

Tree-like structure in large social and information networks

Aaron B. Adcock
Department of Electrical Engineering
Stanford University
Stanford, CA 94305
Email: aadcock@stanford.edu

Blair D. Sullivan
Department of Computer Science
North Carolina State University
Raleigh, NC 27695
Email: blair_sullivan@ncsu.edu

Michael W. Mahoney
Department of Mathematics
Stanford University
Stanford, CA 94305
Email: mmahoney@cs.stanford.edu

Abstract—Although large social and information networks are often thought of as having hierarchical or tree-like structure, this assumption is rarely tested. We have performed a detailed empirical analysis of the tree-like properties of realistic informatics graphs using two very different notions of tree-likeness: Gromov’s δ -hyperbolicity, which is a notion from geometric group theory that measures how tree-like a graph is in terms of its metric structure; and tree decompositions, tools from structural graph theory which measure how tree-like a graph is in terms of its cut structure. Although realistic informatics graphs often do not have meaningful tree-like structure when viewed with respect to the simplest and most popular metrics, *e.g.*, the value of δ or the treewidth, we conclude that many such graphs do have meaningful tree-like structure when viewed with respect to more refined metrics, *e.g.*, a size-resolved notion of δ or a closer analysis of the tree decompositions. We also show that, although these two rigorous notions of tree-likeness capture very different tree-like structures in worst-case, for realistic informatics graphs they empirically identify surprisingly similar structure. We interpret this tree-like structure in terms of the recently-characterized “nested core-periphery” property of large informatics graphs; and we show that the fast and scalable k -core heuristic can be used to identify this tree-like structure.

I. INTRODUCTION

Although large informatics graphs such as social and information networks are often thought of as having hierarchical or tree-like structure, this assumption is rarely tested, and it has proven difficult to exploit this idea in practice. Moreover, given recent work demonstrating that large informatics graphs have properties that are very different than small social networks and graphs that arise in other machine learning and data analysis applications (*e.g.*, they have strong expander-like properties at large size scales), it is not clear whether such structure can be exploited for improved graph mining and machine learning, even assuming it exists. Here, we describe the results of a detailed empirical analysis of the tree-like structure in a wide range of small to moderately-large social and information networks. We consider two traditional notions of the manner in which a graph or network can be tree-like: Gromov’s δ -hyperbolicity, a notion from geometric group theory that measures how tree-like a graph is in terms of its metric or distance structure; and tree decompositions, tools from structural graph theory that measure how tree-like a graph is in terms of its cut or partitioning structure.

By both of these measures, we find that, when compared with a suite of idealized graphs representing low-dimensional

structures, constant-degree expanders, popular generative models, etc., realistic social and information networks do have meaningful large-scale tree-like structure; but we also find that characterizing this structure with just these traditional tools is nontrivial and quite brittle, suggesting that exploiting this tree-like structure for common machine learning and data analysis tasks will be challenging. Relatedly, this tree-like structure is not evident when considering the simplest and most popular hyperbolicity and tree decomposition metrics, *e.g.*, the value of δ or the treewidth; but it is evident when considering more refined metrics, *e.g.*, a size-resolved notion of δ or a closer analysis of the tree decomposition.

Interestingly, although δ -hyperbolicity and tree decompositions attempt to capture very different ways in which a graph can be tree-like, we observe that, for most realistic informatics graphs, they empirically identify surprisingly similar structure; and that this structure is consistent with the recently-characterized “nested core-periphery” structure that is common to many large social and information networks [30]. To understand this phenomenon in greater detail, we used the method of k -core decompositions, which is a fast heuristic that can be applied to networks with millions (or more) of nodes. Interpreting the output of our k -core computations in light of the tree-like structure we identified with the more theoretically-principled δ -hyperbolicity and tree decomposition computations provides a much more detailed and principled way to describe the structural properties of larger networks.

II. BACKGROUND AND RELATED WORK

A. Preliminaries on Graph Theory and Metric Spaces

Let $G = (V, E)$ be a graph with vertex set V and edge set $E \subseteq V \times V$. We will model social and information networks by undirected graphs. The *degree* of a vertex v , denoted $d(v)$, is defined as the number of vertices that are adjacent to v . The average degree is denoted \bar{d} . A graph is called *connected* if there exists a path between any two vertices. A graph is called a *tree* if it is connected and has no cycles. A vertex in a tree is called a *leaf* if it has degree 1. Unless otherwise specified, our analyses will always consider the largest connected subgraph of G , the *giant component*.

A graph G naturally defines a metric l on its vertex set V , where the distance between two vertices is the length of a shortest path between the vertices. The geodesics are shortest paths between vertices and are generally not unique. We will use the term *quadruplet* to refer to a set of four nodes. As

any two nodes in a network can have multiple geodesic paths between them, a quadruplet may be associated with many different sets of geodesics. To refer to a choice of one geodesic between each pair of nodes in a quadruplet, we will use the term *quadruped*.

B. Metric Tree-likeness: Gromov’s δ -hyperbolicity

The notion of δ -hyperbolicity was introduced by Gromov in the context of geometric group theory and the geometry of negatively curved metric spaces [12], [14], [21]. Informally, it may be thought of as characterizing how tree-like a metric space is in terms of its distance or metric structure [21]. There are several ways of characterizing δ -hyperbolic metric spaces, all of which are equivalent up to constant factors [12], [14], [21]. Here, we present the *Gromov 4-point δ -hyperbolicity*.

Definition 1: Let $0 \leq \delta < \infty$. (X, ℓ) is called *4-point δ -hyperbolic* if and only if for all $x, y, u, v \in X$, ordered such that $\ell(x, y) + \ell(u, v) \geq \ell(x, u) + \ell(y, v) \geq \ell(x, v) + \ell(y, u)$, the following condition holds:

$$(\ell(x, y) + \ell(u, v)) - (\ell(x, u) + \ell(y, v)) \leq 2\delta. \quad (1)$$

In particular, the δ of a graph G is the minimum value of δ such that G is δ -hyperbolic. More generally, a space is non-hyperbolic if no finite δ exists such that the 4-point condition is true. For integral distance functions, δ is valued in multiples of $1/2$ for the 4-point condition. Intuitively, δ provides a bound on the maximum distance between any points on two different sides of a geodesic triangle, causing the sides of a triangle to “bend-in” towards the center of the triangle as it becomes large. This is in contrast to Euclidean space where the distance between the sides of a triangle becomes arbitrarily large as the triangle grows.

We have referred to δ -hyperbolicity as a metric notion of tree-likeness. Recall that n -point δ -hyperbolic spaces can be approximated by a tree metric with an additive error of $O(\delta \log_2 n)$ [21], and consider several additional examples. The first example is, as expected, that trees are 0-hyperbolic. The second, perhaps unexpected, example is that cliques are also 0-hyperbolic. It may seem strange that a clique would be considered tree-like, but in the sense of metric structure, it is. Another example is the cycle graph with n vertices. For large n , it is approximately $\frac{n}{4}$ -hyperbolic and it can be shown that this is the largest possible δ -hyperbolicity on a finite graph of n vertices. Finally, the Poincaré half-space is $\log_2 3$ -hyperbolic and discretizations of the hyperbolic plane (including d -regular trees or the ring trees in [13]) have constant hyperbolicity. This is in contrast to a square planar grid, where the hyperbolicity grows with the size of grid considered.

