr

──────── **Abstract** ────────

In this paper, we study Gromov hyperbolicity and related parameters, that represent how close (locally) a metric space is to a tree from a metric point of view. The study of Gromov hyperbolicity for geodesic metric spaces can be reduced to the study of *graph hyperbolicity*. Our main contribution in this note is a new characterization of hyperbolicity for graphs (and for complete geodesic metric spaces). This characterization has algorithmic implications in the field of large-scale network analysis, which was one of our initial motivations. A sharp estimate of graph hyperbolicity is useful, e.g., in embedding an undirected graph into hyperbolic space with minimum distortion [Verbeek and Suri, SoCG'14]. The hyperbolicity of a graph can be computed in polynomial-time, however it is unlikely that it can be done in *subcubic* time. This makes this parameter difficult to compute or to approximate on large graphs. Using our new characterization of graph hyperbolicity, we provide a simple factor 8 approximation algorithm for computing the hyperbolicity of an n-vertex graph $G = (V, E)$ in optimal time $O(n^2)$ (assuming that the input is the distance matrix of the graph). This algorithm leads to constant factor approximations of other graph-parameters related to hyperbolicity (thinness, slimness, and insize). We also present the first efficient algorithms for exact computation of these parameters. All of our algorithms can be used to approximate the hyperbolicity of a geodesic metric space.

1   Introduction

Understanding the geometric properties of complex networks is a key issue in network analysis and geometry of graphs. One important such property is the *negative curvature* [29], causing the traffic between the vertices to pass through a relatively small core of the network – as if the shortest paths between them were curved inwards. It has been empirically observed, then formally proved [14], that such a phenomenon is related to the value of the *Gromov hyperbolicity* of the graph. In this paper, we propose exact and approximation algorithms to compute hyperbolicity of a graph and its relatives (the approximation algorithms can be applied to geodesic metric spaces as well).

A metric space (X, d) is δ-*hyperbolic* [3, 9, 26] if for any four points w, v, x, y of X, the two largest of the distance sums $d(w, v) + d(x, y)$, $d(w, x) + d(v, y)$, $d(w, y) + d(v, x)$ differ by at most $2\delta \geq 0$. A graph $G = (V, E)$ endowed with its standard graph-distance d_G is δ-*hyperbolic* if the metric space (X, d_G) is δ-hyperbolic. In case of geodesic metric spaces and graphs, δ-hyperbolicity can be defined in other equivalent ways, e.g., via thin or slim geodesic triangles. The hyperbolicity $\delta(X)$ of a metric space X is the smallest $\delta \geq 0$ such that X is δ-hyperbolic. The hyperbolicity $\delta(X)$ can be viewed as a local measure of how close X is to a tree: the smaller the hyperbolicity is, the closer the metrics of its 4-point subspaces are close to tree-metrics.

The study of hyperbolicity of graphs is motivated by the fact that many real-world graphs are tree-like from a metric point of view [1, 2, 6] or have small hyperbolicity [28, 29, 32]. This is due to the fact that many of these graphs (including Internet application networks, web networks, collaboration networks, social networks, biological networks, and others) possess certain geometric and topological characteristics. Hence, for many applications, including the design of efficient algorithms (cf., e.g., [6, 11–15, 19, 22, 34]), it is useful to know the hyperbolicity $\delta(G)$ of a graph G.

Related work.   For an n-vertex graph G, the definition of hyperbolicity directly implies a simple brute-force $O(n^4)$ algorithm to compute $\delta(G)$. This running time is too slow for computing the hyperbolicity of large graphs that occur in applications [1, 6, 7, 24]. On the theoretical side, it was shown that relying on matrix multiplication results, one can improve the upper bound on time-complexity to $O(n^{3.69})$ [24]. Moreover, roughly quadratic lower bounds are known [7, 17, 24]. In practice, however, the best known algorithm still has an $O(n^4)$-time worst-case bound but uses several clever tricks when compared to the brute-force algorithm [6]. Based on empirical studies, an $O(mn)$ running time is claimed, where m is the number of edges in the graph. Furthermore, there are heuristics for computing the hyperbolicity of a given graph [16], and there are investigations whether one can compute hyperbolicity in linear time when some graph parameters take small values [18, 23].

Perhaps it is interesting to notice that the first algorithms for testing graph hyperbolicity were designed for Cayley graphs of finitely generated groups (these are infinite vertex-transitive graphs of uniformly bounded degrees). Gromov gave an algorithm to recognize Cayley graphs of hyperbolic groups and estimate the hyperbolicity constant δ. His algorithm is based on the theorem that in Cayley graphs, the hyperbolicity "propagates", i.e., if balls of an appropriate fixed radius induce a δ-hyperbolic space, then the whole space is δ'-hyperbolic for some $\delta' > \delta$ (see [26], 6.6.F and [20]). Therefore, in order to check the hyperbolicity of a Cayley graph, it is enough to verify the hyperbolicity of a sufficiently big ball (all balls of a given radius in a Cayley graph are isomorphic to each other). For other algorithms deciding if the Cayley graph of a finitely generated group is hyperbolic, see [8, 30]. However, similar methods do not help when dealing with arbitrary graphs.

By a result of Gromov [26], if the four-point condition in the definition of hyperbolicity holds for a fixed basepoint w and any triplet x, y, v of X, then the metric space (X, d) is 2δ-hyperbolic. This provides a factor 2 approximation of hyperbolicity of a metric space on n points running in cubic $O(n^3)$ time. Using fast algorithms for computing (max,min)-matrix products, it was noticed in [24] that this 2-approximation of hyperbolicity can be implemented in $O(n^{2.69})$ time. In the same paper, it was shown that any algorithm computing the hyperbolicity for a fixed basepoint in time $O(n^{2.05})$ would provide an algorithm for (max, min)-matrix multiplication faster than the existing ones. In [21], approximation algorithms are given to compute a $(1 + \epsilon)$-approximation in $O(\epsilon^{-1}n^{3.38})$ time and a $(2 + \epsilon)$-approximation in $O(\epsilon^{-1}n^{2.38})$ time. As a direct application of the characterization of hyperbolicity of graphs via a cop and robber game and dismantlability, [11] presents a simple constant factor approximation algorithm for hyperbolicity of G running in optimal $O(n^2)$ time. Its approximation ratio is huge (1569), however it is believed that its theoretical performance is much better and the factor of 1569 is mainly due to the use in the proof of the definition of hyperbolicity via linear isoperimetric inequality. This shows that the question of designing fast and (theoretically certified) accurate algorithms for approximating graph hyperbolicity is still an important and open question.

Our contribution.   In this paper, we tackle this open question and propose a very simple (and thus practical) factor 8 algorithm for approximating the hyperbolicity $\delta(G)$ of an n-vertex graph G running in optimal $O(n^2)$ time. As in several previous algorithms, we assume that the input is the distance matrix D of the graph G. Our algorithm picks a basepoint w, a Breadth-First-Search tree T rooted at w, and considers only geodesic triangles of G with one vertex at w and two sides on T. For all such sides in T, it computes the maximum over all distances between the two preimages of the centers of the respective tripods. This maximum $\rho_{w,T}$ can be easily computed in $O(n^2)$ time and provides an 8-approximation for $\delta(G)$. For complete geodesic spaces (X, d), we show that we can always define a geodesic spanning tree T based at any point w of X, and that the same relationships between $\rho_{w,T}$ and the hyperbolicity of (X, d) hold, thus providing a new characterization of hyperbolicity. Perhaps it is surprising that hyperbolicity that is originally defined via quadruplets and can be 2-approximated via triplets (i.e., via pointed hyperbolicity), can be finally defined and approximated only via pairs (and an arbitrary fixed BFS-tree). We hope that this new characterization can be useful in establishing that graphs and simplicial complexes occurring in geometry and in network analysis are hyperbolic.

