September 2011 No. 4

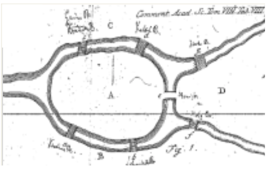

THE EULER NEWSLETTER

In this issue

Internet curvature and possible implications on routing system design

In 2004, Shavitt and Tankel [1] identified a new characteristic of the Internet graph, its curvature so as to better represent the Internet distance map in a geometric space. They show empirically that the Internet topology (namely, the AS graph) embeds with better accuracy (smaller average distortion) into a low-dimensional hyperbolic space than into a Euclidean space with comparable dimension. In their seminal work, C. Papadimitriou and D. Ratajczak [2] introduced the notion of greedy embedding of the graph $G = (V, E)$ defined by the physical network topology in the Euclidean plane.

A greedy embedding of an undirected graph G in a metric space (X, d) is a mapping $f : V(G) \rightarrow X$ with the following property: for every pair of distinct vertices $s, t \in V(G)$ there exists a vertex u adjacent to s such that $d(f(u), f(t)) < d(f(s), f(t))$ (distance preserving). Following this definition, a greedy embedding of a graph G in a metric space (X, d) is a mapping of G in (X, d) such that a distance decreasing path exists between every pair of vertices in G. A distance decreasing path from s to t in a greedy embedding of G is a path $(s = v_1, v_2, \ldots, t = v_k)$ such that $d(v_i, t) > d(v_i+1, t)$, for $i = 1, 2, \ldots, k-1$. Greedy routing is thus a natural abstraction of geometric routing in which nodes are assigned virtual coordinates in a metric space, and these coordinates are used as addresses to perform point-to-point routing in this space. The main challenge here consists in finding the appropriate mapping function f together with a polynomial-time algorithm for embedding $V(G)$ in the space X (and maintaining it) so as to allow for greedy routing using the metric d associated to that space (without degradation). Introduced by R. Kleinberg in 2005 [3], greedy embedding in hyperbolic metric space was a crucial step in search of means to overcome known limitations of geometric routing on unidirected graph G embedded in the Euclidean space.

Work on Internet topology embedding to hyperbolic spaces was further progressed by R. Kleinberg who demonstrated in 2007 [4] that every connected finite graph has a greedy embedding in the hyperbolic plane (though reconstructing such embedding upon topology modification still requires $O(n)$ operations per modification). A hyperbolic n-space, denoted H^n, is the maximally symmetric, simply connected, n-dimensional Riemannian manifold with constant sectional curvature -1. A classical example of hyperbolic geometry is the Poincaré disk model $D = \{(x, y) \in \mathbb{R}^2 \,|\, x^2 + y^2 < 1\}$ which represents the hyperbolic plane H^2 as the interior surface of an open unit disk whose points lie inside the unit circle $\partial D = \{(x, y) \in \mathbb{R}^2 \,|\, x^2 + y^2 = 1\}$ (the circle itself representing the boundary of the disk at infinity), and the geodesics are the arcs of circles orthogonal to the boundary of the unit disk ∂D. Sectional curvature describes the curvature of Riemannian manifolds and controls the rate of geodesic deviation. A metric space (X, d) is called geodesic if any two points $x, y \in X$ in can be joined by a path (a geodesic $[x, y]$) whose length coincides with the distance $d(x, y)$ between these two points.

In Fig. (a), L1,L2;L3 and P1;P2;P3 are examples of geodesics. L1;L2;L3 intersect to form triangle ABC. The sum of its angles a + b + c < π (as opposed to Euclidean geometry where the sum of the angles of a triangle equals to π). Also, there are infinitely many geodesic lines (examples are P1;P2;P3) that are parallel to line L1 and go through a point C that does not belong to L1.