C. Cut-based Tree-likeness: Tree Decompositions

Originally associated with the proof of the Graph Minors Theorem [35], tree decompositions find applications in graph theory, dynamic programming, and linear algebra [22].

Definition 2: A *tree decomposition* of $G = (V, E)$ is a pair $(\{X_i : i \in I\}, T = (I, F))$, with each $X_i \subseteq V$, and T a tree with the following properties: (1) $\cup_{i \in I} X_i = V$; (2) For all $(v, w) \in E, \exists i \in I$ with $v, w \in X_i$; (3) For all $v \in V, \{i \in I : v \in X_i\}$ forms a connected subtree of T .

The X_i are the *bags* of the tree decomposition. The *width* of a tree decomposition is $\max_{i \in I} |X_i| - 1$ and the *treewidth* of G , denoted $tw(G)$ is the minimum width over all tree decompositions of G . The *length* of a tree decomposition is $\max_{i \in I, x, y \in X_i} \ell(x, y)$. The *treelength* of G , denoted $tl(G)$, is the minimum length over all tree decompositions of G .

Finding optimal tree decompositions, e.g., optimizing treewidth, is \mathcal{NP} -hard [37], [22]. Although several approximation algorithms exist [11], [7], it is known that these are not practical [36], and thus in practice one of several widely-used heuristics can be used. For example, the so-called MINDEGREE heuristic [31] and the AMD heuristic [6] are based on making greedy decisions based on node degrees; and nested dissection, or the ND heuristic [19], [20], works by repeatedly finding a balanced separator for the network. We used two of these heuristics (and others which do not scale well), relying on METIS [25] to find the separators; and we have also used several methods for computing lower bounds on the treewidth. We were able to run the ND heuristic on most of our networks, but it does not complete for some Facebook networks.

We emphasize that the tree-like properties captured by a tree decomposition of low width are generally very different than those in δ -hyperbolic spaces. For example, removing the vertices in X_i (or the vertices in $X_i \cap X_j$) from G separates the graph G into two or more disconnected pieces [17]. A tree decomposition thus captures the ways in which the cut-structure of a network behaves like a tree. By this definition, trees have the minimum possible treewidth of 1, but (in contrast to δ -hyperbolicity) a clique is now the least tree-like graph with a treewidth of $n - 1$. In fact, the only valid tree decompositions of a clique have all vertices in a single bag. On the other hand, a cycle graph of size n , has a treewidth of 2. A square, $n \times n$ planar grid has treewidth n and is not tree-like by this measure.

D. k -Core Decompositions

A k -core decomposition provides a way to decompose a graph that, while less explicitly tree-like than δ -hyperbolicity or tree decompositions, is faster for much larger graphs.

Definition 3: Given a set $C \subseteq V$, $G[C]$ is the k -core of G if and only if $G[C]$ is the maximal induced subgraph with the property that for all $v \in C$, the degree $d_{G[C]}(v) \geq k$.

Let G_k denote the k -core of G . We note that the cores G_k are nested, i.e. if $i \leq j$ then $G_j \subseteq G_i$. The maximum k -core of G is the largest k such that $G_k \neq \emptyset$. A node v has *core number* k if v belongs to the k -core, but not the $(k + 1)$ -core. The k -shell of G , denoted S_k , is the set of all nodes with core number k . The k -periphery of a graph is all nodes and edges in a network that are not contained in the k -core. In Section VI, we use an algorithm from [9], [10] that calculates the k -core decomposition in $O(|V| + |E|)$ time.

E. Additional Related Work

Recent work has focused on using ideas related to δ -hyperbolicity for the study of large graphs [38], [27], [29]. This complements results demonstrating that many real-world networks have a highly connected core with “whiskers” or “tendrils” that are connected by short paths through the core [30], [38]. An important practical point is that the usual

definition of hyperbolicity was developed for continuous and infinite metric spaces and leads to trivial results when applied to discrete, finite metric spaces (since the maximum length of the side of any triangle or quadruplet is bounded by the diameter, δ is always finite for a finite graph). Methods to deal with this on finite graphs include [13], [34], [24], [23].

Tree decompositions gained popularity in the theoretical computer science community due to linear time algorithms for \mathcal{NP} -hard optimization problems for graphs of bounded treewidth using dynamic programming [8] and have been used for inference in probabilistic graphical models [28]. One previous study considers heuristics for the hyperbolicity/treewidths of autonomous systems and internet router networks, suggesting that treewidth is large for these networks [33].

Many methods have used k -core decomposition ideas to analyze real networks. For example, the core number has been shown to be a better measure of importance than either node degree or betweenness centrality for some applications [5], [26], and it has the advantage that it is relatively easy to compute [10]. It has also been used in a variety of visualization algorithms for large networks [4], [5], [16].

III. NETWORK DATASETS

We have examined a broad set of real-world social and information networks as well as a wide range of synthetic graphs chosen to illustrate various extreme cases for δ -hyperbolicity and tree decompositions. See Table I for a summary. Since exactly computing the δ -hyperbolicity of a graph is extremely computationally intensive (see [3] for details on these challenges), many of the synthetic networks are rather small, *e.g.*, consisting of fewer than 10,000 nodes. For the real-world networks, we considered both small networks of roughly the same size as well as much larger networks (for which we could perform tree decompositions and/or k -core computations).

Erdős-Rényi (ER) graphs. Although often viewed as a straw-man, this model has several interesting properties for our tree-like analysis when *extremely* sparse values of the connection probability are considered. In the extremely sparse regime of $1/n < p < \log(n)/n$, (w.h.p.) ER graphs are not fully-connected, they have nodes with degree roughly $\log(n)$, and they have small clusters of size roughly $\log(n)$ that are connected to the rest of the graph by one or a small constant number of edges [18], [30]. Here, we set the target number of nodes to $n = 5000$, and we choose $p = \frac{d}{n}$ for various values of d from 1.6 to 32.0. We denote these networks using ER(d).

Power Law (PL) graphs. Here, we consider the Chung-Lu model [15], an ER-like random graph model parameterized to have a power law degree distribution (in expectation) with power law parameter γ , which we vary between 2 and 3.

SNAP Networks. We selected several representative networks from the SNAP website [1] that were used in the large-scale evaluation of clustering and community structure in large networks [30]. These networks represent a range of application domains of social and information network analysis, and include autonomous system snapshots (AS20000102, OREGON1), collaboration networks (CA-GRQC, CA-ASTROPH), an email network from Enron (EMAILENRON), a Gnutella peer-to-peer network (GNUTELLA09), and a Stanford Web network (STANFORDWEB).

Facebook Networks. We selected several representative Facebook graphs out of the 100 Facebook graphs from various American universities [39]. These range in size from ca. 700 nodes (CALTECHFB) to ca. 30,000 nodes (TEXAS84FB, and others). Since these networks can be thought of as being multiple snapshots of graphs constructed with similar generative processes, in general these networks exhibit very similar properties. One notable characteristic that is evident from Table I is that these Facebook networks have a much higher average degree than the networks drawn, *e.g.*, from the SNAP website, that were considered previously [30].

Miscellaneous Networks. Finally, we also selected several miscellaneous real and synthetic networks, designed to illustrate several points and to serve as “controls” on our empirical evaluations: a network of the the Western United States Power Grid (POWERGRID) [40]; a network of US political bloggers (POLBLOGS) [2]; a discretization of part of the 2-dimensional lattice \mathbb{Z}^2 (PLANARGRID); and a random geometric graph (RANDGRID) from the two-dimensional Euclidean plane.