The way $\rho_{w,T}$ is computed is closely related to how hyperbolicity is defined via slimness, thinness, and insize of its geodesic triangles. Similarly to the hyperbolicity $\delta(G)$, one can define slimness $\varsigma(G)$, thinness $\tau(G)$, and insize $\iota(G)$ of a graph G. As a direct consequence of our

algorithm for approximating $\delta(G)$ and the relationships between $\delta(G)$ and $\varsigma(G), \tau(G), \iota(G)$, we obtain constant factor $O(n^2)$ time algorithms for approximating these parameters. On the other hand, an *exact* computation, in polynomial time, of these geometric parameters has long stayed elusive. This is due to the fact that $\varsigma(G), \tau(G), \iota(G)$ are defined as minima of some functions over all the geodesic triangles of G, and that there may be exponentially many such triangles. In this paper we provide the first polynomial time algorithms for computing $\varsigma(G), \tau(G)$, and $\iota(G)$. Namely, we show that the thinness $\tau(G)$ and the insize $\iota(G)$ of G can be computed in $O(n^2m)$ time and the slimness $\varsigma(G)$ of G can be computed in $\widehat{O}(n^2m + n^4/\log^3 n)$ time[1]. However, we show that the minimum value of of $\rho_{w,T}$ over all basepoints w and all BFS-trees T cannot be approximated with a factor strictly better than 2 unless P = NP.

2 Gromov hyperbolicity and its relatives

2.1 Gromov hyperbolicity

Let (X, d) be a metric space and $w \in X$. The *Gromov product*[2] of $y, z \in X$ with respect to w is $(y|z)_w = \frac{1}{2}(d(y, w) + d(z, w) - d(y, z))$. A metric space (X, d) is δ-*hyperbolic* [26] for $\delta \geq 0$ if $(x|y)_w \geq \min\{(x|z)_w, (y|z)_w\} - \delta$ for all $w, x, y, z \in X$. Equivalently, (X, d) is δ-hyperbolic if for any $u, v, x, y \in X$, the two largest of the sums $d(u, v) + d(x, y)$, $d(u, x) + d(v, y)$, $d(u, y) + d(v, x)$ differ by at most $2\delta \geq 0$. A metric space (X, d) is said to be δ-*hyperbolic with respect to a basepoint* w if $(x|y)_w \geq \min\{(x|z)_w, (y|z)_w\} - \delta$ for all $x, y, z \in X$.

▶ **Proposition 1** ([3, 9, 25, 26]). *If (X, d) is δ-hyperbolic with respect to some basepoint, then (X, d) is 2δ-hyperbolic.*

Let (X, d) be a metric space. An (x, y)-*geodesic* is a (continuous) map $\gamma : [0, d(x, y)] \to X$ from the segment $[0, d(x, y)]$ of $\mathbb{R}^1$ to X such that $\gamma(0) = x, \gamma(d(x, y)) = y$, and $d(\gamma(s), \gamma(t)) = |s - t|$ for all $s, t \in [0, d(x, y)]$. A *geodesic segment* with endpoints x and y is the image of the map γ (when it is clear from the context, by a geodesic we mean a geodesic segment and we denote it by $[x, y]$). A metric space (X, d) is *geodesic* if every pair of points in X can be joined by a geodesic. A *real tree* (or an $\mathbb{R}$-*tree*) [9, p.186] is a geodesic metric space (T, d) such that

(1) there is a unique geodesic $[x, y]$ joining each pair of points $x, y \in T$;
(2) if $[y, x] \cap [x, z] = \{x\}$, then $[y, x] \cup [x, z] = [y, z]$.

Let (X, d) be a geodesic metric space. A *geodesic triangle* $\Delta(x, y, z)$ with $x, y, z \in X$ is the union $[x, y] \cup [x, z] \cup [y, z]$ of three geodesics connecting these points. A geodesic triangle $\Delta(x, y, z)$ is called δ-*slim* if for any point u on the side $[x, y]$ the distance from u to $[x, z] \cup [z, y]$ is at most δ. Let m_x be the point of $[y, z]$ located at distance $\alpha_y := (x|z)_y$ from y. Then, m_x is located at distance $\alpha_z := (y|x)_z$ from z because $\alpha_y + \alpha_z = d(y, z)$. Analogously, define the points $m_y \in [x, z]$ and $m_z \in [x, y]$ both located at distance $\alpha_x := (y|z)_x$ from x; see Fig. 1 for an illustration. We define a tripod $T(x, y, z)$ consisting of three solid segments $[x, m], [y, m]$, and $[z, m]$ of lengths α_x, α_y, and α_z, respectively. The function mapping the vertices x, y, z of $\Delta(x, y, z)$ to the respective leaves of $T(x, y, z)$ extends uniquely to a function $\varphi : \Delta(x, y, z) \to T(x, y, z)$ such that the restriction of φ on each side of $\Delta(x, y, z)$ is an isometry. This function maps the points m_x, m_y, and m_z to the center m of $T(x, y, z)$. Any

---

[1] The $\widehat{O}(\cdot)$ notation hides polyloglog factors.
[2] Informally, $(y|z)_w$ can be viewed as half the detour you make, when going over w to get from y to z.
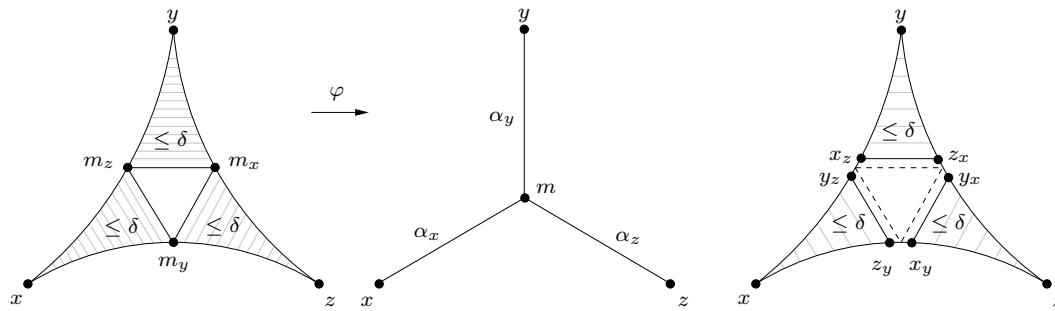
Figure 1 Insize and thinness in geodesic spaces and graphs.

other point of $T(x, y, z)$ is the image of exactly two points of $\Delta(x, y, z)$. A geodesic triangle $\Delta(x, y, z)$ is called δ-*thin* if for all points $u, v \in \Delta(x, y, z)$, $\varphi(u) = \varphi(v)$ implies $d(u, v) \leq \delta$. The *insize* of $\Delta(x, y, z)$ is the diameter of the preimage $\{m_x, m_y, m_z\}$ of the center m of the tripod $T(x, y, z)$. Below, we remind that the hyperbolicity of a geodesic space can be approximated by the maximum thinness and slimness of its geodesic triangles.

For a geodesic metric space (X, d), one can define the following parameters:

- *hyperbolicity* $\delta(X) = \min\{\delta : X \text{ is } \delta\text{-hyperbolic}\}$,
- *pointed hyperbolicity* $\delta_w(X) = \min\{\delta : X \text{ is } \delta\text{-hyperbolic with respect to a basepoint } w\}$,
- *slimness* $\varsigma(X) = \min\{\delta : \text{any geodesic triangle of } X \text{ is } \delta\text{-slim}\}$,
- *thinness* $\tau(X) = \min\{\delta : \text{any geodesic triangle of } X \text{ is } \delta\text{-thin}\}$,
- *insize* $\iota(X) = \min\{\delta : \text{the insize of any geodesic triangle of } X \text{ is at most } \delta\}$.

▶ **Proposition 2** ([3,9,25,26,33]). *For a geodesic metric space (X, d), $\delta(X) \leq \iota(X) = \tau(X) \leq 4\delta(X)$, $\varsigma(X) \leq \tau(X) \leq 4\varsigma(X)$, and $\delta(X) \leq 2\varsigma(X) \leq 3\delta(X)$.*

Due to Propositions 1 and 2, a geodesic metric space (X, d) is called *hyperbolic* if one of the numbers $\delta(X), \delta_w(X), \varsigma(X), \tau(X), \iota(X)$ (and thus all) is finite. Notice also that a geodesic metric space (X, d) is 0-hyperbolic if and only if (X, d) is a real tree [9, p.399] (and in this case, $\varsigma(X) = \tau(X) = \iota(X) = \delta(X) = 0$).