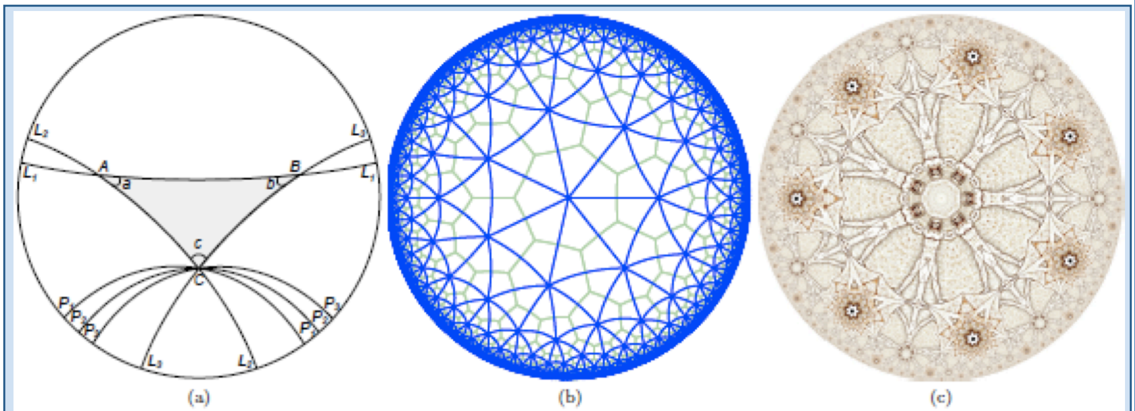

Poincaré disk model

To a certain extend hyperbolicity measures the deviation from tree-likeness. Several authors have thus tried to determine the hyperbolicity of the observable Internet topology with the expectation to exploit routing schemes known to provide a competitive memory-stretch tradeoff if the graph determined by the topology shows bounded hyperbolicity. De Montgolfier et al. [5] check empirically the Gromov's four-points condition [7] on the AS graph and on various router level graphs. A metric space (X, d) is δ-hyperbolic when it verifies the four-point condition, i.e., when the Gromov product $(x \cdot y)_o = [d(o, x) + d(o, y) - d(x, y)]/2$ of any two points $x, y \in X$ with respect to a fixed reference point $o \in X$ satisfies the inequality $(x \cdot y)_o \geq \min\{(x \cdot z)_o, (y \cdot z)_o\} - \delta$, for some $\delta \geq o$, and any four points o, x, y, z in X. Following the four-point condition, the hyperbolicity of a simply connected unweighted graph $G = (V, E)$ equipped with metric $d_G(x, y)$ is the minimum value of $\delta \geq 0$ such that the metric space $D_G = (V, d_G)$ is δ-hyperbolic. Note also that the case $\delta = o$ coincides with the family of metric trees (if o is the root of the tree, then $(x \cdot y)_o$ corresponds to the distance from o to the least common ancestor of x and y): trees are o-hyperbolic. V. Ramasubramanian et al. [6] show empirically that Internet delays follow a relaxed version of four-point condition. Following this finding, they propose to actually build a tree representation where the nodes can be embedded (i.e., mapped onto leaves of a virtual tree, whose inner nodes are virtual points and whose edge weights are carefully selected so as to represent original graph measurements). This allows to estimate latencies and to solve natural tasks such as selection of servers with short latency.

Subscribe to electronic EULER newsletter: https://sympa.inria.fr/sympa/info/euler-news

As an interesting algorithmic consequence, a δ-hyperbolic graph $G = (V, E)$ of n vertices can be approximated in $O(n^2)$ time by a tree with additive error of $2\delta(\log n)$ [7]. C.Gavoille and O.Ly [8] use this property to derive a distance labeling scheme (DLS) which assigns to each vertex of any bounded hyperbolicity graph G a label of $O(\log^2 n)$ bits such that distances can be approximated with an additive error of $O(\delta \log n)$. Given a finite set S of points of *a δ-hyperbolic* metric space, Chepoi et al. [9] propose linear-time algorithms (in the number of links) for computing the diameter of S with an additive error 2δ and an approximate radius and center of a smallest enclosing ball for S with an additive error 3δ. They also provide a simple linear-time construction of distance approximating trees of δ-*hyperbolic* graphs G on n vertices with an additive error $O(\delta \log n)$. In a following paper [10], the authors show that every unweighted δ-*hyperbolic* graphs has an additive $O(\delta \log n)$-spanner with $O(\delta n)$ edges. Moreover, the authors propose an $O(\delta \log n)$-additive error routing labelling scheme with $O(\delta \log^2 n)$-bits labels.

However, according to the best available measurement data [11], as the Internet topology grows, the number links connecting nodes of similar degrees (tangential links) increases faster than the number of links connecting low-degree to high-degree nodes (radial links), i.e., its deviation from tree-likeness increases. Topology measurements fail to detect some tangential links interconnecting medium-to-low degree nodes since many of these links belong to none of the spanning trees rooted at the vantage points in the core. This excess of tangential links is responsible for a slightly higher assortativity of the AS-graph and the increase of the number of BGP routing paths even if the network size itself doesn't increase. This observation suggests that capturing hyperbolicity upper bounds becomes critical as tree approximation errors are proportional to δ.