IV. METRIC TREE-LIKENESS: δ -HYPERBOLICITY

In this section, we will summarize our main empirical results on the hyperbolic properties of a suite of synthetic and realistic networks, with an eye on characterizing how real networks are tree-like in terms of their metric structure.

A. Main Empirical Results for δ -hyperbolicity

Our main empirical results consist of the computation of the δ -hyperbolicity on a suite of idealized graphs (representing low-dimensional structures, constant-degree expanders, popular generative models, etc.) and real-world networks (drawn from several domains of social and information network analysis). As mentioned in Section II-E, the maximum value of δ in a finite graph can be difficult to interpret. A finite graph has a finite δ , limited by the diameter of the graph, and results not presented show that there are a very small number of quadruplets which achieve the maximum δ . Thus, we consider a “scaled” and “size-resolved” version of δ (from which the original δ can be recovered). By scaled, we mean that we divide by the length of the longest side of the quadruplet defined by vertices x, y, u, v ; and by size-resolved, we mean that we look at this as a function of the longest side of the quadruplet.

In more detail, consider the δ value of a given quadruplet: by this, we mean the value on the left side of the inequality in Equation (1), for a given set of nodes x, y, u, v . We will write $\delta(\square_{xyuv})$ to denote this quantity for the quadruplet \square_{xyuv} with vertices x, y, u, v . We define δ_D to be:

$$\delta_D = \max_{\square_{xyuv} \in Q_D} \frac{\delta(\square_{xyuv})}{D},$$

where Q_D is the set of quadruplets with maximum pairwise distance D (the “diameter” of the quadruplet) [34], [24]. This allows us to look at how δ behaves over different distance scales in the network. Then, we will consider plots of δ_D against D , which allows us to see the relative hyperbolicity of the networks at each size scale present in the network.

See Figure 1 for our main results on δ_D as a function of d , a summary of which is presented in Table II. In Figure 1,

Network	n	e	n_c	e_c	\bar{d}	\bar{C}	D	\bar{D}	Description
ER Random Graphs									
ER(1.6)	5000	3996	3210	3471	2.16	0.00	38	15.8	ER graph with $p = 1.6/n$
ER(1.8)	5000	4486	3617	4118	2.28	9.30×10^{-4}	34	12.7	ER graph with $p = 1.8/n$
ER(2)	5000	4986	4001	4783	2.39	9.11×10^{-4}	30	11.9	ER graph with $p = 2/n$
ER(4)	5000	9881	4879	9878	4.05	8.96×10^{-3}	15	6.80	ER graph with $p = 4/n$
ER(8)	5000	20102	4998	20102	8.04	1.59×10^{-3}	7	4.81	ER graph with $p = 8/n$
ER(16)	5000	40215	5000	40215	16.1	3.13×10^{-3}	5	3.86	ER graph with $p = 16/n$
ER(32)	5000	80258	5000	80258	32.1	6.39×10^{-3}	4	3.05	ER graph with $p = 32/n$
PL Random Graphs									
PL(2.25)	5000	5790	3393	5634	3.32	.0131	16	5.51	PL graph with $\gamma = 2.25$
PL(2.50)	5000	7238	4895	6802	2.78	2.46×10^{-3}	18	6.65	PL graph with $\gamma = 2.50$
PL(2.75)	5000	6236	4650	5641	2.43	6.99×10^{-4}	22	8.20	PL graph with $\gamma = 2.75$
PL(3.00)	5000	5363	4071	4556	2.24	1.18×10^{-3}	29	10.1	PL graph with $\gamma = 5.00$
SNAP Graphs									
AS20000102	6474	12572	6474	12572	3.88	.399	9	4.34	Snapshot of autonomous systems network
CA-GRQC	5241	14484	4158	13422	6.46	.665	17	6.74	Collaboration network, general relativity
CA-ASTROPH	18771	198050	17903	196972	22.0	.669	14	4.77	Collaboration network, astrophysics
GNUTELLA09	8114	26013	8104	26008	6.42	.0137	10	5.22	Peer-to-peer filesharing network
EMAILENRON	36692	183831	33696	180811	10.7	.708	13	4.72	E-mail network of Enron
OREGON1	11174	23409	11174	23409	4.19	.453	10	4.28	AS peering information
FB Graphs									
LEHIGHFB	5075	198347	5073	198346	78.2	.270	6	3.19	Facebook friend network from Lehigh
VANDERBILTFB	8096	427832	8063	427829	106	.255	7	3.18	Facebook friend network from Vanderbilt
STANFORDFB	11621	568330	11586	568309	98.1	.252	9	3.35	Facebook friend network from Stanford
Miscellaneous Graphs									
POWERGRID	4941	6594	4941	6594	2.67	.107	46	24.2	Western US power grid
POLBLOGS	1224	16715	1222	16714	27.4	.360	8	3.43	Political blogs network
PLANARGRID	2500	4900	2500	4900	3.92	0.00	98	73.0	50-by-50 planar grid
RANDGRID(3)	2500	3808	114	205	3.60	.510	34	21.4	Random planar graph, average degree 3
RANDGRID(7)	2500	8679	2480	8656	6.98	.596	68	55.7	Random planar graph, average degree 7

TABLE I: Statistics of analyzed networks: number of nodes n ; number of edges e ; nodes in giant component n_c ; edges in giant component e_c ; average degree $\bar{d} = 2E/N$; average clustering coefficient \bar{C} ; diameter D ; and average path length \bar{D} .

each subplot shows a different class of network: ER graphs with different connection probabilities p ; PL graphs with different heavy-tailed parameters γ ; real and synthetic graphs that are low-dimensional or planar; and several realistic social and information graphs. In all subfigures, the vertical axis corresponds to δ_D , where in all cases the maximum achievable δ_D is 0.5; and the horizontal axis shows the quadruplet diameter d . Note that, because of the variations in diameter among the networks shown, the scale is not the same for each plot. The planar networks, in particular, have significantly larger diameters than the other networks investigated and thus a significantly expanded horizontal axis. The main feature of these plots in which we will be interested is the overall increase or decrease in δ as quadruplet diameter increases. Although they can be confusing, the “multiple-of-two variations,” *i.e.*, the fluctuations in δ values for even and odd values of quadruplet diameter, are not directly relevant for understanding the large-scale hyperbolicity properties of these graphs; they will be discussed in the long version of this paper.

The plots for the ER graphs indicate that there are two regimes (each with very different tree-like properties) of interest: one where the graphs are dense enough to be fully connected (even though they are still quite sparse); and one where the graphs are so extremely sparse that they are not even fully connected (in which case we are considering the properties of the giant component). In the former case, the diameter is extremely small, and (aside from the multiple-of-two effect mentioned above) the maximal values of δ_D of 0.5

are obtained for all values of D (indicating that the networks are *not* meaningfully hyperbolic). Interestingly, though, when the connection probability is decreased, the diameter of the graph increases, and the plots exhibit a decreasing δ_D after a certain size scale, indicating the presence of nontrivial large-scale hyperbolic structure. The phenomenon responsible for this is a type of core-periphery structure, where a quadruplet has “legs” in the periphery (or whiskers [30]), and where the “high degree nodes” (where for the extremely sparse ER graphs, this means only logarithmically, or slightly, larger than the expected degree due to random fluctuations and lack of measure concentration) provide a scaffolding in the core for short paths to go through. The same phenomenon is seen in the PL graphs; and, in these cases, the large-scale hyperbolicity properties also seem to be due to the lack of measure concentration caused by the degree variability.