2.2 Hyperbolicity of graphs

All graphs $G = (V, E)$ occurring in this paper are undirected and connected, but not necessarily finite (in algorithmic results they will be supposed to be finite). For any two vertices $x, y \in V$, the *distance* $d(x, y)$ is the minimum number of edges in a path between x and y. Let $[x, y]$ denote a shortest path connecting vertices x and y in G; we call $[x, y]$ a *geodesic* between x and y. The *interval* $I(u, v) = \{x \in V : d(u, x) + d(x, v) = d(u, v)\}$ consists of all vertices on (u, v)-geodesics. There is a strong analogy between the metric properties of graphs and geodesic metric spaces, due to their uniform local structure. Any graph $G = (V, E)$ gives rise to a geodesic space (X_G, d) (into which G isometrically embeds) obtained by replacing each edge xy of G by a segment isometric to $[0, 1]$ with ends at x and y. X_G is called a *metric graph*. Conversely, by [9, Proposition 8.45], any geodesic metric space (X, d) is $(3, 1)$-quasi-isometric to a graph $G = (V, E)$. This graph G is constructed in the following way: let V be an open maximal $\frac{1}{3}$-packing of X, i.e., $d(x, y) > \frac{1}{3}$ for any $x, y \in V$ (that exists by Zorn's lemma). Then two points $x, y \in V$ are adjacent in G if and only if $d(x, y) \leq 1$. Since hyperbolicity is preserved (up to a constant factor) by quasi-isometries, this reduces the computation of hyperbolicity for geodesic spaces to the case of graphs.

The notions of geodesic triangles, insize, δ-slim and δ-thin triangles can also be defined in case of graphs with the single difference that for graphs, the center of the tripod is not necessarily the image of any vertex on the sides of $\Delta(x, y, z)$. For graphs, we "discretize" the notion of δ-thin triangles in the following way. We say that a geodesic triangle $\Delta(x, y, z)$ of a graph G is δ-*thin* if for any $v \in \{x, y, z\}$ and vertices $a \in [v, u]$ and $b \in [v, w]$ ($u, w \in \{x, y, z\}$, and u, v, w are pairwise different), $d(v, a) = d(v, b) \leq (u|w)_v$ implies $d(a, b) \leq \delta$. A graph G is δ-*thin*, if all geodesic triangles in G are δ-thin. Given a geodesic triangle $\Delta(x, y, z) :=$ $[x, y] \cup [x, z] \cup [y, z]$ in G, let x_y and y_x be the vertices of $[z, x]$ and $[z, y]$, respectively, both at distance $\lfloor (x|y)_z \rfloor$ from z. Similarly, one can define vertices x_z, z_x and vertices y_z, z_y; see Fig. 1. The *insize* of $\Delta(x, y, z)$ is defined as $\max\{d(y_z, z_y), d(x_y, y_x), d(x_z, z_x)\}$. An interval $I(x, y)$ is said to be κ-thin if $d(a, b) \leq \kappa$ for all $a, b \in I(x, y)$ with $d(x, a) = d(x, b)$. The smallest κ for which all intervals of G are κ-thin is called the *interval thinness* of G and denoted by $\kappa(G)$. Denote also by $\delta(G)$, $\delta_w(G)$, $\varsigma(G)$, $\tau(G)$, and $\iota(G)$ respectively the hyperbolicity, the pointed hyperbolicity with respect to a basepoint w, the slimness, the thinness, and the insize of a graph G. We will need the following inequalities between $\varsigma(G)$, $\tau(G)$, $\iota(G)$, and $\delta(G)$. They are known to be true for all geodesic spaces (see [3, 9, 25, 26, 33]):

▶ **Proposition 3.** $\delta(G) - \frac{1}{2} \leq \iota(G) = \tau(G) \leq 4\delta(G)$, $\varsigma(G) \leq \tau(G) \leq 4\varsigma(G)$, $\delta(G) - \frac{1}{2} \leq 2\varsigma(G) \leq 6\delta(G) + 1$, and $\kappa(G) \leq \min\{\tau(G), 2\delta(G), 2\varsigma(G)\}$.

3 **Geodesic spanning trees**

In this section, we outline the proof that any complete geodesic metric space (X, d) has a geodesic spanning tree rooted at any basepoint w. We hope that this general result will be useful in other contexts. For graphs this is well-known and simple, and such trees can be constructed in various ways, for example via Breadth-First-Search. The existence of BFS-trees in infinite graphs has been established by Polat [31]. However for complete geodesic spaces this result seems to be new (and not completely trivial) and we consider it as one of the main results of the paper. A *geodesic spanning tree rooted at a point w* (a *GS-tree* for short) of a geodesic space (X, d) is a union of geodesics $\Gamma_w := \bigcup_{x \in X} \gamma_{w,x}$ with one end at w such that $y \in \gamma_{w,x}$ implies that $\gamma_{w,y} \subseteq \gamma_{w,x}$. Finally recall that a metric space (X, d) is called *complete* if every Cauchy sequence of points in (X, d) has a limit in X.

▶ **Theorem 4.** *For any complete geodesic metric space (X, d) and for any basepoint w one can define a geodesic spanning tree $\Gamma_w = \bigcup_{x \in X} \gamma_{w,x}$ rooted at w and a real tree $T = (X, d_T)$ such that any $\gamma_{w,x} \in \Gamma_w$ is the unique (w, x)-geodesic of T.*

The first assertion of the theorem immediately follows from the following proposition:

▶ **Proposition 5.** *For any complete geodesic metric space (X, d), for any pair of points $x, y \in X$ one can define an (x, y)-geodesic $\gamma_{x,y}$ such that for all $x, y \in X$ and for all $u, v \in \gamma_{x,y}$, we have $\gamma_{u,v} \subseteq \gamma_{x,y}$.*

Proof. Let $\preceq$ be a well-order on X. For any $x, y \in X$ we define inductively two sets $P_{x,y}^{\prec v}$ and $P_{x,y}^v$ for any $v \in X$:

$$P_{x,y}^{\prec v} = \{x, y\} \cup \bigcup_{u \prec v} P_{x,y}^u,$$

$$P_{x,y}^v = \begin{cases} P_{x,y}^{\prec v} \cup \{v\} & \text{if there is an } (x, y)\text{-geodesic } \gamma \text{ with } P_{x,y}^{\prec v} \cup \{v\} \subseteq \gamma, \\ P_{x,y}^{\prec v} & \text{otherwise.} \end{cases}$$

We set $P_{x,y} = \bigcup_{u \in X} P_{x,y}^u$. Using transfinite induction, we prove that there exists an (x,y)-geodesic $\gamma_{x,y}^{\prec v}$ (respectively, $\gamma_{x,y}^v$, $\gamma_{x,y}$) containing $P_{x,y}^{\prec v}$ (respectively, $P_{x,y}^v$, $P_{x,y}$), and we show that $P_{x,y}$ is an (x,y)-geodesic. By the definition of $P_{x,y}$ and $P_{x,u}$, we can show that $P_{x,u} = P_{x,y} \cap B(x, d(x,u))$ for any $u \in P_{x,y}$, and it follows that $P_{u,v} \subseteq P_{x,v} \subseteq P_{x,y}$ for any $u, v \in P_{x,y}$ such that $d(x,u) \leq d(x,v)$. ◂

Consequently, $\Gamma_w = \bigcup_{x \in X} \gamma_{w,x}$ is a geodesic spanning tree of (X,d) rooted at w. Using Γ_w, we can define a real tree T as follows. For any $x \in X$, denote by $[w,x]$ the geodesic segment between w and x which is the image of the geodesic map $\gamma_{w,x}$. From the definition of Γ_w, if $x' \in [w,x]$, then $[w,x'] \subseteq [w,x]$. From the continuity of geodesic maps and the definition of Γ_w it follows that for any two geodesics $\gamma_{w,x}, \gamma_{w,y} \in \Gamma_w$ the intersection $[w,x] \cap [w,y]$ is the image $[w,z]$ of some geodesic $\gamma_{w,z} \in \Gamma_w$. Call z the *lowest common ancestor* of x and y (with respect to the root w) and denote it by $\mathrm{lca}(x,y)$. Define d_T by setting $d_T(w,x) := d(w,x)$ and $d_T(x,y) := d(w,x) + d(w,y) - 2d(w,z) = d(x,z) + d(z,y)$ for any two points $x, y \in X$. We prove that $T = (X, d_T)$ is a real tree. To do so, we show in particular that for any $x, y \in X$, $[x,z] \cup [z,y]$ is the unique (x,y)-geodesic in T where $z = \mathrm{lca}(x,y)$, $[x,z]$ is the portion of the geodesic segment $[w,x]$ between x and z, and $[z,y]$ is the portion of the geodesic segment $[w,y]$ between z and y.