Besides verifying the four-point condition, equivalent definitions of δ-hyperbolicity involving different but comparable values of δ can be considered. The space X is δ-hyperbolic if all the geodesic triangles in X are δ-slim or δ-thin for some fixed $\delta \geq 0$. Verifying the δ-slim triangles condition means that there exists a constant $\delta \geq 0$ such that for any geodesic triangle in X (a geodesic triangle with vertices $x, y, z \in X$ is the union $[x, y] \cup [y, z] \cup [z, x]$ of three geodesic segments connecting these vertices) and for any point $w \in [x, y]$, the distance $d(w, [y, z] \cup [z, x]) \leq \delta$ (and similarly for the other sides $[y, z]$, $[z, x]$ of the triangle). Verifying the δ-thin triangles condition means that there exists a constant $\delta \geq 0$ such that for any geodesic triangle in X and any pair of points v, w lying on its sides and equidistant from one of its vertices (e.g., $v \in [x, z]$, $w \in [x, y]$ and $d(x, v) = d(x, w)$), the distance $d(v, w) \leq \delta$. Despite its simplicity proving such condition holds globally results in a major computational challenge, in particular, if we assume that topological information is not globally available.

The seminal work of M.A.Serrano et al. [12] opens new perspectives in approaching geometric routing. Indeed, these authors found in 2007 that self-similarity of clustering in real complex networks (such as the Internet) provides strong empirical evidence that hidden metric spaces underlie the observable topology of these networks. They provided further evidence that this metric space explanation is plausible by introducing a class of network models and by finding that networks generated by elements of this class reproduce all the self-similar effects that can be empirically observed in real networks. Subsequently, these authors found that these hidden metric spaces are hyperbolic. Consequently, research has shifted towards finding construction rules of hyperbolic spaces that best reproduce the hidden properties of observable topologies (hence, the resulting scheme doesn't require maintenance of a global network-wide structure to perform the embedding) so as to allow for greedy routing using the standard metric for hyperbolic space.

References

https://www-sop.inria.fr/mascotte/EULER/wiki/pmwiki.php/Main/News

How the EULER project is progressing?

The EULER project, started on October 1st, 2010, is completing its first year of the 3-year project duration.

EULER organized the second edition of the **TERANET** (Toward Evolutive Routing Algorithms for scale-free/internet-like NETworks) International Workshop. Co-located with the 25th International Symposium on DIStributed Computing (DISC), the workshop took place on September 19, 2011 in Rome, Italy. This yearly event focuses on current research dedicated to new routing paradigms, models, and algorithms for distributed and dynamic routing systems applicable to the Internet and its continuous evolution. The agenda of the workshop as well as the extended abstract and the presentations of the six invived speakers are available at the following URL (EULER wiki): **http://www-sop.inria.fr/mascotte/EULER/wiki/pmwiki.php/Events/TERANET2011.** The next preliminary project meeting will take place on November 7-9, 2011 at UPC, Barcelona, Spain.

Regarding the WPs activities, task T2.1 of WP2 is steadily progressing the specification of the routing system architecture. The work comprises two parts: Part 1 dedicated to the *Routing system architecture* itself and Part 2 dedicated to *Routing system models.* In Part 1, research work is dedicated to the definition of a hierarchical decomposition of the routing functional area by means of Enhanced Functional Flow Block Diagram. Particularly, we have identified and defined the sub-functions (such as discovery or determination of routing path) composing the routing function, and the spatial and temporal distribution existing among them. The work is also dedicated to the specification of the information (extended entity-relationship) model specifying the properties and the organisation of the routing related information, including their relation and their interaction. The second part of the work is dedicated to the specification of the routing procedures and the data structures by means of class diagrams and sequence diagrams (representating message exchanges between interacting entities).

Concerning WP3, major efforts have been concentrated in the edition of the deliverable D3.1 and deliverable D3.2. D3.1 collects, describes and compares the state-of-the-art graph models and the desired properties known to be useful to either compact or greedy routing schemes on hyperbolic metric spaces. D3.1 also includes experimental work concerning graph generation and property tests. D3.2 focuses on the topology analysis and modelling based on measurements. It presents a set of measurable quantities of Internet, the tools able to collect such quantities and the datasets available so far for the project.

In the context of WP4, task T4.2 continues the activity on the documentation of experimental methodology, scenarios, and tools. A (non exhaustive) list of existing tools and tools developed in the course of the EULER project has been defined. The term tool includes topology/graph generators, topology/graph properties testers, routing model simulators, and routing protocol emulators.

For more information: **http://www.euler-fire-project.eu.**

Call for Papers

IEEE Communications Magazine **01/11/2011**
Information-Centric Networking

44th ACM Symposium on Theory of Computing **02/11/2011**
http://cs.nyu.edu/~stoc2012/
May 19-22, 2012, New York, USA

12th joint ACM Sigmetrics/Performance Conf. **04/11/2011**
http://www.sigmetrics.org/sigmetrics2012/
June 11-15, 2012, London, UK

32nd Int. Conf. on Distributed Computing Systems **08/11/2011**
http://icdcs-2012.org/
June 18-21, 2012, Macau, China

Computer Communications **01/12/2011**
Reliable Network-based Services

11th IFIP Networking 2012 **10/12/2011**
http://networking2012.cvut.cz/
May 21-25, 2012, Prague, Czech Republic