A second example of nonhyperbolic networks is seen in PLANARGRID. In this case, the δ_D increases to 0.5 and (aside from the multiple-of-two fluctuations) remains flat. Interestingly, although the POWERGRID “looks like” a two-dimensional mesh in terms of its clustering and isoperimetric structure [30], the hyperbolicity plots indicate that in terms of its metric structure, the variability induced by considering either a RANDGRID graph or the POWERGRID is sufficient to make the graphs have substantial large-scale tree-like structure when viewed in terms of δ_D , though this size-scale is much larger than the other real networks studied.

In light of these results, the real networks do have prop-

Network	D	D_{drop}	δ
Erdős-Rényi			
ER(1.6)	38	21	10
ER(1.8)	34	18	8.0
ER(2)	30	16	7.5
ER(4)	15	9	4.0
ER(8)	7	7	3.0
ER(16)	5	5	2.0
ER(32)	4	4	2.0
Chung-Lu			
PL(2.25)	16	8	3.5
PL(2.50)	18	9	4.0
PL(2.75)	22	11	5.0
PL(3.00)	29	13	6
Planar Networks			
PLANARGRID	98	98	49
RANDGRID(5)	124	40	20
RANDGRID(7)	68	20	18
POWERGRID	46	19	10
Real Networks			
AS20000102	9	6	2.5
CA-GRQC	17	8	3.5
GNUTELLA09	10	7	3.0
POLBLOGS	8	5	2.0
VANDERBILTFB	7	5	2.0

TABLE II: Summary of $\max \delta_D$ Results.

erties consistent with tree-like metric structure. The VANDERBILTFB network is denser and exhibits only a slight decay in δ_D before the diameter of the network is reached, thus indicating slightly hyperbolic properties; but the much sparser CA-GRQC (as well as to a lesser extent other networks such as GNUTELLA09) exhibits a much broader range of decreasing δ_D values. On the whole, these results suggest that the real social and information graphs are nontrivially tree-like in terms of their metric structure when compared with the expander graphs (*e.g.*, moderately dense ER) and low-dimensional graphs (*e.g.*, the PLANARGRID), both of which are well-known to be non-tree-like.

The tendency for δ to begin to decrease as the diameter of the quadruplets increases has been noted in previous work [34]. The drop-off, which happens starting at a size-scale denoted D_{drop} in Table II, indicates that relatively large structures are hyperbolic while there are non-hyperbolic structures at small size scales. This is analogous to the finite- δ definition of hyperbolicity used in infinite spaces [21].

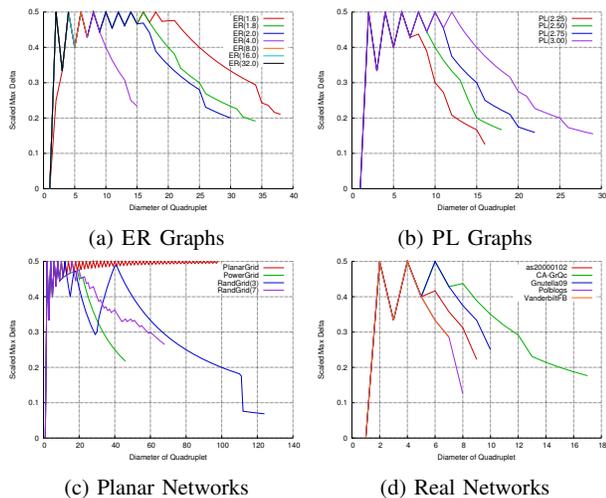


Fig. 1: Max δ_D vs Diameter of Quadruplet.

B. Interpretation of Empirical Results for δ -hyperbolicity

In order to understand better the metric structure of the network, we sampled quadruplets (five from each δ value, and as well as some additional large diameter quadruplets) as we

computed δ , and we kept all of the geodesic paths associated with those quadruplets. See Figures 2 and 3 for representative examples. The examples shown are quadruplets with diameters in the range of the hyperbolic drop-off D_{drop} and relatively large δ values. In all cases, the four square nodes in the figure (which are not visible in some of the denser networks, *e.g.*, VANDERBILTFB), are the four nodes defining the quadruplet. The coloring is based on the relative k -core of the nodes within in the quadruplet (red is high, blue is low) and the layout was generated using the GraphViz command `neato`.

Figure 2 shows all of the geodesics paths between the four nodes in several synthetic networks. Similarities are seen between the very sparse ER(1.6) network, the low-heterogeneity PL(2.75), and POWERGRID, and all sparse networks with small amounts of degree heterogeneity. All of these networks show very little variation in the k -core among the quadruplets, though the peripheral nodes are often in the (slightly) lower cores. This looks quite different from PLANARGRID, where all quadruplets show a very regular structure. The dense ER(16) generates quadruplets with a very flat core structure (most nodes are in the deep core of the network). PL(2.25) shows the most variation in core structure, and the quadruplets sampled suggest that the deeper core nodes make up the central portions of the geodesics in most quadruplets.

In the real networks, we see the (expected) nonzero clustering coefficient and deeper core nodes towards the center of the quadruplets, although for CA-GRQC there are also short paths that don't include deep core nodes. In the Facebook networks, if the diameter of the quadruplet is more than 2, the uninformative visualization we see for VANDERBILTFB is typical. If the diameter is increased to 5 or 6, a similar picture is obtained, except one of the nodes in the quadruplet is connected to the central cluster by a path of length 2 or 3.

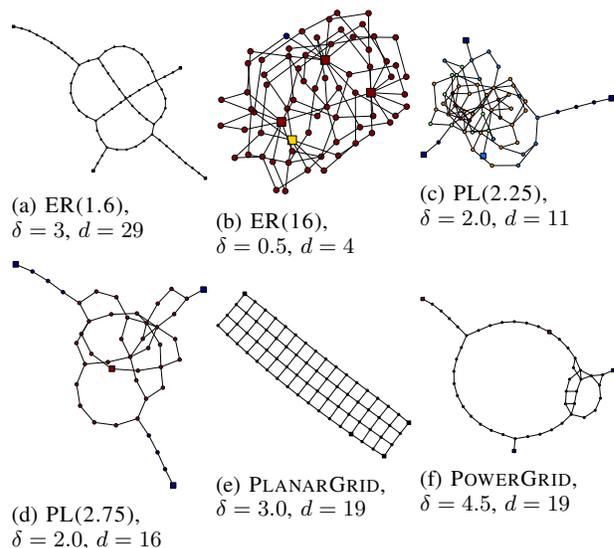


Fig. 2: Selected quadruplets from synthetic networks with all geodesics shown.

We chose quadruplets with large δ values and relatively large diameters because they show a pattern which explains the hyperbolicity drop-off with size-scale seen in Figure 1.

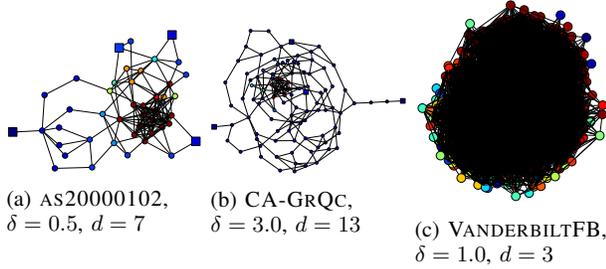


Fig. 3: Selected quadruplets from real networks with all geodesics shown.