4 Fast approximation

In this section, we introduce a new parameter ρ of a graph G (or of a geodesic space X). This parameter depends on an arbitrary fixed BFS-tree of G (or a GS-tree of X). It can be computed efficiently and it provides constant-factor approximations for $\delta(G)$, $\varsigma(G)$, and $\tau(G)$. In particular, we obtain a very simple factor 8 approximation algorithm for the hyperbolicity $\delta(G)$ of an n-vertex graph G running in optimal $O(n^2)$ time (assuming that the input is the distance matrix of G).

4.1 Fast approximation of hyperbolicity

Consider a graph $G = (V, E)$ or a complete geodesic space (X, d) and an arbitrary BFS-tree T or GS-tree T, respectively, rooted at some vertex or point w (see Section 3). Denote by x_y the point of $[w,x]_T$ at distance $\lfloor (x|y)_w \rfloor$ (resp., $(x|y)_w$) from w and by y_x the point of $[w,y]_T$ at distance $\lfloor (x|y)_w \rfloor$ (resp., $(x|y)_w$) from w. In case of graphs, x_y and y_x are vertices of G. Let $\rho_{w,T} := \sup\{d(x_y, y_x) : x, y \in X\}$. In some sense, $\rho_{w,T}$ can be seen as the insize of G with respect to w and T: the differences between $\rho_{w,T}$ and $\iota(G)$ are that we consider only geodesic triangles $\Delta(w, x, y)$ containing w where the geodesics $[w,x]$ and $[w,y]$ belong to T, and we consider only $d(x_y, y_x)$, instead of $\max\{d(x_y, y_x), d(x_w, w_x), d(y_w, w_y)\}$. Using T, we can also define the thinness of G with respect to w and T: let $\mu_{w,T} = \sup\{d(x', y') : \exists x, y \text{ such that } x' \in [w,x]_T, y' \in [w,y]_T \text{ and } d(w,x') = d(w,y') \leq (x|y)_w\}$. Using the same ideas as in the proofs of Propositions 2 and 3 establishing that $\iota(X) = \tau(X)$ and $\iota(G) = \tau(G)$, we can show that these two definitions give rise to the same value.

▶ **Proposition 6.** *For any geodesic space X and any GS-tree T rooted at a point w, $\rho_{w,T} = \mu_{w,T}$. Analogously, for any graph G and any BFS-tree T rooted at w, $\rho_{w,T} = \mu_{w,T}$.*

In the following, when w and T are clear from the context, we denote $\rho_{w,T}$ by ρ. The next theorem is the main result of this paper. It establishes that 2ρ provides an 8-approximation of the hyperbolicity of $\delta(G)$ or $\delta(X)$, and that in the case of a finite graph G, ρ can be computed in $O(n^2)$ time when the distance matrix D of G is given.

▶ **Theorem 7.** *Given a graph G (respectively, a geodesic space X) and a BFS-tree T (respectively, a GS-tree T) rooted at w,*

(1) $\delta(G) \leq 2\rho_{w,T} + 1 \leq 8\delta(G) + 1$ *(respectively, $\delta(X) \leq 2\rho_{w,T} \leq 8\delta(X)$).*

(2) *If G has n vertices, given the distance matrix D of G, $\rho_{w,T}$ can be computed in $O(n^2)$ time. Consequently, an 8-approximation (with an additive constant 1) of the hyperbolicity $\delta(G)$ of G can be found in $O(n^2)$ time.*

Proof. We prove the first assertion of the theorem for graphs (for geodesic spaces, the proof is similar). Let $\rho := \rho_{w,T}$, $\delta := \delta(G)$ and $\delta_w := \delta_w(G)$. By Gromov's Proposition 1, $\delta \leq 2\delta_w$. We proceed in two steps. In the first step, we show that $\rho \leq 4\delta$. In the second step, we prove that $\delta_w \leq \rho + \frac{1}{2}$. Hence, combining both steps we obtain $\delta \leq 2\delta_w \leq 2\rho + 1 \leq 8\delta + 1$.

The first assertion follows from Proposition 3 and from the inequality $\rho \leq \iota(G) = \tau(G)$. To prove that $\delta_w \leq \rho + \frac{1}{2}$, for any quadruplet x, y, z, w containing w, we show the four-point condition $d(x,z) + d(y,w) \leq \max\{d(x,y) + d(z,w), d(y,z) + d(x,w)\} + (2\rho + 1)$. Assume without loss of generality that $d(x,z) + d(y,w) \geq \max\{d(x,y) + d(z,w), d(y,z) + d(x,w)\}$ and that $d(w, x_y) = d(w, y_x) \leq d(w, y_z) = d(w, z_y)$. From the definition of ρ, $d(x_y, y_x) \leq \rho$ and $d(y_z, z_y) \leq \rho$. Consequently, by the definition of x_y, y_x, y_z, z_y and by the triange inequality, we get

$$
\begin{aligned}
d(y,w) + d(x,z) &\leq d(y,w) + d(x, x_y) + d(x_y, y_x) + d(y_x, y_z) + d(y_z, z_y) + d(z_y, z) \\
&\leq (d(y, y_z) + d(y_z, w)) + d(x, x_y) + \rho + d(y_x, y_z) + \rho + d(z_y, z) \\
&= d(y, y_z) + d(w, z_y) + d(x, x_y) + d(y_x, y_z) + d(z_y, z) + 2\rho \\
&= d(y, y_z) + d(x, x_y) + (d(y, y_x) - d(y, y_z)) + d(w, z) + 2\rho \\
&\leq d(x,y) + 1 + d(w, z) + 2\rho,
\end{aligned}
$$

the last inequality following from the definition of x_y and y_x in graphs (in the case of geodesic metric spaces, we have $d(x, x_y) + d(y, y_x) = d(x, y)$). This establishes the four-point condition for w, x, y, z, and proves that $\delta_w \leq \rho + \frac{1}{2}$.

We present now a simple self-contained algorithm for computing ρ in $O(n^2)$ time when $G = (V, E)$ is a graph with n vertices. (Its space complexity can be improved using the algorithm of [4,5] for computing level d ancestors in trees.) For any non-negative integer r, let $x(r)$ be the unique vertex of $[w, x]_T$ at distance r from w if $r < d(w, x)$ and the vertex x if $r \geq d(w, x)$. First, we compute in $O(n^2)$ time a table M with lines indexed by V, columns indexed by $\{1, \ldots, n\}$, and such that $M(x, r)$ is the identifier of the vertex $x(r)$ of $[w, x]_T$ located at distance r from w. To compute this table, we explore the tree T starting from w. Let x be the current vertex and r its distance to the root w. For every vertex y in the subtree of T rooted at x, we set $M(y, r) := x$. Assuming that the table M and the distance matrix $D := (d(u, v) : u, v \in X)$ between the vertices of G are available, we can compute $x_y = M(x, \lfloor (x|y)_w \rfloor)$, $y_x = M(y, \lfloor (x|y)_w \rfloor)$ and $d(x_y, y_x)$ in constant time for each pair of vertices x, y, and thus $\rho = \max\{d(x_y, y_x) : x, y \in V\}$ can be computed in $O(n^2)$ time.   ◀

Theorem 7 provides a new characterization of infinite hyperbolic graphs.

▶ **Corollary 8.** *Consider an infinite graph G and an arbitrary BFS-tree T rooted at a vertex w. The graph G is hyperbolic if and only if $\rho_{w,T} < \infty$.*

The following result shows that the bounds in Theorem 7 are optimal.