Using the 4-point definition of hyperbolicity, it is easy to show that, given a quadruplet and one of the associated quadrupeds (a single choice of geodesics between each of the six pairs of nodes), if there is a section of the geodesics where all six geodesics overlap, then this portion of the network does not contribute to the hyperbolicity. In particular, if all the geodesics overlap, then the length of that section will show up in both the positive and negative terms of definition and will cancel out. Thus, the peripheral “legs” of the quadrupeds, which are particularly prominent in the synthetic networks, do not contribute to the hyperbolicity of the quadruplet. In the real networks (particularly in the Facebook networks), this is much easier to see if a single quadruplet is visualized for a quadruplet, because the increased clustering and larger average degrees results in many paths obscuring the path-like nature of individual quadrupeds.

Clearly, below a certain size scale, there exist non-hyperbolic structures in the core of the network. As the diameter of the quadruplets becomes larger than the diameter of the network core, the geodesic paths converge and form long paths which do not change the hyperbolicity, but they *do* change the scaling factor for the quadruplet, reducing δ_D . Thus we see a drop-off in δ_D as quadruplets grow beyond the core of the network. In other words, the maximum hyperbolicity of these networks is driven by large non-hyperbolic structures in the core. As the quadruplet size grows, these structures eventually stop growing and most quadruplets consist of the smaller non-hyperbolic structures, likely cyclical structures, attached to long legs extending out in the periphery. Note that there could be paths between these legs, but *not* shortest paths. This will become important in Section VI, where we provide additional evidence of this structure and its interactions with the k -core decomposition.

V. CUT-BASED TREE-LIKENESS: TREE DECOMPOSITIONS

In this section, we describe our results illustrating the combinatorial tree-like structures present in many real networks.

A. Main Empirical Results for Tree Decompositions

Our main empirical results consist of using the AMD and ND heuristics with Gavril’s algorithm to compute tree decompositions on our suite of idealized graphs. In addition to reporting the width (W) and the length (L) of the decompositions computed using these heuristics, we also compute several other statistics: the number of bags in the tree decomposition, N_{bags} ; and the diameter of the tree (we call this E , for eccentricity

of the tree, and this shows how far out the tree’s branches extend from its center/bicenter). An additional quantity of interest is how many branches the tree decomposition has. We measure the *branch ratio* of the network with the statistic: $B = \frac{N_{leaves}-2}{N_{bags}-3}$. A summary of our results is presented in Table III. A subscript of a indicates the AMD heuristic, and a subscript of n indicates the ND heuristic. A bar indicates an averaged statistic.

Consider, first, the ER networks, where we see several trends as the parameter p is varied. As expected, the width and average width of the tree decomposition grow with the density of the network indicating the increasingly large number of vertices needed to produce a cut in the network. In the densest of the ER networks, ER(32.0), 80% of the network is in a single, central bag and the length of this bag is equivalent to the diameter. In addition, the number of bags and the diameter of the tree decomposition is smaller in this network. Recall that this is similar to the tree decomposition of a clique (which would contain all nodes in a single bag). In contrast, and importantly, the sparsest ER networks exhibit a more tree-like structure. They have tree decompositions with much smaller widths, longer branches, and lengths that are small compared to the diameter of the network. The basic structure of the tree decomposition is the same, however, with the largest width bag in the center of the tree. This is due to the presence of a dense core in all of these networks, but in the sparsest networks the ‘dense’ core is only slightly more dense than the longer whiskers coming off of it.

For the random PL graphs, there is a negative correlation between the γ exponent and the width of the tree, *i.e.*, the lower γ values yield larger widths. An apparent exception to this is PL(2.25), but this is because the degree heterogeneity of this graph is large enough that the giant component, *i.e.*, the graph we are actually considering, is much smaller. Similar to the ER networks, the trees with large widths also have a larger branch ratio. In contrast to the ER networks, the bag lengths in these trees are small relative to the diameter of the network. The length grows with the power law exponent, suggesting that high width bags become less tightly connected when there is less degree heterogeneity. Again, this phenomena can be explained by a relatively dense core which must be placed in a single bag with tree-like whiskers coming off and forming the branches of the trees. This core-periphery effect is more pronounced in the networks with more degree-heterogeneity, whereas the networks with less degree-heterogeneity begin to look more like the sparse ER networks.

Looking at the real networks, several have similar properties; for example, AS20000102 and GNUTELLA09 both have tree decompositions with a much higher branch ratio compared to any other network. In fact, the number of bags and the diameter of the tree are similar in both of these networks, indicating similar cut structures in the decompositions of the two networks. They also have similar lengths. The widths of the two networks, are very different; GNUTELLA09 has a width of 1628, while AS20000102 has a width of 87. The tree-decompositions show the same structure, a set of central bags with many small width branches coming off of them. In these networks, the large central bags have a great number of leaf bags attached to them, most having a width of 2 or 3, indicating a core-periphery structure which has many single

Network	N_a	N_n	B_a	B_n	E_a	E_n	L_a	L_n	W_a	W_n	\bar{W}_a	\bar{W}_n
ER Networks												
ER(1.6)	3127	2881	.384	.518	43	37	19	21	78	85	1.94	2.55
ER(1.8)	3457	3229	.398	.532	37	39	17	18	156	155	2.52	3.58
ER(2)	3760	3554	.413	.533	37	52	16	17	234	229	3.03	5.11
ER(4)	3777	3683	.489	.545	34	103	9	10	1092	1106	13.0	37.2
ER(8)	2787	2646	.517	.540	28	162	6	6	2207	2244	35.4	148
ER(16)	1856	1734	.539	.539	27	188	5	5	3141	3186	73.2	361
ER(32)	1137	1069	.540	.529	21	163	4	4	3862	3902	137	626
PL Networks												
PL(2.25)	3192	3042	.591	.658	26	30	6	7	199	206	3.47	4.60
PL(2.50)	4672	4377	.537	.629	31	35	7	9	218	255	2.99	3.96
PL(2.75)	4500	4170	.498	.604	38	36	9	10	147	159	2.40	3.05
PL(3.00)	3974	3637	.466	.584	35	32	12	14	95	105	1.89	2.35
Planar Networks												
PLANARGRID	1884	1917	.609	.537	60	33	69	33	70	60	7.11	8.10
POWERGRID	4666	4420	.442	.555	58	32	26	22	20	22	2.21	2.45
Real Networks												
AS20000102	6364	6307	.833	.854	32	39	4	5	87	90	2.592	2.76
CA-GRQC	3014	2829	.566	.627	38	67	7	8	221	225	5.08	7.28
LEHIGHFB	1919	-	.465	-	150	-	4	-	2982	-	194	-
GNUTELLA09	6475	6276	.824	.874	32	167	6	6	1628	1764	13.4	47.3

TABLE III: Network TD Statistics: X_a =AMD, X_n =ND.

or double node whiskers coming off of the core. This large number of leaf nodes creates the high branch ratio.

GRQC and LEHIGHFB both show similarities to some of the synthetic networks. GRQC has a small set of large central bags and relatively thin outer branches. It has a smaller width (and average width) compared to the ER networks of a similar density. It is more similar to the PL networks, both qualitatively and quantitatively than the ER networks. LEHIGHFB has the largest width of the real networks and the second largest width of all the networks with 50% of the network contained in the central bag. Noting that it has the highest average degree of all the networks, this is not surprising, but the fact that it does not have the largest width demonstrates that there are multiple factors which affect the width of a tree decomposition besides just the average degree. The large average degree of LEHIGHFB also slows down the computations, hence we were only able to compute a tree decomposition using the faster AMD heuristic. LEHIGHFB is also distinctive from the other AMD tree-decompositions because of the large eccentricity of the tree decomposition. This is due to a long, wide ‘trunk’ of bags which have narrow branches at each end.

Again, the two planar networks (PLANARGRID and POWERGRID) are different from the other networks, showing a large length relative to the diameter and width of the bag. In fact, we can see that the tree decomposition of PLANARGRID has a length of similar value to the width (which we also saw with moderately dense ER), which occurs as the bags consist of paths (with a few additional nodes) across the grid.

The tree decompositions generated using the ND heuristic show several of the same general trends, though there are a few differences. The widths and lengths of ND are very similar to the AMD widths/lengths in the networks and demonstrate that the basic cut-structure of each network is captured by both heuristics. The main feature difference between ND and AMD seems to be that ND generates long thick trunks in the networks with dense cores.

B. Interpretation of Empirical Results for TDs

In many of the networks, such as GNUTELLA09, the eccentricity of the tree decomposition and the size of the bags indicates that there is a large overlap between bags. By looking at the statistics of nodes in a bag, we can get a preliminary idea of what the interior of a bag looks like. In Figure 4, we provide results for a representative set of real networks, and we provide a least-squares fit in the figure for reference. We see that the nodes that appear in many bags have a large degree, as expected. The POWERGRID network shows the weakest correlation while AS200001202 shows the strongest. Similar results are obtained for ND. In AS20000102, the strong linear correlation occurs because most of the nodes (not bags) in the network have many ‘leaves’ or neighbors of degree 1. Checking the bags, each of these leaf nodes, depending on the heuristic used, tends to generate a bag containing the leaf node and the single high-degree node it is connected to. This causes each node to appear in a number of bags that is approximately the same as its degree. The larger bags then have many leaf bags of small width hanging off of them and each of these leaf bags contains a high degree node and a low degree node.

In the other networks, this does not happen. In these networks (such as CA-GRQC and LEHIGHFB), there is a set of thick central bags with only a few of these (in several networks, only the largest bag) having any leaf bags connected to them. These central bags form a trunk which eventually splinters at several points into many smaller branches. These central bags tend to contain many of the high degree nodes in the network. Performing the same calculations, including the median values of the degree and k -core depth, shows that the high k -core nodes tend to be in many bags in the tree decompositions as well.

These results for the real graphs are consistent with a tree-like structure that includes the presence of a central, relatively well-connected core, as will be discussed in more detail in the next section. As this core is difficult to separate without breaking many edges, these nodes are repeated in the many large central bags. They are also repeated in low-width bags that contain more peripheral nodes, consistent with the

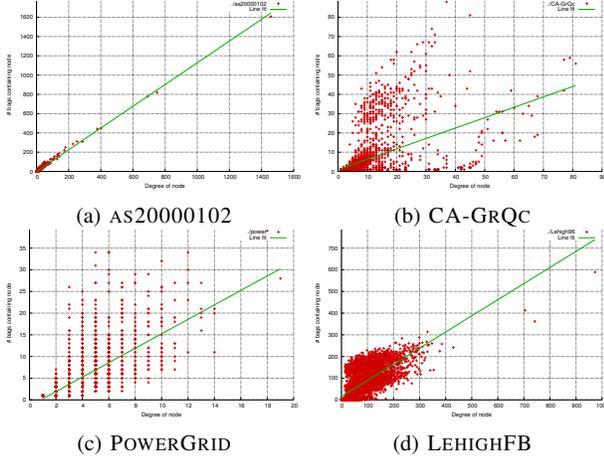


Fig. 4: Bags containing node vs the degree of node in the AMD heuristic.

existence of sparser periphery which consists of both whiskers (such as the low-width leaves in the decomposition) or sparser, peripheral layers in the network.

VI. CORE-PERIPHERY: k -CORE

In this section we look at the k -core decomposition as an easy-to-compute alternative to the previous (more rigorous but more expensive) notions of tree-likeness. We investigate some of the connections between the k -core structure and the tree-like properties of real networks. As a comment on our nomenclature, we will often talk about the “depth” of a node in this section. By depth we will mean a node’s core number (i.e., the shell S_k to which the node belongs). A deep node has a high core number, while a shallow node has a small core number, relative to the range of cores present in the network. This last point is important as there is a large amount of variation among network k -core structure and number of shells in a given network. Also note that we define k_{min} to be the maximum k such that the k -core contains the entire network.

A. Summary of k -core Statistics

Table IV contains the size and depth of the k -core structure, as well as relevant graph statistics such as network size and average degree, for our suite of networks.

Among the synthetic networks, there is a marked difference between expander graphs and low-dimensional graphs versus more realistic graphs; and, in particular, between the ER networks and networks with more degree heterogeneity. With both network types, the depth of the k -core structure increases with the average degree (increased connectivity of the network), as one would expect. With the ER networks, the connectivity is distributed uniformly among all nodes. This uniformity causes k_{min} to increase with the average degree, keeping a smaller difference between the maximum and minimum shells. In contrast, the nodes in the power-law networks have a wide range of degrees with a few nodes having degrees orders of magnitude larger than the others. As a consequence, the k -core decomposition deepens, keeping a non-empty 1-shell while the maximum k -core increases. Also since most nodes have

a small degree, the shallow shells have a larger proportion of the nodes than the deeper shells.

In PLANARGRID, there is a single shell, S_2 , while the random geometric networks have a richer k -core structure, but the core is less prominent than in the real networks. This is to be expected since the ‘core’ is the results of random fluctuations and boundary effects in the generation process.

Similarly, the average degree and the maximum k -core are correlated in the real networks. Broadly speaking, the real networks we have been studying have degree heterogeneity similar to the synthetic PL networks and thus have a similar but richer k -core structure to these synthetic networks, with a few exceptions. The collaboration networks behave in this manner in the shallow shells, but have large gaps in the k -core decomposition with isolated deep cores consisting of a few small, disconnected subgraphs. These cliques are likely due to a few large collaborations and publications associated with a single experiment, and they are not central to the network structure [32]. The real networks with the highest average degree (the Facebook networks, email networks, and the political blogs network) have the deepest k -cores and a larger percentage of their nodes in the deep cores.

These k -core statistics also agree with several of the empirical results in [30]. Cuts of good conductance do not exist in many real networks because the shallow nodes have many connections to the highly connected central core and a good cut cannot pass through this core. Thus, the “whiskers,” or small relatively-dense subgraphs attached to the main core by only a few edges, are the best cuts. Facebook networks, because of their deep, large central cores and the tendency of peripheral nodes to connect to these cores, have very poor conductance cuts, lacking even the whiskers described above. This behavior is different than the classical models (such as dense ER graphs) where the nodes are essentially all in the same core of the network (a different reason for lacking good cuts). In contrast, the planar geometric networks have good cuts at all size scales and also do not have a deep core structure or many connections to the deep cores.