▶ **Proposition 9.** *For any positive integer k, there exists a graph H_k, a vertex w, and a BFS-tree T rooted at w such that $\delta(H_k) = k$ and $\rho_{w,T} = 4k$.*
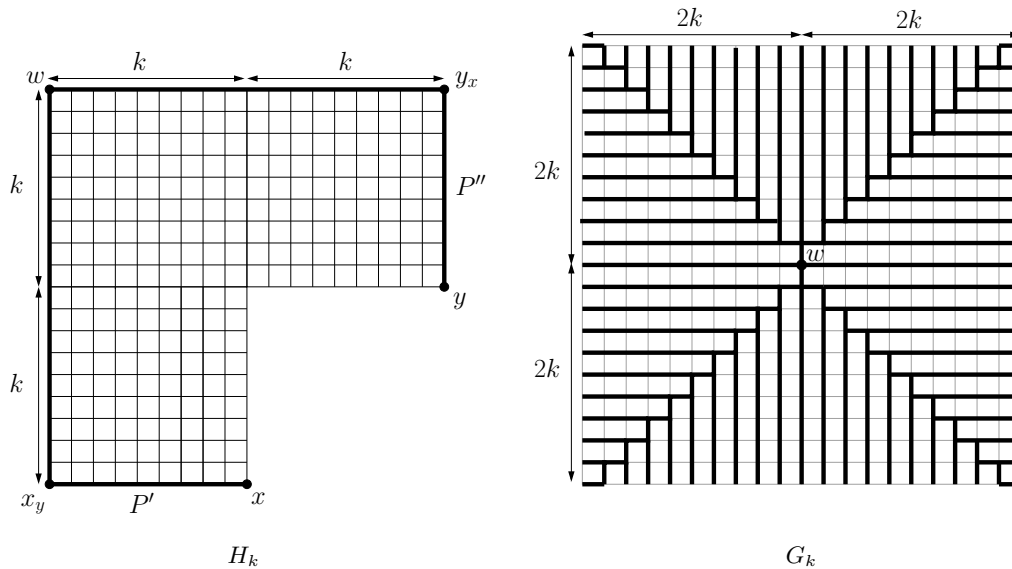
Figure 2 In H_k, $\rho_{w,T} = d(x_y, y_x) = 4\delta(H_k)$, showing that the inequality $\rho \leq 4\delta$ is tight in the proof of Theorem 7. In G_k, $\rho_{w,T} \leq 2k = \frac{1}{2}\delta(G_k)$, showing that (up to an additive factor of 1) the inequality $\delta \leq 2\rho + 1$ is tight in the proof of Theorem 7.

For any positive integer k, there exists a graph G_k, a vertex w, and a BFS-tree T rooted at w such that $\rho_{w,T} \leq 2k$ and $\delta(G_k) = 4k$.

Proof. The graph H_k is the $2k \times 2k$ square grid from which we removed the vertices of the rightmost and downmost $(k-1) \times (k-1)$ square (see Fig. 2, left). The graph H_k is a median graph and therefore its hyperbolicity is the size of a largest isometrically embedded square subgrid [12, 27]. The largest square subgrid of H_k has size k, thus $\delta(H_k) = k$.

Let w be the leftmost upmost vertex of H_k. Let x be the downmost rightmost vertex of H_k and y be the rightmost downmost vertex of H_k. Then $d(x, y) = 2k$ and $d(x, w) = d(y, w) = 3k$. Let P' and P'' be the shortest paths between w and x and w and y, respectively, running on the boundary of H_k. Let T be any BFS-tree rooted at w and containing the shortest paths P' and P''. The vertices $x_y \in P'$ and $y_x \in P''$ are located at distance $(x|y)_w = 2k$ from w. Thus x_y is the leftmost downmost vertex and y_x is the rightmost upmost vertex. Hence $\rho_{w,T} \geq d(x_y, y_x) = 4k$. Since the diameter of H_k is $4k$, we conclude that $\rho_{w,T} = 4k = 4\delta(H_k)$.

Let G_k be the $4k \times 4k$ square grid and note that $\delta(G_k) = 4k$. Let w be the center of G_k. Note that $\delta(G_k) = 4k$. We suppose that G_k is isometrically embedded in the ℓ_1-plane in such a way that w is mapped to the origin of coordinates $(0, 0)$ and the four corners of G_k are mapped to the points with coordinates $(2k, 2k), (-2k, 2k), (-2k, -2k), (2k, -2k)$, We build the BFS-tree T of G_k as follows. First we connect w to each of the corners of G_k by a shortest zigzagging path (see Fig. 2, right). For each $i \leq i \leq k$, we add a vertical path from (i, i) to $(i, 2k)$, from $(i, -i)$ to $(i, -2k)$, from $(-i, i)$ to $(-i, 2k)$, and from $(-i, -i)$ to $(-i, -2k)$. Similarly, for each $i \leq i \leq k$, we add a horizontal path from (i, i) to $(2k, i)$, from $(i, -i)$ to $(2k, -i)$, from $(-i, i)$ to $(-2k, i)$, and from $(-i, -i)$ to $(-2k, -i)$. For any vertex v, the shortest path of G_k connecting w to v in T has the following structure: it starts by a subpath of one of the zigzagging paths until it reaches the vertical or horizontal line containing v and then it continues along this line until v. One can show that $\rho_{w,T} \leq 2k = \frac{1}{2}\delta(G_k)$. ◄

The definition of $\rho_{w,T}$ depends on the choice of the basepoint w and of the BFS-tree T rooted at w. We show below that the best choices of w and T do not improve the bounds in Theorem 7. For a graph G, let $\rho_-(G) = \min\{\rho_{w,T} : w \in V$ and T is a BFS-tree rooted at $w\}$ and call $\rho_-(G)$ the *minsize* of G. On the other hand, the *maxsize* $\rho_+(G) = \max\{\rho_{w,T} : w \in V$ and T is a BFS-tree rooted at $w\}$ of G coincides with its insize $\iota(G)$. Indeed, from the definition, $\rho_+(G) \le \iota(G)$. Conversely, consider a geodesic triangle $\Delta(x, y, w)$ maximizing the insize and suppose, without loss of generality, that $d(x_y, y_x) = \iota(G)$, where x_y and y_x are chosen on the sides of $\Delta(x, y, w)$. Then, if we choose a BFS-tree rooted at w, and such that x_y is an ancestor of x and y_x is an ancestor of y, then one obtains that $\rho_+(G) \ge \iota(G)$. We show in Section 5 that $\rho_+(G) = \tau(G)$ can be computed in polynomial time, and by Proposition 3, it gives a 4-approximation of $\delta(G)$.

On the other hand, the next proposition shows that one cannot get better than a factor 8 approximation of hyperbolicity if instead of computing $\rho_{w,T}$ for an arbitrary BFS tree T rooted at some arbitrary vertex w, we compute the minsize $\rho_-(G)$. Furthermore, we show in Section 5 that we cannot approximate $\rho_-(G)$ with a factor strictly better than 2 unless P = NP. In order to prove the following proposition, we modify slightly the graph H_k of Proposition 9 so that the shortest paths from w to x and y become unique, and then we glue two copies of this modified graph in w.

▶ **Proposition 10.** *For any positive integer k, there exists a graph H_k^* with $\delta(H_k^*) = k + O(1)$ and $\rho_+(H_k^*) \ge \rho_-(H_k^*) \ge 4k - 2$ and a graph G_k^* with $\delta(G_k^*) = 4k$ and $\rho_-(G_k^*) \le 2k$.*

If instead of knowing the distance-matrix D, we only know the distances between the vertices of G up to an additive error k, then we can define a parameter $\widehat{\rho}_{w,T}$ in a similar way as $\rho_{w,T}$ is defined and show that $2\widehat{\rho}_{w,T} + k + 1$ is an 8-approximation of $\delta(G)$ with an additive error of $3k + 1$.

▶ **Proposition 11.** *Given a graph G, a BFS-tree T rooted at a vertex w, and a matrix $\widehat{D}$ such that $d(x, y) \le \widehat{D}(x, y) \le d(x, y) + k$, we can compute in time $O(n^2)$ a value $\widehat{\rho}_{w,T}$ such that $\delta(G) \le 2\widehat{\rho}_{w,T} + k + 1 \le 8\delta(G) + 3k + 1$.*

Interestingly, $\rho_{w,T}$ can also be defined in terms of a distance approximation parameter. Consider a geodesic space X and a GS-tree T rooted at some point w, and let $\rho = \rho_{w,T}$. For a point $x \in X$ and $r \in \mathbb{R}^+$, denote by $x(r)$ the unique point of $[w, x]_T$ at distance r from w if $r < d(w, x)$ and the point x if $r \ge d(w, x)$. For any x, y and $\epsilon \in \mathbb{R}^+$, let $r_{xy}(\epsilon) := \sup\{r : d(x(r'), y(r')) \le \epsilon$ for any $0 \le r' \le r\}$. This supremum is a maximum because the function $r' \mapsto d(x(r'), y(r'))$ is continuous. Observe that by Proposition 6, $\rho = \inf\{\epsilon : r_{xy}(\epsilon) \ge (x|y)_w$ for all $x, y\}$.

Denote by $x_y(\epsilon)$ (respectively, $y_x(\epsilon)$) the point of $[x, w]_T$ (respectively, of $[w, y]_T$) at distance $r_{xy}(\epsilon)$ from w. Let $\widehat{d}_\epsilon(x, y) = d(x, x_y(\epsilon)) + \epsilon + d(y_x(\epsilon), y)$. By the triangle inequality, $d(x, y) \le d(x, x_y(\epsilon)) + d(x_y(\epsilon), y_x(\epsilon)) + d(y_x(\epsilon), y) \le \widehat{d}_\epsilon(x, y)$. Observe that for any ϵ and for any x, y, we have $r_{xy}(\epsilon) \ge (x|y)_w$ if and only if $d(x, x_y(\epsilon)) + d(y_x(\epsilon), y) \le d(x, y)$, i.e., if and only if $d(x, y) \le \widehat{d}_\epsilon(x, y) \le d(x, y) + \epsilon$. Consequently, $\rho = \inf\{\epsilon : d(x, y) \le \widehat{d}_\epsilon(x, y) \le d(x, y) + \epsilon$ for all $x, y\}$.

When we consider a graph G with a BFS-tree T rooted at some vertex w, we have similar results. For a vertex x, we define $x(r)$ as before when r is an integer and for vertices x, y, we define $r_{xy}(\epsilon) := \max\{r \in \mathbb{N} : d(x(r'), y(r')) \le \epsilon$ for any $0 \le r' \le r\}$. Since $\rho = \inf\{\epsilon : r_{xy}(\epsilon) \ge \lfloor(x|y)_w\rfloor$ for all $x, y\}$, we get that $d(x, y) \le \widehat{d}_\rho(x, y) + 1 \le d(x, y) + \rho + 1$.

▶ **Proposition 12.** *If the distance matrix D of a graph G is unknown but the kth power graph G^k of G is given for $k \ge \rho_{w,T}$, then one can approximate the distance matrix D of G in optimal $O(n^2)$ time with an additive term depending only on k.*

Proof. With G^k at hand, for a fixed vertex $x \in X$ the values of $r_{xy}(k)$ and $\widehat{d}_k(x, y)$, for every $y \in X$, can be computed in linear time using a simple traversal of the BFS-tree T. ◄

4.2 Fast approximation of thinness, insize, and slimness

Using Proposition 3 and Theorem 7, we also get the following corollary.

▶ **Corollary 13.** *For a graph G and a BFS-tree T rooted at a vertex w, $\tau(G) \le 8\rho_{w,T} + 4 \le 8\tau(G) + 4$ and $\varsigma(G) \le 6\rho_{w,T} + 3 \le 24\varsigma(G) + 3$. Consequently, an 8-approximation (with additive surplus 4) of the thinness $\tau(G)$ and a 24-approximation (with additive surplus 3) of the slimness $\varsigma(G)$ can be found in $O(n^2)$ time.*

Consider a collection $\mathcal{T} = (T_w)_{w \in V}$ of trees where for each w, T_w is an arbitrary BFS-tree rooted at w, and let $\rho_{\mathcal{T}} = \max_{w \in V} \rho_{w,T_w}$. Since for each w, ρ_{w,T_w} can be computed in $O(n^2)$ time, $\rho_{\mathcal{T}}$ can be computed in $O(n^3)$ time. We stress that for any fixed $w \in V$, $\delta_w(G)$ can be also computed in $O(n^3)$ time. Furthermore, by Proposition 1, $\delta_w(G)$ gives a 2-approximation of the hyperbolicity $\delta(G)$ of G. In what follows, we present similar complexity approximations for $\varsigma(G)$ and $\tau(G)$.

To get a better bound for $\varsigma(G)$, we need to involve one more parameter. Let u and v be arbitrary vertices of G and $T_u \in \mathcal{T}$ be the BFS-tree rooted at u. Let also $(u = u_0, u_1, \ldots, u_\ell = v)$ be the path of T_u joining u with v. Define $\kappa_{T_u}(u, v) := \max\{d(a, u_i) : a \in I(u, v), d(a, u) = i\}$ and $\kappa_{\mathcal{T}} := \max\{\kappa_{T_u}(u, v) : u, v \in V\}$. Note that $\kappa_{\mathcal{T}} \le \kappa(G)$ and that $\kappa_{\mathcal{T}}$ can be computed in $O(n^3)$ time and $O(n^2)$ space. Observe also that for any u, v, $\kappa_{T_u}(u, v) \le \rho_{u,T_u}$ and thus $\kappa_{\mathcal{T}} \le \rho_{\mathcal{T}}$.

▶ **Proposition 14.** *For a graph G and a collection of BFS-trees $\mathcal{T} = (T_w)_{w \in V}$, $\iota(G) = \tau(G) \le \rho_{\mathcal{T}} + 2\kappa_{\mathcal{T}} \le 3\rho_{\mathcal{T}} \le 3\tau(G)$ and $\varsigma(G) \le \rho_{\mathcal{T}} + 2\kappa_{\mathcal{T}} \le 8\varsigma(G)$. Consequently, a 3-approximation of the thinness $\tau(G)$ and an 8-approximation of the slimness $\varsigma(G)$ can be found in $O(n^3)$ time and $O(n^2)$ space.*

Proof. Pick any geodesic triangle $\Delta(x, y, w)$ with sides $[x, y]$, $[x, w]$ and $[y, w]$. Let $[x, w]_T$ and $[y, w]_T$ be the corresponding geodesics of the BFS-tree T for vertex w. Consider the vertices $x_y \in [x, w]_T, y_x \in [w, y]_T$ and vertices $a \in [x, w], b \in [y, w]$ with $d(w, x_y) = d(w, y_x) = d(w, a) = d(w, b) = \lfloor (x|y)_w \rfloor$. We know that $d(x_y, y_x) \le \rho_{\mathcal{T}}$. Since $(x|a)_w = d(a, w)$ and $(y|b)_w = d(b, w)$, $d(a, x_y) \le \kappa_{T_w}(w, x) \le \kappa_{\mathcal{T}}$ and $d(b, y_x) \le \kappa_{T_w}(w, y) \le \kappa_{\mathcal{T}}$. Hence, $d(a, b) \le \rho_{\mathcal{T}} + 2\kappa_{\mathcal{T}}$. Repeating this argument for vertices x and y and their BFS-trees, we get that the insize of $\Delta(x, y, w)$ is at most $\rho_{\mathcal{T}} + 2\kappa_{\mathcal{T}}$. So $\tau(G) \le \rho_{\mathcal{T}} + 2\kappa_{\mathcal{T}} \le 3\rho_{\mathcal{T}} \le \tau(G)$ and by Proposition 3, $\varsigma(G) \le \tau(G) \le \rho_{\mathcal{T}} + 2\kappa_{\mathcal{T}} \le \tau(G) + 2\kappa(G) \le 8\varsigma(G)$. ◄

5 Exact computation

In this section, we provide exact algorithms for computing the slimness $\varsigma(G)$, the thinness $\tau(G)$, and the insize $\iota(G)$ of a given graph G. The algorithm computing $\tau(G) = \iota(G)$ runs in $O(n^2 m)$ time and the algorithm computing $\varsigma(G)$ runs in $\widehat{O}(n^2 m + n^4 / \log^3 n)$ time; both algorithms use $O(n^2)$ space. When the graph is dense (i.e., $m = \Omega(n^2)$), that stays of the same order of magnitude as the best-known algorithms for computing $\delta(G)$ in practice (see [6]), but when the graph is not so dense (i.e., $m = o(n^2)$), our algorithms run in $o(n^4)$ time. In contrast to this result, the existing algorithms for computing $\delta(G)$ exactly are not sensitive to the density of the input. We also show that the minsize $\rho_-(G)$ of a given graph G cannot be approximated with a factor strictly better than 2 unless P = NP. The results of this section are summarized by the following theorem:

▶ **Theorem 15.** *For a graph $G = (V, E)$ with n vertices and m edges, the following holds:*

(1) *the thinness $\tau(G)$ and the insize $\iota(G)$ of G can be computed in $O(n^2 m)$ time;*

(2) *the slimness $\varsigma(G)$ of G can be computed in $\widehat{O}(n^2 m + n^4/\log^3 n)$ time ;*

(3) *Deciding whether the minsize $\rho_-(G)$ of G is at most 1 is NP-complete.*

5.1 Exact computation of thinness and insize

We present an algorithm to compute $\tau(G) = \iota(G)$ that runs in time $O(n^2 m)$. We first compute the distance matrix D of G in time $O(mn)$. To compute $\tau(G)$, we introduce the "pointed thinness" $\tau_w(G)$ of a given vertex w. For a fixed vertex w, let $\tau_w(G) = \max \big\{ d(x', y') : \exists x, y \in V$ such that $x' \in I(w, x), y' \in I(w, y)$, and $d(w, x') = d(w, y') \leq (x|y)_w \big\}$. Observe that for any BFS-tree T rooted at w, we have $\rho_{w,T} \leq \tau_w(G) \leq \tau(G)$, and thus by Corollary 13, $\tau_w(G)$ is an 8-approximation (with additive surplus 4) of $\tau(G)$. Since $\tau(G) = \max_{w \in V} \tau_w(G)$, in order to prove Theorem 15(1), it is sufficient to describe an algorithm computing $\tau_w(G)$ in $O(mn)$. Let $\tau_{w,x}(G) = \max \big\{ d(x', y') : x' \in I(w, x)$ and $\exists y \in V$ such that $y' \in I(w, y)$ and $d(w, x') = d(w, y') \leq (x|y)_w \big\}$ and observe that $\tau_w(G) = \max_{x \in V} \tau_{w,x}(G)$.