B. k -core and Hyperbolicity

We performed several experiments to confirm that the k -core is an important part of the hyperbolic structure of the network. If the intuition of geodesic paths connecting to a more complicated, higher δ structure in the core of the network is true, then removing (or collapsing) these nodes should have several predictable effects. In particular, if deep k -cores are central to the structure of the network, removing the cores from the network should result in a “hole” in the center of the network. The geodesics would now have to travel around this hole, resulting in large cyclical structures which would increase the maximum δ observed in the networks. This would also cause the diameter of the network to increase and the scaled δ_d to fall off later. If, on the other hand, we collapsed the k -cores (i.e., remove the k -core and then assign all edges from nodes outside the removed k -core into nodes in the removed k -core to a single supernode), we should see the hyperbolicity decrease while the distances between the nodes should stay approximately the same as it will modify only a short portion of the geodesics.

Network	n_c	\bar{d}	k_{min}	k_{max}	P_{kmin}	P_{kmax}
ER Networks						
ER(1.6)	3210	2.16	1	2	58.7	41.3
ER(1.8)	3617	2.28	1	2	49.6	50.4
ER(2)	4001	2.39	1	2	41.5	58.5
ER(4)	4879	4.05	1	3	8.28	67.6
ER(8)	4998	8.04	1	5	.300	88.3
ER(16)	5000	16.1	4	11	0.04	88.3
ER(32)	5000	32.1	7	23	0.02	93.8
PL Networks						
PL(2.25)	3393	3.32	1	5	45.7	0.825
PL(2.50)	4895	2.78	1	4	52.2	0.776
PL(2.75)	4650	2.43	1	2	59.0	41.0
PL(3.00)	4071	2.24	1	2	65.9	34.1
Planar Networks						
PLANARGRID	2500	3.92	2	2	100	100
POWERGRID	4941	2.67	1	5	32.1	0.243
Real Networks						
AS20000102	6474	3.88	1	12	37.9	0.324
OREGON1	11174	4.19	1	17	35.3	0.269
LEHIGHFB	5073	78.2	1	62	1.42	15.4
VANDERBILTFB	8063	106.1	1	86	1.98	23.3
STANFORDFB	11586	98.1	1	91	4.36	20.1
CA-GRQC	4158	6.46	1	43	17.9	1.06
CA-ASTROPH	17903	22.0	1	56	5.55	0.318
EMAILENRON	33696	10.7	1	43	28.4	0.816
EMAILEUALL	224832	3.02	1	37	83.9	0.130
POLBLOGS	1222	27.4	1	36	11.3	4.50
STANFORDWEB	255265	15.2	1	71	5.98	0.152

TABLE IV: k -core network statistics: P_{kmin} and P_{kmax} are percentage of nodes in the k_{min} and k_{max} shells, respectively.

We performed all of these experiments, but because of space, we will show only the results involving the maximum hyperbolicity calculations on the k -periphery of several real networks. In Figure 5, we remove cores and plot the resulting maximum hyperbolicity of the giant component of the k -periphery. As predicted, removing cores from the network increases the size-scale at which cyclical structures are still found. The minimum scale at which cycles appear stays the same, indicating that small scale cycles must appear across the network, not only in the deep cores. This provides additional corroboration for the notion of a size-scale above which the network structure is hyperbolic or tree-like, but below which it is not. Additionally, the increasing size of the cyclical structures in the network as cores are removed provides further evidence that the quadrupled centers are contained in the deep k -cores of the network.

C. k -core and Tree Decompositions

In this section, we will demonstrate that the k -core decomposition also interacts with the cut-structure of the networks we studied. To start, we can ask: how is the k -core structure in real networks related to the cut-structure (via the tree-decompositions) in real networks? Figure 6 shows two AMD tree decompositions of AS20000102. The diameters of the bags are scaled by width and colored by the average k -core of the nodes they contain. The smallest peripheral bags are difficult to see, but they have a green or blue coloring (low k -core), while the center bags have a yellow to red coloring (high k -core).

Looking at the average k -core of the bags in the real networks, results not shown demonstrate that a higher k -core is correlated with the central, wide bags of the tree decomposition. But we also see (as Figure 6 shows) that this is a noisy correlation. The central bags are yellow, while bags

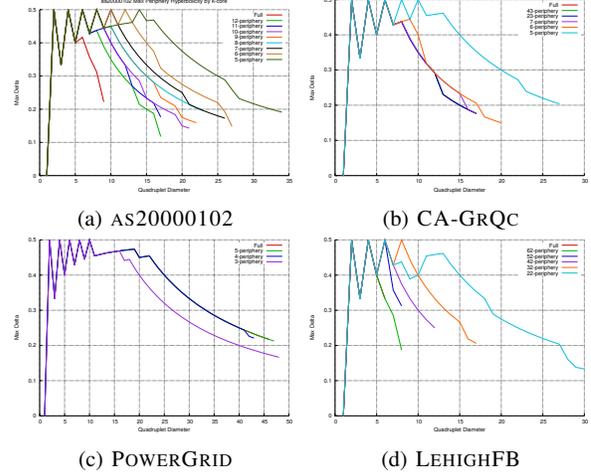


Fig. 5: Maximum Hyperbolicity of Network Periphery.

further towards the periphery are red. There are also hundreds of small bags containing only two or three nodes connected to the central bag. As each of these contains a single node in a deep core and the rest of the bag is a shallow core, these nodes have a high average k -core. In the planar networks, most of the bags contain low k -core nodes, while a few peripheral bags contain the deep core nodes, suggesting that the k -core is not as central to the network structure.

We also removed the k -cores, as in the previous section. This had the effect of reducing the width in the real networks we tested on, with the exception of POWERGRID, where the widths remained flat until the network became disconnected. In CA-GRQC, we see no effect from the deepest cores (the large clique-like physics collaborations), but we see a similar fall-off in width when we begin removing shallower cores. Interestingly, while removing the cores decreases the widths of many of the real networks, it leaves similar tree structures (see Figure 6b). The bags have a smaller width and a lighter coloring, but the tree structure appears to be similar. This would seem to support the idea of a nested core-periphery as reported in [30], but additional investigation is required.

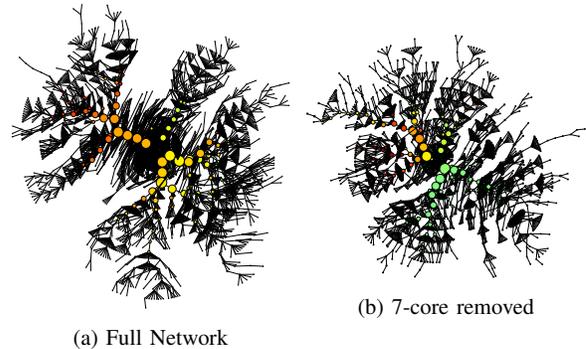


Fig. 6: AMD of AS20000102, bags colored by average core number of nodes

VII. DISCUSSION AND CONCLUSION

Our empirical results (only some of which have been presented here for space reasons) indicate that realistic social and information networks do have meaningful large-scale tree-like structure, both in terms of their cut properties and in terms of their metric properties. Although characterizing this tree-like structure with just the traditional tools of δ -hyperbolicity or tree decomposition theory is nontrivial and quite brittle, our empirical results also indicate that — for the class of realistic social and information networks — these two very different tree-like notions have strong similarities. In addition, these similarities can be interpreted in terms of a nested core-periphery structure, and they can be identified with the fast k -core decomposition heuristic. Clearly, the next step is to try to exploit our insights and this heuristic to do improved graph data mining, machine learning, and statistical inference on much larger realistic social and information networks.