For every ordered pair w, x and every vertex z, let $g_z(w, x) = \max \big\{ d(x', z) : x' \in I(w, x)$ and $d(w, x') = d(w, z) \big\}$. Observe that $\tau_{w,x} = \max \big\{ g_z(w, x) : \exists y$ such that $z \in I(w, y)$ and $d(w, z) \leq (x|y)_w \big\}$.

▶ **Lemma 16.** *For any fixed $w, z \in V$, one can compute the values of $g_z(w, x)$ for all $x \in V$ in $O(m)$ time.*

Proof. In order to compute $g_z(w, x)$, we use the following recursive formula: $g_z(w, x) = 0$ if $d(w, x) < d(w, z)$, $g_z(w, x) = d(x, z)$ if $d(w, x) = d(w, z)$, and $g_z(w, x) = \max \big\{ g_z(w, x') : x' \in N(x)$ and $d(w, x') = d(w, x) - 1 \big\}$ otherwise. Given the distance matrix D, for any $x \in V$, we can compute $\{ x' \in N(x) : d(w, x') = d(w, x) - 1 \}$ in $O(\deg(x))$ time. Therefore, using a standard dynamic programming approach, we can compute the values $g_z(w, x)$ for all $x \in V$ in $O(\sum_x \deg(x)) = O(m)$ time. ◀

Let $h_{w,x}(z) = \max \big\{ (x|y)_w : z \in I(w, y) \big\}$ and observe that $\tau_{w,x}(G) = \max \big\{ g_z(w, x) : d(w, z) \leq h_{w,x}(z) \big\}$. Note that if $w, x \in V$ are fixed, then $h_{w,x}(z)$ satisfies the following recursive formula: $h_{w,x}(z) = \max \big\{ (x|z)_w, h'_{w,x}(z) \big\}$ where $h'_{w,x}(z) = \max \big\{ h_{w,x}(z') : z' \in N(z)$ and $d(w, z') = d(w, z) + 1 \big\}$. If we order the vertices of V by non-increasing distance to w, using dynamic programming, we can compute the values of $h_{w,x}(z)$ for all z in $O(\sum_z \deg(z)) = O(m)$ time.

We can thus compute the values $g_z(w, x)$ and $h_{w,x}(z)$ for all $x, z \in V$ in $O(mn)$ time. Then for every fixed w, x, we can compute $\tau_{w,x}(G) = \max \big\{ g_z(w, x) : d(w, z) \leq h_{w,x}(z) \big\}$ in $O(n)$ time, and consequently we can compute $\tau_w(G) = \max_x \tau_{w,x}(G)$ in $O(mn)$ time.

5.2 Exact computation of slimness

To prove Theorem 15(2), we introduce the "pointed slimness" $\varsigma_w(G)$ of a given vertex w. Formally, $\varsigma_w(G)$ is the least integer k such that, in any geodesic triangle $\Delta(x, y, z)$ such that $w \in [x, y]$, we have $d(w, [x, z] \cup [y, z]) \leq k$. Note that $\varsigma_w(G)$ cannot be used to approximate $\varsigma(G)$ (that is in sharp contrast with $\delta_w(G)$ or $\tau_w(G)$). In particular, $\varsigma_w(G) = 0$ whenever w is a pending vertex of G (or, more generally, a simplicial vertex of G). On the other hand, we have $\varsigma(G) = \max_{w \in V} \varsigma_w(G)$. Therefore, in order to prove Theorem 15(2), it is sufficient to describe an algorithm computing $\varsigma_w(G)$ that runs in $\widehat{O}(nm + n^3/\log^3 n)$ time for every

w. For every $x, z \in V$ we set $p_w(x, z)$ to be the least integer k such that, for every geodesic $[x, z]$, we have $d(w, [x, z]) \leq k$. Observe that $\varsigma_w(G) \leq k$ if and only if for all $x, y \in V$ such that $w \in I(x, y)$, and any $z \in V$, $\min\{p_w(x, z), p_w(y, z)\} \leq k$.

As before, we assume that the distance matrix D of G has been already computed. The algorithm for computing $\varsigma_w(G)$ proceeds in two phases. We first compute $p_w(x, z)$ for every $x, z \in V$. Second, we seek for a triple (x, y, z) such that $w \in I(x, y)$ and $\min\{p_w(x, z), p_w(y, z)\}$ is maximized.

For any fixed $w, x \in V$, observe that $p_w(x, z)$ satisfies the following recursive formula: $p_w(x, x) = d(w, x)$ and for any $z \neq x$, $p_w(x, z) = \min\{d(w, z), p'_w(x, z)\}$ where $p'_w(x, z) = \max\{p_w(x, z') : z' \in N(z) \text{ and } d(x, z') = d(x, z) - 1\}$. Using dynamic programming, one can compute the values $p_w(x, z)$ for all $z \in V$ in $O(m)$ time. Consequently, we can compute the values $p_w(y, z)$ for all $y, z \in V$ in $O(mn)$ time, and then, we can compute $\varsigma_w(G)$ in $O(n^3)$ time by enumerating all possible triples $x, y, z \in V$ such that $w \in I(x, y)$ and keeping one maximizing $\min\{p_w(x, z), p_w(y, z)\}$.

We can improve the running time by reducing the problem to TRIANGLE DETECTION as follows. Given a fixed integer k, the graph $\Gamma_\varsigma^w[k]$ has vertex set $V_1 \cup V_2 \cup V_3$, with every set V_i being a copy of $V \setminus \{w\}$. There is an edge between $x_1 \in V_1$ and $y_2 \in V_2$ if and only if the corresponding vertices $x, y \in V$ satisfy $w \in I(x, y)$. Furthermore, there is an edge between $x_1 \in V_1$ and $z_3 \in V_3$ (respectively, between $y_2 \in V_2$ and $z_3 \in V_3$) if and only if we have $p_w(x, z) > k$ (respectively, $p_w(y, z) > k$). It is easy to see that $\varsigma_w(G) \leq k$ if and only if $\Gamma_\varsigma^w[k]$ is triangle-free. Once the distance matrix of G and the values $p_w(y, z)$ for all $y, z \in V$ have been computed, we can construct $\Gamma_\varsigma^w[k]$ in $O(n^2)$ time. Since TRIANGLE DETECTION can be solved in $\widehat{O}(n^3/\log^4 n)$ time [35], we can decide whether $\varsigma_w(G) \leq k$ in the same time, and by performing binary search, we can compute $\varsigma_w(G)$ in $\widehat{O}(n^3/\log^3 n)$ time.

5.3 Approximating the minsize is hard

We now prove Theorem 15(3). Note that if we are given a BFS-tree T rooted at a vertex w, we can easily check whether $\rho_{w,T} \leq 1$, and thus deciding whether $\rho_-(G) \leq 1$ is in NP. In order to prove that this problem is NP-hard, we do a reduction from SAT. Let Φ be a SAT formula with m clauses $c_1, c_2, \ldots, c_m$ and n variables $x_1, x_2, \ldots, x_n$. Let $X = \{x_1, \overline{x}_1, \ldots, x_n, \overline{x}_n\}$. For simplicity, in what follows, we denote $x_i, \overline{x}_i$ by ℓ_{2i-1}, ℓ_{2i}. Let $C = \{c_1, \ldots, c_m\}$ be the clause-set of Φ. Finally, let w and $V = \{v_1, v_2, \ldots, v_{2n}\}$ be additional vertices. We construct a graph G_Φ with $V(G_\Phi) = \{w\} \cup V \cup X \cup C$ and where $E(G_\Phi)$ is defined as follows:

- $N(w) = V$ and V is a clique,
- for every i, i', v_i and $l_{i'}$ are adjacent if and only if $i = i'$;
- for every i, i', ℓ_i and $\ell_{i'}$ are adjacent if and only if $\ell_{i'} \neq \overline{\ell}_i$;
- for every i, j, v_i and c_j are not adjacent;
- for every i, j, ℓ_i and c_j are adjacent if and only if $\ell_i \in c_j$;
- for every j, j', $c_j, c_{j'}$ are adjacent if and only if $c_j, c_{j'}$ intersect in exactly one literal.

We can show that we can preprocess the formula Φ in polynomial time such that:

(1) for every BFS-tree T rooted at $u \neq w$, $\rho_{u,T} \geq 2$,
(2) for any BFS-tree T rooted at w, for any $t, u \in V$, if $d(t_u, u_t) \geq 2$, then $t, u \in C$,
(3) for every $c, c' \in C$, $d(c, c') \leq 2$.

Now, for every $c, c' \in C$, observe that $\lfloor (c|c')_w \rfloor = 2$ and thus $c_{c'}, c'_c \in X$. Consequently, $\rho_{w,T} \leq 1$ if and only if $d(c_{c'}, c'_c) \leq 1$ for all $c, c' \in C$, i.e., if and only if $c_{c'} \neq \overline{c'}_c$ for all $c, c' \in C$. Therefore, there exists a tree T rooted at w such that $\rho_{w,T} = 1$ if and only if there exists a satisfying assignment for Φ.

── **References** ──────────────────────

1   M. Abu-Ata and F.F. Dragan. Metric tree-like structures in real-world networks: an empirical study. *Networks*, 67(1):49–68, 2016.

2   A.B. Adcock, B.D. Sullivan, and M.W. Mahoney. Tree-like structure in large social and information networks. In *ICDM*, pages 1–10. IEEE Computer Society, 2013.

3   J.M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short. Notes on word hyperbolic groups. In E. Ghys, A. Haefliger, and A. Verjovsky, editors, *Group Theory from a Geometrical Viewpoint*, ICTP Trieste 1990, pages 3–63. World Scientific, 1991.

4   A.M. Ben-Amram. The Euler path to static level-ancestors. *CoRR*, abs/0909.1030, 2009.

5   M.A. Bender and M Farach-Colton. The level ancestor problem simplified. *Theor. Comput. Sci.*, 321(1):5–12, 2004.

6   M. Borassi, D. Coudert, P. Crescenzi, and A. Marino. On computing the hyperbolicity of real-world graphs. In *ESA*, volume 9294 of *Lecture Notes in Computer Science*, pages 215–226. Springer, 2015.

7   M. Borassi, P. Crescenzi, and M. Habib. Into the square: On the complexity of some quadratic-time solvable problems. *Electr. Notes Theor. Comput. Sci.*, 322:51–67, 2016.

8   B.H. Bowditch. Notes on Gromov's hyperbolicity criterion for path-metric spaces. In E. Ghys, A. Haefliger, and A. Verjovsky, editors, *Group Theory from a Geometrical Viewpoint*, ICTP Trieste 1990, pages 64–167. World Scientific, 1991.

9   M.R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1999.

10 J. Chalopin, V. Chepoi, F.F. Dragan, G. Ducoffe, A. Mohammed, and Y. Vaxès. Fast approximation and exact computation of negative curvature parameters of graphs. *CoRR*, abs/1803.06324, 2018.

11 J. Chalopin, V. Chepoi, P. Papasoglu, and T. Pecatte. Cop and robber game and hyperbolicity. *SIAM J. Discrete Math.*, 28(4):1987–2007, 2014.

12 V. Chepoi, F.F. Dragan, B. Estellon, M. Habib, and Y. Vaxès. Diameters, centers, and approximating trees of delta-hyperbolicgeodesic spaces and graphs. In *Symposium on Computational Geometry*, pages 59–68. ACM, 2008.

13 V. Chepoi, F.F. Dragan, B. Estellon, M. Habib, Y. Vaxès, and Yang Xiang. Additive spanners and distance and routing labeling schemes for hyperbolic graphs. *Algorithmica*, 62(3-4):713–732, 2012.

14 V. Chepoi, F.F. Dragan, and Y. Vaxès. Core congestion is inherent in hyperbolic networks. In *SODA*, pages 2264–2279. SIAM, 2017.

15 V. Chepoi and B. Estellon. Packing and covering δ-hyperbolic spaces by balls. In *APPROX-RANDOM*, volume 4627 of *Lecture Notes in Computer Science*, pages 59–73. Springer, 2007.

16 N. Cohen, D. Coudert, and A. Lancin. On computing the Gromov hyperbolicity. *ACM Journal of Experimental Algorithmics*, 20:1.6:1–1.6:18, 2015.

17 D. Coudert and G. Ducoffe. Recognition of C_4-free and 1/2-hyperbolic graphs. *SIAM J. Discrete Math.*, 28(3):1601–1617, 2014.

18 D. Coudert, G. Ducoffe, and A. Popa. Fully polynomial FPT algorithms for some classes of bounded clique-width graphs. In *SODA*, pages 2765–2784. SIAM, 2018.

19 B. DasGupta, M. Karpinski, N. Mobasheri, and F. Yahyanejad. Effect of Gromov-hyperbolicity parameter on cuts and expansions in graphs and some algorithmic implications. *Algorithmica*, 80(2):772–800, 2018.

20 T. Delzant and M. Gromov. Courbure mésoscopique et théorie de la toute petite simplification. *J. Topol.*, 1:804–836, 2008.

21 R. Duan. Approximation algorithms for the Gromov hyperbolicity of discrete metric spaces. In *LATIN*, volume 8392 of *Lecture Notes in Computer Science*, pages 285–293. Springer, 2014.

22 K. Edwards, W.S. Kennedy, and I. Saniee. Fast approximation algorithms for p-centers in large δ-hyperbolic graphs. In *WAW*, volume 10088 of *Lecture Notes in Computer Science*, pages 60–73, 2016.

23 T. Fluschnik, C. Komusiewicz, G.B. Mertzios, A. Nichterlein, R. Niedermeier, and N. Talmon. When can graph hyperbolicity be computed in linear time? In *WADS*, volume 10389 of *Lecture Notes in Computer Science*, pages 397–408. Springer, 2017.

24 H. Fournier, A. Ismail, and A. Vigneron. Computing the Gromov hyperbolicity of a discrete metric space. *Inf. Process. Lett.*, 115(6-8):576–579, 2015.

25 E. Ghys and P. de la Harpe (eds). *Les groupes hyperboliques d'après M. Gromov*, volume 83 of *Progress in Mathematics*. Birkhäuser, 1990.

26 M. Gromov. Hyperbolic groups. In S. Gersten, editor, *Essays in Group Theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

27 M.F. Hagen. Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topology*, 7(2):385–418, 2014.

28 W.S. Kennedy, I. Saniee, and O. Narayan. On the hyperbolicity of large-scale networks and its estimation. In *BigData*, pages 3344–3351. IEEE, 2016.

29 O. Narayan and I. Saniee. Large-scale curvature of networks. *Phys. Rev. E*, 84:066108, 2011.

30 P. Papasoglu. An algorithm detecting hyperbolicity. In *Geometric and computational perspectives on infinite groups (Minneapolis, MN and New Brunswick, NJ, 1994)*, volume 25 of *DIMACS - Series in Discrete Mathematics and Theoretical Computer Science*, pages 193–200. 1996.

31 N. Polat. On infinite bridged graphs and strongly dismantlable graphs. *Discrete Mathematics*, 211:153–166, 2000.

32 Y. Shavitt and T. Tankel. Hyperbolic embedding of internet graph for distance estimation and overlay construction. *IEEE/ACM Trans. Netw.*, 16(1):25–36, 2008.

33 M. Soto. *Quelques propriétés topologiques des graphes et applications à Internet et aux réseaux*. PhD thesis, Université Paris Diderot, 2011.

34 K. Verbeek and S. Suri. Metric embedding, hyperbolic space, and social networks. In *Symposium on Computational Geometry*, pages 501–510. ACM, 2014.

35 H. Yu. An improved combinatorial algorithm for boolean matrix multiplication. In *ICALP(1)*, volume 9134 of *Lecture Notes in Computer Science*, pages 1094–1105. Springer, 2015.