Acknowledgments. The authors would like to acknowledge financial support from the Air Force Office of Scientific Research, the Army Research Office, the Defense Advanced Research Projects Agency, and the National Science Foundation. Part of this work was performed while B. D. S. was an employee at Oak Ridge National Laboratory. This manuscript has been authored by a contractor of the U.S. Government under Contract No. DE-AC05-00OR22725. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

REFERENCES

- [1] Supporting website. <http://snap.stanford.edu/data/index.html>.
- [2] L. A. Adamic and N. Glance. The political blogosphere and the 2004 U.S. election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*, pages 36–43. ACM, 2005.
- [3] A. B. Adcock, B. D. Sullivan, O. R. Hernandez, and M. W. Mahoney. Evaluating OpenMP tasking at scale for the computation of graph hyperbolicity. In *Proceedings of the 9th International Workshop on OpenMP (IWOMP 2013)*, pages 71–83, 2013.
- [4] J. Ignacio Alvarez-Hamelin, L. Dall’Asta, A. Barrat, and A. Vespignani. Large scale networks fingerprinting and visualization using the k -core decomposition. In *Proceedings of NIPS*, 2005.
- [5] J. Ignacio Alvarez-Hamelin, L. Dall’Asta, A. Barrat, and A. Vespignani. K -core decomposition of Internet graphs: hierarchies, self-similarity and measurement biases. *Networks and Heterogeneous Media*, 3(2):371–293, 2008.
- [6] P. R. Amestoy, T. A. Davis, and I. S. Duff. Algorithm 837: AMD, an approximate minimum degree ordering algorithm. *ACM Transactions on Mathematical Software*, 30(3):381–388, 2004.
- [7] E. Amir. Approximation algorithms for treewidth. *Algorithmica*, 56(4):448–479, 2010.
- [8] S. Arnborg and A. Proskurowski. Linear time algorithms for NP-hard problems restricted to partial k -trees. *Discrete Applied Mathematics*, 23(1):11 – 24, 1989.
- [9] V. Batagelj and A. Mrvar. Pajek - analysis and visualization of large networks. In *Graph Drawing Software*, pages 77–103. Springer, 2003.
- [10] V. Batagelj and M. Zaversnik. An $O(m)$ algorithm for cores decomposition of networks. *CoRR*, cs.DS/0310049, 2003.
- [11] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.
- [12] M.R. Bridson and A. Häflicher. *Metric Spaces of Non-Positive Curvature*. Grundlehren Der Mathematischen Wissenschaften. Springer, 2009.
- [13] W. Chen, W. Fang, G. Hu, and M. W. Mahoney. On the hyperbolicity of small-world and tree-like random graphs. *In press at: Internet Mathematics*. Also available at: [arXiv:1201.1717](https://arxiv.org/abs/1201.1717).
- [14] V. Chepoi, F. Dragan, B. Estellon, M. Habib, and Y. Vaxès. Diameters, centers, and approximating trees of δ -hyperbolic geodesic spaces and graphs. In *Proceedings of the 24th Annual Symposium on Computational Geometry*, pages 59–68, 2008.
- [15] F.R.K. Chung and L. Lu. *Complex Graphs and Networks*. American Mathematical Society, 2006.
- [16] P. Colomer-de Simon, A. Serrano, M. G. Beiro, J. Ignacio Alvarez-Hamelin, and M. Boguna. Deciphering the global organization of clustering in real complex networks. *CoRR*, abs/1306.0112, 2013.
- [17] R. Diestel. *Graph Theory*. Graduate Texts in Mathematics. Springer, 2006.
- [18] P. Erdos and Alfred R. On the evolution of random graphs. *Publ. Math. Inst. Hungary. Acad. Sci.*, 5:17–61, 1960.
- [19] J. A. George. Nested dissection of a regular finite element mesh. *SIAM Journal of Numerical Analysis*, 10:345–363, 1973.
- [20] J. R. Gilbert and R. E. Tarjan. The analysis of a nested dissection algorithm. *Numerische Mathematik*, 50(4):377–404, 1986.
- [21] M. Gromov. Hyperbolic groups. In *Essays in Group Theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer-Verlag, 1987.
- [22] I.V. Hicks, A. Koster, and E. Kolotoglu. Branch and tree decomposition techniques for discrete optimization. *INFORMS TutORials in Operation Research*, 2005:1–29, 2005.
- [23] E. A. Jonckheere, P. Lohsoonthorn, and F. Ariaei. Scaled Gromov four-point condition for network graph curvature computation. *Internet Mathematics*, 7(3):137–177, 2011.
- [24] E. A. Jonckheere, P. Lohsoonthorn, and F. Bonahon. Scaled Gromov hyperbolic graphs. *Journal of Graph Theory*, 57(2):157–180, 2008.
- [25] G. Karypis and V. Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput.*, 20(1):359–392, 1998.
- [26] M. Kitsak, L. K. Gallos, S. Havlin, F. Liljeros, L. Muchnik, H. E. Stanley, and H. A. Makse. Identification of influential spreaders in complex networks. *Nature Physics*, 6:888–893, 2010.
- [27] R. Kleinberg. Geographic routing using hyperbolic space. In *Proceedings of the 26th IEEE International Conference on Computer Communications*, pages 1902–1909, 2007.
- [28] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.
- [29] D. V. Krioukov, F. Papadopoulos, M. Kitsak, A. Vahdat, and M. Boguñá. Hyperbolic geometry of complex networks. *CoRR*, abs/1006.5169, 2010.
- [30] J. Leskovec, K.J. Lang, A. Dasgupta, and M.W. Mahoney. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics*, 6(1):29–123, 2009. Also available at: [arXiv:0810.1355](https://arxiv.org/abs/0810.1355).
- [31] H. M. Markowitz. The elimination form of the inverse and its application to linear programming. *Management Science*, 3(3):255–269, 1957.
- [32] T. Martin, B. Ball, B. Karrer, and M. E. J. Newman. Coauthorship and citation in scientific publishing. *CoRR*, abs/1304.0473, 2013.
- [33] F. Montgolfier, M. Soto, and L. Viennot. Treewidth and hyperbolicity of the internet. In *Proceedings of the 10th IEEE International Symposium on Network Computing and Applications (NCA)*, pages 25–32, 2011.
- [34] O. Narayan and I. Sanjeev. Large-scale curvature of networks. *Phys. Rev. E*, 84:066108, Dec 2011.
- [35] N. Robertson and P.D Seymour. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309 – 322, 1986.
- [36] H. Röhrig. Tree decomposition: a feasibility study. Master’s thesis, Universität des Saarlandes, Saarbrücken, Germany, 1998.
- [37] P.D. Seymour and R. Thomas. Call routing and the ratcatcher. *Combinatorica*, 14(2):217–241, 1994.
- [38] Y. Shavit and T. Tanel. Hyperbolic embedding of internet graph for distance estimation and overlay construction. *IEEE/ACM Trans. Netw.*, 16(1):25–36, 2008.
- [39] A. L. Traud, P. J. Mucha, and M. A. Porter. Social structure of Facebook networks. *CoRR*, abs/1102.2166, 2011.
- [40] D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393:440–442, 1998.